

Symantec Backup Exec 2010 R2

Administrator's Guide



The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Documentation version: 2010

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our Web site at the following URL:

www.symantec.com/business/support/

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America supportsolutions@symantec.com

Contents

Technical Support	4	
Chapter 1	Introducing Backup Exec	63
	About Backup Exec	63
	How Backup Exec works	68
	What's new in Backup Exec	70
	What's new in Backup Exec agents and options	75
	Backup Exec agents and options	80
	About Backup Exec media server components	80
	About Backup Exec server protection agents	81
	About Backup Exec application protection agents	82
	About Backup Exec's virtual machine agents	85
	About Backup Exec client protection agents	86
	About Backup Exec media server storage options	87
	About the Administration Console	91
	About the Home view	95
	Configuring the Home view	95
	Restoring the Home view's default configuration	95
	Editing items on the Home view	96
	Help and Technical Support items	96
	Summary items	98
	Detail items	99
Chapter 2	Installing Backup Exec	101
	About installing Backup Exec	102
	Before you install	103
	About the Environment Check	104
	Checking your environment before installing	105
	About the Backup Exec service account	106
	Changing service account information	107
	About changing Windows security	108
	Changing Windows security to back up servers (only) in one domain	109
	Changing Windows security to back up servers and selected workstations in one domain	109

Changing Windows security to back up servers in more than one domain	110
Changing Windows security to back up servers and workstations in more than one domain	111
About Microsoft SQL Server 2005 Express Edition components installed with Backup Exec	111
About Backup Exec's standard features	112
System requirements	114
About the differences between a typical installation and a custom installation	116
Installing a typical installation of Backup Exec	116
Installing a custom installation of Backup Exec	119
Installing additional Backup Exec options to the local media server	123
Special considerations for installing Backup Exec to remote computers	125
Push-installing Backup Exec to remote computers	126
About installing Backup Exec options to remote computers	131
Push-installing the Remote Agent and Advanced Open File Option to remote computers	134
Push-installing the Desktop Agent and DLO Maintenance Service from the media server to remote computers	137
About installing the Remote Agent for Windows Systems	139
Installing the Remote Agent and the Advanced Open File Option to a remote computer in the backup selections list	140
How to install the Remote Agent and Advanced Open File Option in an Active Directory network	140
Using a command prompt to install the Remote Agent on a remote computer	145
Using a command prompt to uninstall the Remote Agent from a remote computer	147
Using a command script to install the Remote Agent and AOFO	148
Using a command script to uninstall the Remote Agent and AOFO	149
Installing the Remote Administrator	150
Running the Remote Administrator	151
Installing Backup Exec using the command line (silent mode)	153
Command line switches for silent mode installation of Backup Exec	154
Installing the Remote Administrator using the command line	162
Uninstalling Backup Exec using the command line	164

Creating installation parameter files	164
Using installation parameter files	165
Installing a trial version of Backup Exec agents and options	166
About the installation log	166
Repairing Backup Exec	167
Starting and stopping Backup Exec services	168
Backup Exec Services Manager options	168
Uninstalling Backup Exec	168
Uninstalling Backup Exec options from the local media server	169
About updating Backup Exec with LiveUpdate	170
About scheduling automatic updates using LiveUpdate	171
Scheduling automatic updates using LiveUpdate	172
Running LiveUpdate manually	173
Viewing installed updates	173
Viewing license information	174
License information options	174
Adding licenses	175
Finding installed licenses in your environment	176
About Backup Exec maintenance contract information	177
Running the Backup Exec Renewal Assistant	178
Viewing Backup Exec maintenance contract information	179
About upgrading from previous versions of Backup Exec	180
Post-installation tasks	182
Chapter 3	Configuring Backup Exec settings and options
	183
About configuring Backup Exec	184
About configuring logon accounts	184
About the default Backup Exec logon account	185
About Backup Exec restricted logon accounts	186
Creating a Backup Exec logon account	187
About the Backup Exec System Logon Account	189
Editing a Backup Exec logon account	190
Changing a Backup Exec logon account password	191
Replacing a Backup Exec logon account	192
Deleting a Backup Exec logon account	192
Changing your default Backup Exec logon account	193
Creating a new Backup Exec System Logon Account	193
About Backup Exec defaults	194
About job priority	195
Changing the default device and media set for jobs	196
Changing application settings for Backup Exec preferences	196
Application settings for Backup Exec preferences	197

Copying configuration settings to another media server	199
Adding multiple destination media servers by importing a list	200
Adding a destination media server in a non-CASO environment	200
Adding a destination media server in a CASO environment	201
Copy Settings options	203
Copying logon account information	204
Copy Logon Account options	204
About audit logs	205
Configuring the audit log	206
Viewing the audit log	206
Removing entries from the audit log	208
Saving the audit log to a file	208
About database maintenance	209
Configuring database maintenance	209
Viewing the location of Backup Exec databases	211
Advanced properties for a media server	212
About the Symantec RSS Reader	213
Viewing an article in the Symantec RSS Reader	213
Customizing the Symantec RSS Reader	214
Removing a default RSS feed from the Symantec RSS reader	215
Hiding columns	215
Showing a hidden column	215
Rearranging columns	216
Sorting column information	216
Viewing properties	217
Chapter 4	
Managing media	219
About media in Backup Exec	219
About media overwrite protection	223
About the default media set	226
About creating media sets	226
Deleting a media set	228
Renaming a media set	228
Associating media with a media set	229
Editing general properties for media sets	230
Media overwrite protection levels	232
About overwriting allocated or imported media	232
How Backup Exec searches for overwritable media	233
Selecting settings for media management	237

Settings for media management	237
Viewing audit log entries for media operations	241
Configuring specific media operations to appear in the audit log	241
Media labeling	242
Renaming a media label	243
Imported media labeling	244
Bar code labeling	244
Bar code rules in mixed media libraries	245
Creating bar code rules in mixed media libraries	245
Editing a bar code rule	245
Deleting a bar code rule	246
Application settings for bar code rules	246
Add Bar Code Rule options	247
About WORM media	247
Creating a new catalog	248
Device options for catalog jobs	249
Creating a restore job while reviewing media or devices	250
Media locations and vaults	250
Creating media vaults	251
Media Vault Properties	251
Configuring vault rules for media sets	252
Properties for vault rules for media sets	252
Deleting a media vault	253
Renaming a media vault	253
Finding media in a location or vault	254
About moving media to a vault or to the offline media location	254
Scanning bar code labels to move media	255
Scheduling a job to move media	255
Using the Vault Wizard to move media	256
Using the Move to vault task to move media	257
Move Media options	257
Move Media to Vault options	258
Drag and drop methods to move media	258
Using drag and drop methods to move media	259
About removing damaged media	259
About deleting media	260
Deleting media	260
General properties for media	261
Statistics properties for media	263
Media rotation strategies	265
Son media rotation strategy	265
Father/son media rotation strategy	266
Grandfather media rotation strategy	267

Chapter 5	Preparing for backup	269
	How to prepare for backup	270
	About backup strategies	270
	How to choose a backup strategy	270
	How to determine your backup schedule	271
	How to determine the amount of data to back up	271
	How to determine a schedule for data storage	272
	How to determine which devices to back up	272
	How to determine the number of resources to back up in a job	272
	About the archive bit and backup methods	273
	About backup methods	274
	About using the Windows NTFS Change Journal to determine changed files	280
	About selecting data to back up	280
	About using fully qualified computer names in backup selections	282
	About the Computer Name node in the backup selections list	282
	About the Favorite Resources node in the backup selections list	284
	Adding a Windows system to the Favorite Resources node in the backup selections list	285
	Deleting a Windows system from the Favorite Resources node in the backup selections list	286
	About the Domains node in the backup selections list	287
	Adding an Active Directory domain to the Active Directory Domains node	288
	Deleting an Active Directory domain from the Active Directory Domains node	289
	Manage Active Directory Domains options	289
	About the User-defined Selections node in the backup selections list	290
	Adding a user-defined selection to the User-defined Selections node	290
	Deleting a user-defined selection from the User-defined Selections node	292
	User-defined Selections options	292
	About managing Microsoft Virtual Hard Disk (VHD) files in Backup Exec	293
	How to back up user-defined Microsoft Windows Distributed File System data	294
	About selection lists	295
	Creating selection lists	296

Merging selection lists	300
Replacing selection lists	301
Copying selection lists	302
Holding jobs that back up a selection list	303
Deleting selection lists	303
Editing selection lists	304
Editing the Excludes selection list	305
About priority and availability windows for selection lists	306
Setting default priority and availability windows for all selection lists	307
Setting priority and availability windows for selection lists	307
Creating separate selection lists for each computer or resource	309
Creating a custom filter for backup selection lists	310
Filtering backup selection lists	313
Searching selection lists	314
Viewing the history for backup selection lists	314
Viewing a summary for a selection list	315
About resource discovery	316
Using resource discovery to search for new resources	316
About the Backup Exec Shadow Copy Components file system	320
How to restore individual items by using Granular Recovery Technology	321
Recommended devices for backups that use Granular Recovery Technology	324
About requirements for jobs that use Granular Recovery Technology	325
Chapter 6 Backing up data	329
How to back up data	329
Required user rights for backup jobs	331
Creating a backup job by using the Backup Wizard	331
Preventing the Backup Wizard from launching from the Backup button	332
Configuring the Backup Wizard to launch from the Backup button	332
Creating a backup job by setting job properties	332
Selections options for backup jobs	336
Resource Credentials options	337
Resource Order Backup options	338
Device and media options for backup jobs and templates	339
General options for backup jobs and templates	342

Advanced options for backup jobs	348
Pre/post commands for backup or restore jobs	352
Job Summary properties	355
How to include or exclude files for backup	355
About scheduling jobs	356
Scheduling jobs	356
About the scheduling calendar	359
Scheduling a job to run on specific days	359
Scheduling a job to run on recurring week days	360
Scheduling a job to run on recurring days of the month	361
Scheduling a job to run on a day interval	362
Setting the effective date for a job schedule	363
About time windows	364
Setting the time window for a scheduled job	364
Restarting a job during a time interval	365
Excluding dates from a schedule	366
Configuring default schedule options	366
About the full backup method for backing up and deleting files	367
Backing up and deleting files	368
About duplicating backed up data	369
Duplicating backed up data	369
How to copy data directly from a virtual tape library to a physical tape device	381
Verifying a backup	382
Selections properties for verify jobs	383
Device properties for verify jobs	384
General properties for verify jobs	384
About test run jobs	385
Creating a test run job	386
Setting test run default options	387
 Chapter 7	
Customizing backup options	391
Setting default backup options	391
Default Backup options	392
About pre/post commands	399
Setting default pre/post commands	400
About specifying backup networks	402
About using IPv4 and IPv6 in Backup Exec	404
Setting default backup network and security options	404
About using Backup Exec with Symantec Endpoint Protection	408
About using Backup Exec with firewalls	409
Backup Exec Ports	411

Backup Exec Listening Ports	412
Backup Exec Desktop and Laptop Option ports	413
Browsing systems through a firewall	414
About enabling a SQL instance behind a firewall	414
About encryption	415
About software encryption	415
About hardware encryption	416
Encryption keys	416
About restricted keys and common keys in encryption	417
About pass phrases in encryption	417
About encryption key management	418
Creating an encryption key	420
Replacing an encryption key	421
About deleting an encryption key	421
Deleting an encryption key	422
About restoring encrypted data	422
About cataloging media that contains encrypted backup sets	423
About configuring DBA-initiated job settings	423
Creating a template for DBA-initiated jobs	424
Editing DBA-initiated jobs	434
Deleting a job template for DBA-initiated jobs	435
About preferred server configurations	435
Creating preferred server configurations	436
Deleting preferred server configurations	438
Editing settings for preferred server configurations	438
Designating a default preferred server configuration	438
Removing the default status for a preferred server configuration	439
Chapter 8	
About devices	441
About storage devices	441
About the Configure Devices Assistant	443
Configuring storage devices by using the Configure Devices Assistant	444
About sharing storage	444
Managing shared storage	445
Pausing a media server	446
Resuming a media server	446
Pausing storage devices	446
Resuming storage devices	447
Renaming storage devices	447

About inventorying media	448
Inventorying media in a device	448
Erasing media	449
Chapter 9	
Managing tape drives and robotic libraries	451
About tape drives and robotic libraries	451
About the Virtual Tape Library Unlimited Drive Option	452
About the Library Expansion Option	453
About configuring tape devices by using the Tape Device Configuration Wizard	453
About adding or replacing devices by using the Hot-swappable Device Wizard	453
Adding or replacing devices by using the Hot-swappable Device Wizard	454
About installing Symantec tape device drivers	455
Installing Symantec tape device drivers by running tapeinst.exe	455
Installing Symantec tape device drivers by using the Tape Device Configuration Wizard	456
Changing the preferred block size, buffer size, buffer count, and high water count for devices	456
Enabling hardware compression for devices	457
Specifying read and write operations on types of media	457
Viewing storage device properties	458
General properties for devices	458
Configuration properties for devices	460
SCSI information for devices	463
Statistics properties for devices	463
Cleaning properties for devices	465
Media type properties for devices	466
About robotic libraries in Backup Exec	467
Requirements for setting up robotic library hardware	468
Troubleshooting the display of robotic library devices	469
Initializing robotic libraries when the Backup Exec service starts	470
Enabling bar code rules for robotic libraries	470
Defining a cleaning slot	471
Configuration properties for robotic libraries	471
Statistics properties for robotic libraries	472
Properties for robotic library slots	472
About robotic library partitions	475
About creating utility jobs to help manage devices and media	480

Utility jobs for virtual tape libraries and simulated tape libraries	482
General options for utility jobs	482
Inventorizing robotic libraries when Backup Exec services start	483
Creating a job to initialize a robotic library	484
Retensioning a tape	484
Formatting media in a drive	485
Labeling media	486
Ejecting media from a drive	487
Creating a cleaning job	488
About importing media to a robotic library	489
Exporting media from a robotic library	490
About exporting expired media from a robotic library	491
Locking the robotic library's front panel	493
Unlocking the robotic library's front panel	494
Chapter 10 Managing backup-to-disk folders	495
About backup-to-disk folders	496
Requirements for creating a backup-to-disk folder	497
Requirements for creating a removable backup-to-disk folder	498
Creating a backup-to-disk folder by using the Backup-to-Disk Wizard	498
Creating a backup-to-disk folder by setting properties	499
About sharing backup-to-disk folders	505
Sharing an existing backup-to-disk folder	506
Changing the path of a backup-to-disk folder	506
Deleting a backup-to-disk folder	507
Recreating a backup-to-disk folder and its contents	507
Changing the status of a device to online	508
Renaming a backup-to-disk file	508
Deleting a backup-to-disk file	509
Recreating a deleted backup-to-disk file	509
Erasing backup-to-disk files	510
Recommendations for using backup-to-disk folders with backup jobs that use Granular Recovery Technology	511
How to reclaim disk space for backup jobs that use Granular Recovery Technology	513

Chapter 11	Managing device pools	515
	About device pools	515
	Creating device pools	516
	Device pool options	517
	Adding devices to a device pool	517
	Setting priorities for devices in a device pool	518
	Removing devices from a device pool	519
	Deleting device pools	519
	Retarget Job options	520
	Device Pool Properties	520
Chapter 12	Policies and templates	521
	About policies and templates	521
	Creating a new policy	522
	Creating a new policy using the Policy Wizard	523
	Editing a policy	525
	Deleting a policy	526
	Using an example policy	526
	Re-creating example policies	528
	About using templates in policies	529
	Adding a backup template to a policy	530
	About the verify backup sets templates	533
	Adding a verify backup sets template to a policy	534
	About export media templates	536
	Adding an export media template to a policy	537
	Importing a template into a policy	538
	Editing a template in a policy	539
	Deleting a template from a policy	539
	About template rules	540
	Setting template rules	542
	Changing template rules	542
	Deleting template rules	543
	About creating jobs using policies and selection lists	544
	Creating new jobs for a policy	544
	Creating new jobs for a selection list	545
	Viewing the policies that are designated to back up selection lists	545
	Viewing the selection lists that are designated for backup by policies	546
	Editing the next occurrence of a policy-based job	546
	Deleting a job created from a policy	546
	Renaming a job created from a policy	547

	About duplicate backup set templates	548
	Adding a duplicate backup template to a policy	550
Chapter 13	Administrating Backup Exec	553
	About administrating Backup Exec	553
	Copying jobs, selection lists, or policies	554
	Copy to Media Server options	555
	Viewing the job log for a copy to media server job	556
	Editing job properties	556
	Job Monitor options	557
	Viewing properties for active jobs	557
	Searching for text in the job history or job properties	561
	Canceling an active job	562
	Placing all scheduled occurrences of an active job on hold	563
	Removing the hold on a scheduled job	563
	Active job statuses	563
	Scheduled job statuses	565
	Running a scheduled job immediately	568
	Placing a scheduled job on hold	568
	Removing the hold on a scheduled job	569
	Placing the job queue on hold	569
	Removing the hold on the job queue	570
	Changing the priority for a scheduled job	570
	Running a test job for a scheduled job	570
	Deleting scheduled jobs	571
	Viewing the properties for completed jobs	572
	Viewing the history of a job, policy, or selection list	575
	Deleting completed jobs	576
	Linking from the job log to the Symantec Technical Support Web site	577
	Completed job statuses	577
	Configuring default job log options	580
	About using job logs with vertical applications	581
	Filtering jobs	582
	About managing custom filters	582
	Creating a custom filter for jobs	583
	Creating a custom filter for current jobs	583
	Creating a custom filter for jobs in the job history	585
	Deleting custom filters	587
	Editing custom filters	588
	Viewing the job workload for a media server from the Calendar tab	588

Viewing jobs for specific days on the calendar	589
Managing jobs from the Calendar tab	589
Viewing the Symantec Endpoint Protection Security Summary	590
About error-handling rules	590
Creating a custom error-handling rule	591
Custom error-handling rule for recovered jobs	594
Cluster failover error-handling rule	595
How thresholds are used to stall, fail, and recover jobs	595
Setting thresholds to recover jobs	596
Application settings for job status and recovery options	597

Chapter 14

Restoring data	599
About restoring data	599
Restore jobs and the catalog	600
Changing application settings for the catalogs	601
Catalog levels	603
Restoring data by using the Restore Wizard	604
Preventing the Restore Wizard from launching from the Restore button	604
Configuring the Restore Wizard to launch from the Restore button	605
Restoring data by setting job properties	605
Selections options for restore jobs	609
Device options for restore jobs	611
General options for restore jobs	612
Advanced options for restore jobs	613
Network and security restore options	617
Running pre and post commands for restore jobs	618
About restoring file permissions	619
About System State	620
Restoring System State	621
About restoring Shadow Copy Components	622
About restoring utility partitions	623
About performing redirected restores of utility partitions	623
About restoring media created with other backup software	624
About restoring data from ARCserve media	624
Restoring data from ARCserve media	625
About selecting data to restore	626
Creating a restore selection list	628
Changing and testing resource credentials for restore jobs	630
Searching for files to restore	631
About redirecting restore jobs	634

	File Redirection restore options	634
	About redirecting restore jobs to native Microsoft Virtual Hard Disk (VHD) files	636
	Using redirected restore for Active Directory, Active Directory Application Mode for Windows Server 2003/2008	636
	Setting defaults for restore jobs	638
	Default restore options	638
	Canceling a restore job	641
Chapter 15	Alerts and notifications	643
	About alerts and notifications	644
	About alert views	645
	Active Alerts view and Alert History view	646
	Viewing alerts	648
	Filtering alerts	648
	Creating custom filters for alerts	649
	Editing custom filters for alerts	650
	Deleting custom filters for alerts	650
	Viewing alert properties	651
	Viewing the job log from an alert	653
	Responding to active alerts	654
	About automatic responses for alert categories	654
	Configuring automatic responses for alert categories	655
	Clearing informational alerts from the Active Alerts pane	657
	Alert response options	657
	Configuring alert category properties	658
	Configure Alert Categories options	659
	Enabling or disabling alerts from the Active Alerts pane	661
	Deleting alerts from the Alert History	661
	Setting up notification for alerts	661
	Configuring SMTP for email or mobile phone text message notification	662
	Configuring MAPI email for notification	664
	Configuring VIM email for notification	665
	Configuring a pager for alert notification	666
	Configure Recipients options	666
	Configuring SMTP email or mobile phone text messaging for a person recipient	667
	Configuring MAPI mail for a person recipient	668
	Configuring VIM mail for a person recipient	670
	Configuring a pager for a person recipient	671
	Configuring a Net Send recipient	674

Configuring a printer recipient	675
Configuring a group recipient	677
Scheduling notification for recipients	677
Editing recipient notification properties	678
Editing recipient notification methods	679
Removing recipients	679
Assigning recipients to alert categories for notification	680
Assign Recipients to Alert Categories options	680
Stopping alert notification for a recipient	681
Sending a notification when a job completes	682
Sending a notification when a selection list is used in a job	682
Notification options for jobs	682
About SNMP notification	683
Installing and configuring the SNMP system service	686
Installing the Windows Management Instrumentation performance counter provider	687
Installing the Windows Management Instrumentation provider for SNMP	687
Uninstalling the Windows Management Instrumentation performance counter provider	688
Uninstalling the Windows Management Instrumentation provider for SNMP	688

Chapter 16

Reports in Backup Exec	689
About reports in Backup Exec	690
Viewing the list of available reports	691
Running a report	691
Additional settings for standard reports	692
Available groups for creating reports	693
Running a new report job	694
General options for a new report job	695
Saving a report	696
Saving a report to a new location	696
Printing a report from the Backup Exec Report Viewer	697
Printing a report that is saved in PDF format	697
Printing a report that is saved in HTML format	697
Deleting a report from Job History	698
About scheduling report jobs and setting notification recipients	698
About custom reports in Backup Exec	699
Creating a custom report	699
Custom report name and description options	700
Field options for custom reports	701

About grouping fields in custom reports	702
Sorting fields in custom reports	704
Setting graph options in custom reports	706
Example graphs for custom reports	708
Previewing custom reports	712
Setting filters for custom reports	712
Filter expressions for defining custom reports	715
Copying custom reports	718
Editing custom reports	718
Deleting custom reports	719
Editing application settings for reports	719
Application settings for reports	720
Viewing report properties	721
General properties for reports	721
Available reports	722
Active Alerts Report	729
Active Alerts by Media Server Report	730
Alert History Report	730
Alert History by Media Server Report	731
Application Event Log Report	732
Audit Log Report	732
Backup Job Success Rate Report	733
Backup Recommendations Report	733
Backup Resource Success Rate Report	734
Backup Set Details by Resource Report	734
Backup Sets by Media Set Report	735
Backup Size By Resource Report	736
Configuration Settings Report	737
Current Job Status Report	738
Daily Device Utilization Report	738
Deduplication device summary	739
Deduplication summary	740
Device Summary Report	740
Device Usage by Policy	741
Error-Handling Rules Report	742
Event Recipients Report	743
Failed Backup Jobs Report	744
Job Distribution by Device Report	745
Jobs Summary Report	745
Machines Backed Up Report	746
Managed Media Servers Report	747
Media Audit Report	748
Media Errors Report	749

Media Required for Recovery Report	749
Media Set Report	750
Media Vault Contents Report	751
Missed Availability Report	752
Move Media to Vault Report	752
Operations Overview Report	753
Overnight Summary Report	755
Policy Jobs by Resource Summary Report	756
Policy Jobs Summary Report	757
Policy Properties Report	758
Policy Protected Resources	759
Problem Files Report	759
Recently Written Media Report	760
Resource Backup Policy Performance Report	761
Resource Risk Assessment Report	761
Resources Protected by Policy report	762
Restore Set Details by Resource Report	762
Retrieve Media from Vault Report	763
Robotic Library Inventory Report	764
Scheduled Server Workload	765
Scratch Media Availability Report	766
Selection Lists Report	766
Test Run Results Report	767
Archive Job Success Rate report	768
Archive Selections by Archive Rules and Retention Categories report	768
Exchange Mailbox Group Archive Settings report	769
Failed Archive Jobs report	770
File System Archive Settings report	770
Overnight Archive Summary report	771
Vault Store Usage Details report	772
Vault Store Usage Summary Report	773

Chapter 17	Disaster preparation and recovery	775
	About disaster preparation	775
	About key elements of a disaster preparation plan (DPP)	776
	Returning to the last known good configuration	777
	Creating a hardware profile copy	778
	About creating an emergency repair disk (Windows 2000 computers only)	779
	About manual disaster recovery of Windows computers	780

	About a manual disaster recovery of a local Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)	780
	Running a manual disaster recovery of a local Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)	781
	About a disaster recovery operation of a remote Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)	785
	Running a disaster recovery operation on a remote Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)	786
Chapter 18	Troubleshooting	789
	Troubleshooting hardware-related issues	789
	How to get more information about alerts and error messages	794
	Troubleshooting backup issues	795
	About cluster sizes for NTFS partitions	796
	Troubleshooting restore issues	797
	How to improve Backup Exec's performance	797
	About the Symantec Knowledge Base	801
	Searching the Symantec Knowledge Base	802
	How to contact Technical Support	802
	About Backup Exec diagnostic tools	802
	About the Symantec Backup Exec Support Tool	803
	About the Backup Exec diagnostic application	804
	How to use the Symantec Gather Utility for troubleshooting	808
	Running the begather utility to troubleshoot Backup Exec components on Linux servers	809
	Using the Backup Exec Debug Monitor for troubleshooting	810
	About the Backup Exec debug tool	810
Chapter 19	Using Symantec Backup Exec with Server Clusters	813
	About Backup Exec and clusters	814
	Requirements for clustering Backup Exec in a Microsoft Cluster Server environment	815
	How Backup Exec works in a Microsoft Cluster Server environment	816
	Requirements for installing Backup Exec on a Microsoft Cluster Server	816

Installing Backup Exec on a Microsoft Cluster Server	818
Upgrading Backup Exec on a Microsoft cluster	819
Installing additional Backup Exec options on a Microsoft cluster	820
Uninstalling Backup Exec from a Microsoft cluster	820
Creating device pools for Microsoft Cluster Servers	821
Using checkpoint restart on Microsoft Cluster Server failover	822
Enabling or disabling checkpoint restart	824
Specifying a different failover node	824
Designating a new SAN SSO primary server and central administration server in a Microsoft Cluster Server	825
Configurations for Backup Exec and Microsoft Cluster Servers	827
Two-node cluster with locally attached storage devices	828
Two-node cluster with tape devices on a shared SCSI bus	828
Configuring a shared SCSI bus for tape devices	829
Multi-node clusters on a fibre channel SAN with the SAN SSO	832
Using the Central Admin Server Option with Microsoft clusters and SAN SSO	835
About backing up Microsoft Cluster Servers	836
Backing up local disks in a Microsoft cluster	837
Backing up shared disks in a Microsoft cluster	838
Backing up database files in a Microsoft cluster	838
Backing up Windows 2008 R2 cluster shared volumes	839
About restoring data to a Microsoft cluster	840
Restoring the cluster quorum for Windows Server 2003/2008 computers to a Microsoft cluster	840
Specifying a new drive letter for the cluster quorum disk	841
Using Backup Exec with Veritas Cluster Server	842
Requirements for installing Backup Exec with the CASO option on a Veritas Cluster Server	843
Installing Backup Exec with the CASO option on a Veritas Cluster Server	844
Requirements for clustering Backup Exec using Veritas Cluster Server	844
Clustering Backup Exec using Veritas Cluster Server	845
About backing up Veritas Cluster Servers	846
About backing up Windows 2000 and Windows Server 2003/2008 features in a Veritas cluster	847
Backing up local disks in a Veritas cluster	848
Backing up shared disks in a Veritas cluster	848
Backing up database files in a Veritas cluster	849
About restoring data to Veritas Cluster Servers	849

About backup job failover with Veritas Cluster Servers	850
Disaster recovery of a cluster	850
Using IDR to prepare for disaster recovery of a cluster	851
Recovering nodes on the cluster using IDR	851
Recovering Backup Exec on a Microsoft cluster using IDR	852
Recovering the entire cluster using a manual disaster recovery procedure	853
Restoring the Microsoft Cluster Server data files	854
Recovering all shared disks in a Microsoft cluster	855
Recovering all shared disks in a Veritas cluster	856
Recovering Backup Exec in a Microsoft cluster	856
Troubleshooting clusters	857
Changing the Quorum disk signature	859
Manually joining two cluster disk groups and resynchronizing volumes	860
Chapter 20 Using Backup Exec Retrieve	861
About Backup Exec Retrieve	861
How Backup Exec Retrieve works	862
What end users can do with Backup Exec Retrieve	864
Before you install Backup Exec Retrieve	866
Requirements for installing Backup Exec Retrieve on a Web server	866
Requirements for using Backup Exec Retrieve on end users' computers	868
About deploying the Silverlight run time in your organization	869
Upgrading from Backup Exec Retrieve that runs under Backup Exec System Recovery Manager 8.5	869
Installing Backup Exec Retrieve	869
About configuring Backup Exec Retrieve	871
Adding a data source	871
Editing a data source	873
Deleting a data source	874
Editing application settings for Backup Exec Retrieve	874
Application settings for Backup Exec Retrieve	874
Uninstalling Backup Exec Retrieve	876
Troubleshooting Backup Exec Retrieve	876

Appendix A	Symantec Backup Exec Active Directory Recovery Agent	879
	About the Active Directory Recovery Agent	880
	Requirements for the Active Directory Recovery Agent	880
	About installing the Active Directory Recovery Agent	881
	How the Active Directory Recovery Agent works	882
	How Granular Recovery Technology works with Active Directory and ADAM/AD LDS backups	883
	Editing defaults for Active Directory and ADAM/AD LDS backup and restore jobs	883
	Microsoft Active Directory default options	884
	Backing up Active Directory	885
	Backing up ADAM/AD LDS	886
	Active Directory Recovery Agent backup job options	887
	About restoring individual Active Directory and ADAM/AD LDS objects	888
	Restoring individual objects from an Active Directory backup	890
	Restoring individual objects from an ADAM/AD LDS backup	891
	About recreating purged Active Directory and ADAM/AD LDS objects	892
	Recreating purged Active Directory objects	893
	Recreating purged ADAM/AD LDS objects	894
	Resetting the Active Directory computer object and the computer object account	895
Appendix B	Symantec Backup Exec Advanced Disk-based Backup Option	897
	About the Advanced Disk-based Backup Option	898
	About installing the Advanced Disk-based Backup Option	898
	About the synthetic backup feature	899
	What you can back up with synthetic backup	901
	Requirements for synthetic backup	901
	Best practices for synthetic backup	902
	About collecting additional information for synthetic backup and true image restore	904
	Methods for creating a synthetic backup	904
	Creating a synthetic backup by using the Policy Wizard	905
	About creating a synthetic backup by copying the example policy	906

Creating a synthetic backup by adding templates to a policy	907
Creating template rules to run job templates for synthetic backup	910
General options for synthetic backup templates	911
Advanced options for synthetic backup templates	911
About true image restore	912
Requirements for true image restore	915
Best practices for true image restore	915
Enabling backups for true image restore	916
About true image catalogs	916
About restoring a backup set enabled for true image restore	917
Selecting backup sets that are enabled for true image restore	918
Troubleshooting tips for true image restore	918
About offhost backup	919
Requirements for offhost backup	921
Requirements for offhost backup when using the Veritas Storage Foundation for Windows Provider	922
Best practices for offhost backup	923
Browsing remote computers for installed snapshot providers	925
Setting offhost backup options for a backup job	925
Backup options for the Advanced Disk-based Backup Option	926
Setting default options for offhost backup jobs	928
Configuring a GRT-enabled offhost backup for Exchange resources	928
About restoring offhost backup data	929
Troubleshooting the offhost backup	929
Offhost backup failures when using VSW as a provider	933
Offhost backup issues when using a hardware provider	935

Appendix C

Symantec Backup Exec Advanced Open File Option	937
About the Advanced Open File Option	937
About supported snapshot technologies	940
Requirements for using Advanced Open File Option	941
How to install the Advanced Open File Option	941
Installing the Advanced Open File Option to remote Windows computers using the command line	942
Setting default options for the Advanced Open File Option	943
About Snap Start on a Veritas Storage Foundation volume	944
Using Snap Start on a Veritas Storage Foundation volume	944

Best practices for using the Symantec Volume Snapshot Provider	945
About the Symantec Volume Snapshot Provider cache file location	946
How to adjust the Symantec Volume Snapshot Provider cache file size	947
Configuring the Advanced Open File Option for backup jobs	948
Advanced Open File options	949
About the job log and the Advanced Open File Option	951

Appendix D	Symantec Backup Exec Agent for DB2 on Windows Servers	953
	About the Backup Exec DB2 Agent	953
	Requirements for the DB2 Agent	954
	Configuring the DB2 Agent on Windows computers	955
	Adding information about DB2 servers to the media server's authentication credentials list	955
	Configuring database access for DB2 operations on Windows computers	959
	Adding a DB2 instance to the DB2 Agent on Windows computers that run the Remote Agent Utility	963
	Editing a DB2 instance by using the Remote Agent Utility	963
	Deleting a DB2 instance by using the Remote Agent Utility	963
	Backing up DB2 resources	964
	DB2 backup options	966
	Restoring DB2 data	967
	DB2 restore options	968
	Redirecting a restore of DB2 data	970
	About using DB2 to run DBA-initiated jobs	972
	About using the DB2 database archive logging methods	974
	About the db2.conf file	975
	Editing a db2.conf file	975
	Example db2.conf file	976
	Troubleshooting DB2	978

Appendix E	Symantec Backup Exec Agent for Enterprise Vault	981
	About the Enterprise Vault Agent	982
	Requirements for the Enterprise Vault Agent	987
	About installing the Enterprise Vault Agent	988
	About backup methods for Enterprise Vault backup jobs	988
	Enterprise Vault backup options	992

Setting a default backup method for Enterprise Vault backup jobs	993
About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases	993
Backing up an Enterprise Vault open partition	994
Backing up Enterprise Vault closed partitions	995
Backing up Enterprise Vault 8.x ready partitions	997
Backing up the Enterprise Vault Directory database	998
Backing up the Enterprise Vault Monitoring database	999
Backing up an Enterprise Vault vault store database	1000
Backing up the Enterprise Vault 8.x Audit database	1001
Backing up the Enterprise Vault 8.x FSA Reporting database	1002
Backing up the Enterprise Vault 8.x Fingerprint database	1003
Backing up the Enterprise Vault 8.x Compliance Accelerator Configuration database and Compliance Accelerator customer databases	1004
Backing up the Enterprise Vault 8.x Discovery Accelerator Configuration database and Discovery Accelerator customer databases	1005
Backing up the Discovery Accelerator Custodian database	1006
Backing up an Enterprise Vault vault store	1007
About backing up an Enterprise Vault 7.x server and an Enterprise 8.x site	1009
Backing up an Enterprise Vault 7.x server	1009
Backing up an Enterprise Vault site	1010
Backing up Enterprise Vault index locations	1011
About restoring Enterprise Vault	1012
About automatic redirection of Enterprise Vault components under an Enterprise Vault server	1014
Restoring the Enterprise Vault Directory database	1015
Restoring the Enterprise Vault Monitoring database	1016
Restoring Enterprise Vault partitions	1017
Restoring an Enterprise Vault vault store database	1019
Restoring an Enterprise Vault 8.x Audit database	1021
Restoring the Enterprise Vault 8.x FSA Reporting database	1022
Restoring the Enterprise Vault 8.x Fingerprint database	1023
Restoring the Compliance Accelerator Configuration database	1024
Restoring the Compliance Accelerator Customer database	1025
Restoring the Discovery Accelerator Configuration database	1026
Restoring the Discovery Accelerator Custodian database	1027
Restoring the Discovery Accelerator Customer database	1028

About restoring individual files and folders with the Enterprise Vault Agent	1029
Restoring individual files from partitions by using the Enterprise Vault Agent	1030
Restoring individual folders from an Enterprise Vault index backup	1032
Restoring an Enterprise Vault 7.x server to its original location	1033
Enterprise Vault restore options	1034
Redirecting an Enterprise Vault restore job	1036
Redirecting the restore of Enterprise Vault databases to a different Microsoft SQL Server computer	1039
Best practices for the Enterprise Vault Agent	1041
About the Backup Exec Migrator for Enterprise Vault	1041
Backup Exec Migrator for Enterprise Vault requirements	1041
How the Backup Exec Migrator works	1042
Configuring the Backup Exec Migrator	1049
About the Restore view of migrated Enterprise Vault data	1057
About retrieving migrated Enterprise Vault data	1058
About the Partition Recovery Utility	1059
Best practices for using the Backup Exec Migrator	1061
Troubleshooting Backup Exec Migrator and Partition Recovery Utility issues	1062
 Appendix F	
Symantec Backup Exec Agent for Lotus Domino Server	1063
About the Agent for Lotus Domino Server	1064
Lotus Domino Agent requirements	1064
About installing the Lotus Domino Agent on the media server	1066
About the Lotus Domino Agent and the Domino Attachment and Object Service (DAOS)	1066
Best practices for restoring the missing .nlo files	1067
Viewing Lotus Domino databases that are created while Backup Exec is running	1068
Viewing Lotus Domino databases that are on the local server	1068
Viewing Lotus Domino databases that are on remote computers	1069
Configuring default Lotus Domino options	1069
Lotus Domino default options	1070
About backing up Lotus Domino databases	1071
About automatic exclusion of Lotus Domino files during volume-level backups	1073
About supported Lotus Domino database configurations	1073

About Lotus Domino transaction logs	1074
About selecting Lotus Domino databases for backup	1075
Selecting Lotus Domino databases for backup	1076
Selecting backup options for Lotus Domino databases	1076
Lotus Domino backup job options	1076
Restoring Lotus Domino databases	1078
About selecting Lotus Domino databases for restore	1079
Selecting restore options for Lotus Domino databases	1082
Lotus Domino restore options	1082
Redirecting restore jobs for Lotus Domino databases	1083
Redirecting the restore of DAOS NLO files	1084
How to prepare for disaster recovery on a Lotus Domino server	1086
Recovering a Lotus Domino server from a disaster	1087
About disaster recovery of a Lotus Domino server using archive logging	1090
Recovering a Lotus Domino server that uses circular logging	1090
Recovering the Lotus Domino server, databases, and transaction logs when archive logging is enabled	1091
 Appendix G	
Symantec Backup Exec Agent for Microsoft Exchange Server	1093
About the Backup Exec Exchange Agent	1094
Requirements for using the Exchange Agent	1095
About installing the Exchange Agent	1099
Recommended configurations for Exchange	1100
Requirements for accessing Exchange mailboxes	1101
Backup strategies for Exchange	1102
Automatic exclusion of Exchange data during volume-level backups	1105
About the circular logging setting for Exchange	1106
How Granular Recovery Technology works with the Exchange Information Store	1106
About Backup Exec and Microsoft Exchange Web Services	1107
Snapshot and offhost backups with the Exchange Agent	1108
Troubleshooting Exchange Agent snapshot and offhost jobs	1109
Configuring a snapshot backup for Exchange resources	1110
Configuring an offhost backup with the Exchange Agent	1111
About continuous protection for Exchange data	1112
Requirements for installing components for CPS Exchange backup jobs	1113

Requirements for configuring continuous protection for Exchange data	1115
Best practices for continuous protection of Exchange	1117
About managing the CPS Exchange backup job for Exchange data	1118
About reviewing disk space availability for CPS Exchange backup jobs	1119
Stopping CPS Exchange backup jobs temporarily	1120
Viewing the CPS console from Backup Exec	1121
About using recovery points to restore individual Exchange items to a point in time	1122
Troubleshooting CPS Exchange backup jobs	1123
Setting default backup and restore options for Exchange data	1123
Default backup and restore options for Exchange	1123
About backing up Exchange 2003/2007	1129
About backing up Exchange Server 2010 Databases	1130
Adding an Exchange Server 2010 forest to backup selections	1130
Managing an Exchange Server 2010 forest	1131
Backing up Exchange	1132
Microsoft Exchange backup options	1134
About selecting individual Exchange mailboxes for backup	1140
Backing up individual Exchange mailboxes	1143
About restoring Exchange data	1144
Requirements for restoring Exchange 2000 or later	1145
Configuring a database in Exchange	1146
Dismounting Exchange databases that are being restored	1146
About restoring data using the Exchange 2003/2007 recovery storage group or Exchange Server 2010 recovery database	1147
Restoring a database to an Exchange 2007 recovery storage group	1148
About restoring Exchange data from snapshot backups	1149
About restoring Exchange data from continuous protection backups	1150
About restoring Exchange mailboxes and public folders from mailbox backups	1152
Restoring individual Exchange public folder messages from tape by duplicating backup sets to disk	1153
Restoring Exchange data	1154
About redirecting Exchange restore data	1159
About redirecting Exchange storage group and database restores	1160

	About redirecting restores of mailboxes or public folders that were backed up separately from the Exchange Information Store	1160
	Redirecting Exchange restore data	1162
	How to prepare for disaster recovery of Exchange Server	1165
	Recovering from a disaster for Exchange 2000 or later	1166
Appendix H	Symantec Backup Exec Agent for Microsoft Hyper-V	1169
	About the Agent for Microsoft Hyper-V	1169
	About installing the Agent for Microsoft Hyper-V	1170
	Requirements for using the Agent for Microsoft Hyper-V	1171
	About upgrading from the Agent for Microsoft Virtual Servers	1173
	About backup selections for Microsoft Hyper-V	1173
	How Backup Exec automatically protects new virtual machines during a backup job	1174
	Backing up data by using the Agent for Microsoft Hyper-V	1175
	Microsoft Hyper-V backup options	1175
	Virtual Machine Application Granular Recovery Technology Settings	1176
	How Granular Recovery Technology works with the Agent for Microsoft Hyper-V	1177
	How Backup Exec protects Microsoft Exchange, SQL, and Active Directory data on virtual machines	1178
	Requirements for protecting Microsoft Exchange, SQL, and Active Directory data on virtual machines	1179
	About restore selections for Microsoft Hyper-V	1180
	Restoring data to the Hyper-V host	1182
	Microsoft Hyper-V restore options	1183
	Restoring a virtual machine to a different host	1184
	Microsoft Hyper-V Redirection options	1184
	Setting default backup and restore options for the Agent for Microsoft Hyper-V	1186
	Microsoft Hyper-V default options	1186
	About backing up and restoring highly available virtual machines	1188
Appendix I	Symantec Backup Exec Agent for Microsoft SharePoint	1189
	About the SharePoint Agent	1190
	Requirements for the SharePoint Agent	1190
	About installing the SharePoint Agent	1191

Adding a SharePoint server farm to the backup selections list	1191
Add Server Farm options	1193
Manage SharePoint Server Farms options	1193
Server Farm Properties	1194
How Backup Exec automatically protects new SharePoint resources	
during a backup job	1194
Changing the name of a SharePoint server farm	1195
Deleting a farm from the Microsoft SharePoint Server Farms	
node	1195
Disabling or enabling communication between a SharePoint Web	
server and Backup Exec	1196
Setting default options for SharePoint Server 2003 and later	1196
Microsoft SharePoint default options	1197
About using the SharePoint Agent with SharePoint Server 2010 and	
Windows SharePoint Foundation 2010	1200
Backing up a farm for Microsoft Office SharePoint Server 2010	
or Windows SharePoint Foundation 2010	1202
Backing up individual SharePoint 2010 Web applications in a	
Microsoft SharePoint server farm	1203
About restoring resources for SharePoint Server 2010 and	
SharePoint Foundation 2010	1203
Restoring resources for SharePoint Server 2010 and SharePoint	
Foundation 2010	1204
Restoring individual SharePoint 2010 objects from full database	
backups to their original locations	1205
Restoring a Microsoft Office SharePoint Server 2010 Web	
application to its original location	1207
Redirecting a restore job for SharePoint 2010	1208
Redirecting the restore of individual SharePoint 2010 objects to	
a file path	1210
Redirecting the restore of a Microsoft Office SharePoint Server	
2010 Web application	1211
Redirecting the restore of a Microsoft Office SharePoint Server	
2010 database to an alternate SQL instance	1213
Recovering SharePoint 2010 data after a disaster	1214
About using the SharePoint Agent with SharePoint Server 2007 and	
Windows SharePoint Services 3.0	1218
About adding a SharePoint 2007 server farm to the backup	
selections list	1219
Backing up a farm for Microsoft Office SharePoint Server 2007	
or a Windows SharePoint Services 3.0	1219
Backing up individual SharePoint 2007 Web applications in a	
Microsoft SharePoint server farm	1220

About restoring resources for SharePoint Server 2007 and SharePoint Services 3.0	1221
Restoring resources for SharePoint Server 2007 and SharePoint Services 3.0	1222
Restoring individual SharePoint 2007 objects from full database backups to their original locations	1223
Restoring a Microsoft Office SharePoint Server 2007 Shared Services Provider	1224
Restoring a Microsoft Office SharePoint Server 2007 Web application to its original location	1225
Redirecting a restore job for SharePoint 2007	1227
Redirecting the restore of individual SharePoint 2007 objects to a file path	1228
Redirecting the restore of a Microsoft Office SharePoint Server 2007 Web application	1229
Redirecting the restore of a Microsoft Office SharePoint Server 2007 database to an alternate SQL instance	1230
Recovering SharePoint 2007 data after a disaster	1231
About using the SharePoint Agent with SharePoint Portal Server 2003 and Windows SharePoint Services 2.0	1235
About selecting SharePoint Server 2003 resources for backup	1236
Backing up resources from SharePoint 2003	1236
About selecting SharePoint 2003 resources for restore	1237
Restoring SharePoint 2003 resources	1238
Restoring individual SharePoint 2003 objects (Microsoft SQL Server-based) from full database backups	1238
Restoring SharePoint 2003 document libraries (Web storage system-based)	1240
Restoring previous versions of SharePoint 2003 documents from document library (Web storage system-based) backups	1241
Redirecting a restore job for SharePoint 2003	1241
Redirecting the restore of SharePoint 2003 document library (Web storage system-based) data to another document library	1242
Redirecting the restore of individual SharePoint 2003 objects to a file path	1243
Redirecting the restore of a Microsoft Office SharePoint Server 2003 database to an alternate SQL instance	1244
Microsoft SharePoint backup options	1246
Microsoft SharePoint restore options	1247
Microsoft SharePoint redirection options	1250

Appendix J	Symantec Backup Exec Agent for Microsoft SQL Server	1253
	About the Agent for Microsoft SQL Server	1254
	Requirements for using the SQL Agent	1255
	About installing the SQL Agent	1256
	How to use Backup Exec logon accounts for SQL resources	1256
	About backup strategies for SQL	1258
	SQL backup strategy recommendations	1259
	About consistency checks for SQL	1261
	How to use snapshot technology with the SQL Agent	1262
	How to use AOFO with the SQL Agent	1263
	How to use ADBO with the SQL Agent	1264
	Setting default backup and restore options for SQL	1265
	Microsoft SQL default options	1265
	Setting backup options for SQL	1273
	SQL backup options	1273
	About automatic exclusion of SQL data during volume level backup	1280
	About backing up SQL databases	1281
	About backing up SQL filegroups	1282
	Displaying SQL filegroups on the backup selections pane	1284
	Backing up SQL filegroups	1284
	How to back up SQL transaction logs	1285
	About SQL 2005 or later database snapshots	1286
	Creating SQL database snapshots	1288
	Setting restore options for SQL	1289
	SQL restore options	1289
	About restoring SQL databases and file groups	1293
	About restoring encrypted SQL databases	1295
	Restoring from SQL database backups	1295
	How to restore from SQL transaction logs up to a point in time	1296
	How to restore from SQL transaction logs up to a named transaction	1297
	About restoring from SQL filegroup backups	1298
	Restoring an entire SQL database, a missing primary filegroup, or a filegroup containing a deleted or changed table	1299
	Restoring a missing or corrupted SQL nonprimary filegroup	1300
	About restoring the SQL master database	1301
	Restarting SQL using database copies	1302
	Restoring the master database	1304

	About redirecting restores for SQL	1305
	Redirecting restores for SQL	1305
	About reverting SQL 2005 or later databases using database snapshots	1309
	About disaster recovery of a SQL Server	1310
	How to prepare for disaster recovery of SQL	1311
	Requirements for SQL disaster recovery	1311
	Disaster recovery of SQL	1312
Appendix K	Symantec Backup Exec Agent for Oracle on Windows or Linux Servers	1315
	About the Backup Exec Oracle Agent	1315
	About installing the Oracle Agent	1316
	Upgrading the Backup Exec Oracle Agent	1317
	Configuring the Oracle Agent on Windows computers and Linux servers	1318
	Configuring an Oracle instance on Windows computers	1319
	Viewing an Oracle instance on Windows computers	1321
	Editing an Oracle instance on Windows computers	1322
	Deleting an Oracle instance on Windows computers	1323
	Enabling database access for Oracle operations on Windows computers	1323
	Configuring an Oracle instance on Linux servers	1324
	Viewing an Oracle instance on Linux servers	1326
	Editing an Oracle instance on Linux servers	1326
	Deleting an Oracle instance on Linux servers	1327
	Enabling database access for Oracle operations on Linux servers	1327
	About authentication credentials on the media server	1329
	Setting authentication credentials on the media server for Oracle operations	1330
	Editing authentication credentials on the media server for Oracle operations	1332
	Deleting an Oracle server from the media server's list of authentication credentials	1333
	About Oracle instance information changes	1333
	Setting application defaults for Oracle	1334
	Oracle default options	1334
	About backing up Oracle resources	1335
	About backing up Oracle RAC resources	1337
	Backing up Oracle resources	1337
	About performing a DBA-initiated backup job for Oracle	1340

About restoring and recovering Oracle resources	1341
About DBA-initiated restore and recovery for Oracle	1343
Restoring Oracle data	1343
About redirecting a restore of Oracle data	1347
Restoring from a legacy GRFS Oracle Agent database backup	1348
Requirements for recovering the complete Oracle instance and database using the original Oracle server	1349
Recovering the complete Oracle instance and database using the original Oracle server	1350
Requirements for recovering the complete Oracle instance or database to a computer other than the original Oracle server	1351
Recovering the complete Oracle instance or database to a computer other than the original Oracle server	1352
Troubleshooting the Oracle Agent	1353
Changing the SqlplusTimeout for Oracle instances on Windows computers	1356
Changing the SqlplusTimeout for Oracle instances on Linux computers	1357
Changing the time-out for an automatic RMAN channel for Oracle instances on Windows computers	1357
Changing the time-out for an automatic RMAN channel for Oracle instances on Linux computers	1358
Updating the online redo log file path	1358

Appendix L

Symantec Backup Exec Agent for SAP Applications	1361
About the SAP Agent	1362
How the SAP Agent works	1362
About using the SAP Agent with RMAN	1363
Requirements for using the SAP Agent	1364
About installing the SAP Agent	1365
About SAP Agent security and privileges	1365
About encrypting SAP data	1366
About generating SAP Agent alerts	1366
About preserving the integrity of the SAP Agent catalog	1366
Before backing up SAP data	1367
Configuring biparam.ini for the SAP Agent	1368
Configuring DBA-initiated job settings for SAP	1369
About system level SAP backup jobs	1371
About backing up and restoring with the SAP Agent	1371

	Requirements for submitting jobs from remote computers by using the SAP Agent	1372
	Restoring data with BRRESTORE and the SAP Agent	1372
	About redirecting SAP restore jobs	1373
	Backing up SAP data with RMAN	1373
	Restoring SAP data with RMAN	1375
	Migrating the SAP Agent catalog from _backint.mdb to _backint.xml	1376
	About backing up a clustered SAP database on Microsoft Cluster Server	1377
	About backing up MaxDB databases by using the SAP Agent	1378
	Preparing MaxDB databases for backup	1379
	Backing up MaxDB databases	1379
	Restoring MaxDB databases by using the SAP Agent	1380
	About performing disaster recovery using the SAP Agent	1380
	SAP disaster recovery prerequisites	1381
	Recovering a remote SAP database server from a disaster	1381
	Recovering a combination SAP database server and media server	1382
Appendix M	Symantec Backup Exec Agent for VMware Virtual Infrastructure	1385
	About the Agent for VMware	1386
	Requirements for using the Agent for VMware	1386
	About installing the Agent for VMware	1387
	Adding VMware vCenter and ESX servers	1387
	Deleting VMware vCenter and ESX servers	1388
	About backing up VMware resources	1388
	How Backup Exec automatically protects new virtual machines during a backup job	1390
	Creating a full backup of VMware resources	1390
	VMware backup options	1391
	Virtual Machine Application Granular Recovery Technology Settings	1394
	Creating an incremental or a differential backup of VMware resources	1395
	How Granular Recovery Technology works with the Agent for VMware	1396
	How Backup Exec protects Exchange, SQL, and Active Directory data on virtual machines	1397
	Requirements for protecting Exchange, SQL, and Active Directory data on virtual machines	1398

About protecting databases and applications with the Symantec VSS	
Provider	1398
Changing the log truncation setting of the Symantec VSS	
Provider	1399
About restoring VMware resources	1400
About selecting VMware resources for restore	1400
Restoring VMware resources	1401
Redirecting the restore of a VMware virtual machine	1404
VMware Redirection options	1404
Setting default backup and restore options for the Agent for	
VMware	1406
VMware default options	1407
Appendix N	
Symantec Backup Exec Archiving Option	1413
About the Archiving Option	1414
Requirements for both the Exchange Mailbox Archiving Option and	
the File System Archiving Option	1415
Requirements for the Backup Exec Exchange Mailbox Archiving	
Option	1416
Requirements for the Backup Exec File System Archiving	
Option	1418
About granting permissions on the Exchange Server for the	
Backup Exec service account in the Archiving Option	1419
About Enterprise Vault services for the Archiving Option	1422
How to calculate disk space requirements for the Exchange	
Mailbox Archiving Option	1423
How to calculate disk space requirements for the File System	
Archiving Option	1425
Installing the Backup Exec Archiving Option	1428
About uninstalling or reinstalling the Archiving Option	1428
About installing Enterprise Vault on a media server on which	
the Archiving Option is installed	1429
How the Archiving Option works	1429
Types of data not included in Archiving Option archive jobs	1431
About Archiving Option operation entries in the audit log	1432
How Archiving Option end users retrieve archived data by using	
Backup Exec Retrieve	1432
Best practices for the Archiving Option	1433
About creating an Archiving Option archive job	1434
Creating an Archiving Option archive job by setting job	
properties	1435
About vault stores in the Archiving Option	1445

Creating a vault store in the Archiving Option	1446
New vault store options	1446
Editing or viewing vault store properties	1447
Vault store properties	1448
Vault store selections	1449
About deleting an Archiving Option vault store	1450
About vault store partitions in the Archiving Option	1451
Creating a vault store partition	1451
Editing vault store partition properties	1452
Vault store partition properties	1452
About archives in the Archiving Option	1453
Editing archive properties	1454
Archive properties	1454
Deleting an archive	1455
About archive settings in the Archiving Option	1455
Archive settings options	1456
About retention categories for archived items	1457
Applying different archive settings to file system share and folder selections for archive jobs	1459
Include/Exclude Selections options for archive jobs	1460
About Exchange mailbox groups in archive jobs	1461
Mailbox group options	1462
Managing Exchange mailbox groups	1463
About searching for data in the archives	1464
Searching for data in the archives	1465
Search Archives options	1465
Search Archives options for file system selections	1466
Search archives options for Exchange selections	1467
About restoring items from the archives	1468
Restoring items from archives	1468
Selections options for restoring items from archives	1470
General options to restore items from archives	1470
Microsoft Exchange options to restore items from archives	1471
File redirection options to restore items from archives	1472
Microsoft Exchange redirection options to restore items from archives	1473
About deleting items from the archives	1474
Deleting items from the archives	1475
About deleting archived data from its original location	1477
About backing up Archiving Option components	1477
About consistency checks for Archiving Option databases	1480
About disabling backup mode for Archiving Option components	1480

Backing up Archiving Option components	1481
Backup job properties for archive jobs	1482
About restoring an Archiving Option component	1483
Restoring an Archiving Option component	1483
About backing up and restoring the Archiving Option components	
from a remote media server	1491
Editing backup job default settings for Archiving Option	
components from a remote media server	1492
Backup job default settings for the Archiving Option	1492
Preventing the deletion of expired archived items from an	
archive	1493
About synchronizing archive permissions and settings	1493
About single instance storage of archived items	1494
Enabling single instance storage of archived items	1494
Editing default settings for archive jobs	1494
Archive job default settings	1495
About moving Archiving Option components to a new location	1498
Troubleshooting archive jobs	1499
Viewing the Enterprise Vault event log for Archiving Option	
events	1500
Reports for the Archiving Option	1500

Appendix O

Symantec Backup Exec Central Admin Server	
Option	1503
How CASO works	1504
How CASO and the Shared Storage Option work together	1508
Requirements for installing CASO	1508
How to choose the location for CASO device and media data	1509
Installing the CASO central administration server	1512
Installing a managed media server from the central administration	
server in CASO	1513
About installing a CASO managed media server across a firewall	
.....	1518
Changing the dynamic port on the SQL Express instance in CASO	
to a static port	1518
Creating an alias for a managed media server when a SQL Express	
instance is used	1519
Opening a SQL port in CASO for a SQL 2005 or 2008	
instance	1520
Creating an alias for a managed media server when a SQL 2005	
or SQL 2008 instance is used	1521
About upgrading an existing CASO installation	1521

Upgrading an existing CASO central administration server	1522
Upgrading an existing CASO managed media server	1523
Changing a Backup Exec media server to a central administration server	1525
Changing a media server to a managed media server	1526
Changing a managed media server to a stand-alone media server	1527
Running the Backup Exec Utility for CASO operations	1527
Uninstalling Backup Exec from the central administration server in CASO	1528
Uninstalling Backup Exec from a managed media server	1528
About configuring CASO	1529
About reducing network traffic in CASO	1530
Setting defaults for managed media servers	1531
Setting communication thresholds and active job status updates for CASO	1533
What happens when CASO communication thresholds are reached	1536
Copying logs and histories to the central administration server	1536
How alerts work in CASO	1538
About alerts and notification in CASO	1540
Enabling managed media servers to use any available network interface card	1540
About CASO catalog locations	1541
Changing the CASO catalog location	1542
About job delegation in CASO	1544
How to use media server pools in CASO	1545
Restricting the backup of a selection list to specific devices in CASO	1546
Creating a media server pool in CASO	1547
Adding managed media servers to a media server pool in CASO	1547
Renaming a media server pool in CASO	1548
Deleting a media server pool in CASO	1548
Removing a managed media server from a media server pool in CASO	1548
Viewing general properties for a media server pool in CASO	1549
Viewing active job and alert statistics for a media server pool in CASO	1549
Applying settings to all managed media servers in a pool in CASO	1550
About copying jobs instead of delegating jobs in CASO	1551

Requirements for duplicate backup data and synthetic backup jobs in CASO	1551
How centralized restore works in CASO	1552
How CASO restores data that resides on multiple devices	1553
Best practices for centralized restore in CASO	1555
Restoring from the CASO central administration server	1555
Media Servers view in CASO	1557
About managing jobs in CASO	1559
About recovering failed jobs in CASO	1560
Pausing a managed media server in CASO	1561
Resuming a paused managed media server in CASO	1562
How paused storage devices appear on the Devices view in CASO	1562
Disabling communications in CASO	1562
Enabling communications in CASO	1563
Stopping Backup Exec services for CASO	1563
Starting Backup Exec services for CASO	1563
Connecting to a remote managed media server	1564
Viewing managed media server properties	1564
Disaster Recovery in CASO	1565
 Appendix P	
Symantec Backup Exec Deduplication Option	1567
About the Deduplication Option	1568
Deduplication methods for Backup Exec agents	1570
Requirements for the Deduplication Option	1572
About installing the Deduplication Option	1573
About OpenStorage devices	1574
Adding an OpenStorage device	1574
Viewing properties for OpenStorage devices	1578
About deduplication storage folders	1579
Adding a deduplication storage folder	1580
Deleting a deduplication storage folder	1584
Viewing properties of a deduplication storage folder	1584
Changing the password for the logon account for a deduplication storage folder	1588
Sharing a deduplication device between multiple media servers	1589
About client-side deduplication	1590
Configuring client-side deduplication	1591
Configuring a Remote Agent for Deduplication	1592
Viewing properties of a Remote Agent for Deduplication	1594
About backup jobs for deduplication	1595
About optimized duplication	1595

Setting up optimized duplication	1596
About copying deduplicated data to tapes	1597
About using deduplication with encryption	1597
About restoring deduplicated data	1597
About disaster recovery of a deduplication storage folder	1597
Preparing for disaster recovery of a deduplication storage folder	1598
Performing disaster recovery of a deduplication storage folder	1598
About disaster recovery of OpenStorage devices	1599

Appendix Q

Symantec Backup Exec Desktop and Laptop Option	1601
About the Desktop and Laptop Option	1603
About the components of DLO	1604
Before you install DLO	1605
System requirements for the DLO Administration Console	1609
About installing the Backup Exec Desktop and Laptop Option	1610
How to deploy the Desktop Agent	1611
Customizing the Desktop Agent installation	1612
Preparing for a manual push-deployment of the Desktop Agent	1614
About setting a recovery password	1616
Checking data integrity	1616
Data Integrity Scanner Options	1617
Changing DLO service credentials	1618
Service Account Information options	1618
About administrator accounts in DLO	1618
Administrator Account Management options	1619
Adding an administrator account	1620
Editing an administrator account	1621
Removing an administrator account	1621
About automated permissions management in DLO	1622
About limited restore in DLO	1622
Using a list of individual accounts to manage DLO permissions	1623
Using domain groups to manage DLO permissions	1623
Permissions options	1624
About default DLO settings	1625
Changing default DLO profile settings	1625
Changing default DLO backup selection settings	1626
Changing default DLO global settings	1626
Global Settings options	1627

Desktop Agent Interval options	1628
User Activity Settings options	1630
LiveUpdate options	1630
Configuring DLO to use a specific port for database access	1631
About using Backup Exec Retrieve with DLO	1632
About updating DLO	1632
Updating the DLO Administration Console	1633
Updating the Desktop Agent	1633
Performing a silent upgrade of the Desktop Agent	1635
About upgrading DLO to Windows Vista	1636
Upgrading From NetBackup Professional to DLO	1636
Starting the DLO Administration Console from Backup Exec	1638
About the DLO Overview view	1638
Connecting to DLO on a different Backup Exec Media Server	1639
Connect to Media Server options for DLO	1640
How to configure DLO	1640
Starting the Configuration Wizard	1641
About DLO profiles	1641
Creating a new DLO profile	1642
Copying a DLO profile	1657
Modifying a DLO profile	1657
About backup selections in DLO	1658
About default backup selections in DLO	1659
Removing default DLO backup selections from a profile	1660
Adding a DLO backup selection to a profile	1660
General options for DLO backup selections	1661
Including and excluding files or folders from a DLO backup selection	1662
About revision control in DLO	1663
About file grooming in DLO	1664
Revision Control options for DLO backup selections	1664
Setting options for a DLO backup selection	1666
How to use DLO macros in backup selections	1667
Modifying a DLO backup selection	1670
Deleting DLO backup selections	1670
About Delta File Transfer	1671
Requirements for Delta File Transfer	1671
Maintenance Server technical information and tips	1671
How to enable Delta File Transfer for a backup selection	1672
Adding a new Maintenance Server	1672
Configuring a maintenance server for delegation	1673
Confirming that the desktop user's account is configured for delegation	1674

Confirming that the server process account is trusted for delegation	1674
Changing the default maintenance server	1675
Reassigning a file server	1675
About DLO Storage Locations	1676
Supported Storage Location configurations	1676
How to use hidden shares as Storage Locations	1677
Creating DLO Storage Locations	1678
Configuring a remote Windows share or NAS device for DLO Storage Locations	1680
Configuring a remote Windows share or NAS device for DLO Storage Locations using non-administrator case	1681
Deleting DLO Storage Locations from a remote Windows share or NAS device	1682
Deleting DLO Storage Locations	1682
About Automated User Assignments	1683
Creating Automated User Assignments	1684
Modifying Automated User Assignments	1686
Changing the priority of Automated User Assignments	1686
Viewing Automated User Assignment properties	1686
Deleting Automated User Assignments	1687
About configuring global exclude filters in DLO	1687
Specifying files and folders to exclude from all DLO backups	1688
Excluding email from all DLO backups	1690
Excluding files and folders from compression	1692
Excluding files and folders from encryption	1693
Excluding files and folders from Delta File Transfer	1694
About excluding files that are always open	1695
About using DLO macros to define global excludes	1696
About managing Desktop Agent users	1696
Manually creating new network user data folders	1697
Adding a single desktop user to DLO	1697
Importing multiple desktop users who have existing network storage	1699
Changing the profile for a Desktop Agent user	1699
Enabling or disabling DLO access for a desktop user	1700
Deleting a user from DLO	1700
Moving Desktop Agent users to a new network user data folder	1701
Migrating a desktop user to a new computer	1702
Viewing a list of Desktop Agent users	1703
Modifying computer properties	1703

Enabling or disabling a desktop computer	1704
Deleting a desktop computer from DLO	1704
Backing up a desktop from the DLO Administration Console	1705
Setting blackout windows	1705
Deleting a blackout window schedule	1706
Restoring files and folders from the DLO Administration	
Console	1707
Restore options	1708
Restore Summary options	1710
Searching for files and folders to restore with DLO	1710
Restore search options	1711
About DLO emergency restore and recovery passwords	1711
About changing recovery passwords	1712
What happens when a user is deleted by the DLO Administration	
Console	1712
Recovering data for a single user by using DLO Emergency	
Restore	1713
Recovering data for a single user without using DLO Emergency	
Restore	1713
Recovering a media server or a file server if a non-system disk	
fails or is otherwise corrupted	1714
Recovering a media server if the hard drive fails or the computer	
needs to be replaced	1714
Recovering a file server if the hard drive fails or the computer	
needs to be replaced	1714
Computer History pane options and Job History pane options	1715
Viewing history logs	1717
Setting filters for the job history view	1718
Searching history logs	1719
About monitoring alerts on the DLO Administration Console	1720
Alert categories	1721
DLO informational alerts	1721
DLO warnings	1722
DLO alerts	1723
Configuring alerts	1724
Managing DLO alerts	1725
Clearing DLO alerts	1726
About configuring notification methods for DLO alerts	1727
Configuring notification methods for DLO alerts	1728
About configuring recipients for notification in DLO	1728
Enabling a person to receive DLO alert notifications by SMTP	
mail	1729

Enabling a person to receive DLO alert notifications by MAPI mail	1729
Enabling a person to receive DLO alert notifications by VIM mail	1730
Enabling a person to receive DLO alert notifications by pager	1730
Enabling SNMP Trap to receive DLO alert notifications	1730
Enabling Net Send to receive DLO alert notifications	1731
Enabling a printer to receive DLO alert notifications	1731
Enabling a group to receive DLO alert notifications	1732
Scheduling notification for recipients in DLO	1732
Changing information about a recipient in DLO	1733
Changing the notification method for a recipient in DLO	1733
Removing recipients for DLO alerts	1734
About DLO reports	1734
Running a DLO report	1736
Viewing DLO report properties	1736
About maintaining the DLO database	1737
About clustering the Desktop and Laptop Option	1738
Installing Backup Exec and the Desktop and Laptop Option to an existing cluster	1738
Upgrading an existing Backup Exec 9.x or 10.x cluster that includes DLO	1739
Upgrading an existing Backup Exec 9.x or 10.x cluster and adding DLO to the cluster	1739
Reconnecting a Desktop Agent to a cluster node after you uncluster DLO	1740
Moving a Storage Location in a DLO cluster environment before taking DLO out of the cluster	1740
About the DLO command syntax	1740
About remote server options for the command line	1741
DLO commands in detail	1741
About the -AssignSL Command	1742
About the -EnableUser Command	1743
About the -ChangeServer Command	1744
About the -KeyTest Command	1745
About the -ListProfile Command	1746
About the -ListSL Command	1747
About the -ListUser Command	1748
About the -LogFile Command	1748
About the -Update Command	1749
About the -EmergencyRestore Command	1752
About the -SetRecoveryPwd Command	1752

About the -NotifyClients Command	1752
About the -InactiveAccounts Command	1753
About the -RenameDomain Command	1753
About the -RenameMS Command	1753
About the -LimitAdminTo Command	1754
About the -IOProfile Command	1754
About the Desktop Agent	1755
Desktop Agent terminology	1756
Features and benefits of the Desktop Agent	1756
System requirements for the Desktop Agent	1757
Installing the Desktop Agent	1758
How to configure the Desktop Agent	1759
About connecting from the Desktop Agent to the media server	1759
Alternate Credentials options	1760
About using local accounts on desktop computers	1761
Resetting dialog boxes and account information in DLO	1762
Changing your connection status	1762
Enabling the Desktop Agent	1763
Disabling the Desktop Agent	1763
About the Desktop Agent Console	1763
About using the Desktop Agent to back up your data	1765
About revisions	1766
Modifying backup selections in the Desktop Agent's standard view	1767
Adding backup selections in the Desktop Agent's advanced view	1768
Modifying backup selections in the Desktop Agent's advanced view	1768
Deleting backup selections in the Desktop Agent's advanced view	1769
About using DLO to back up Outlook PST files incrementally	1769
About backing up Lotus Notes NSF files incrementally	1770
Configuring the Desktop Agent for incremental backup of Lotus Notes files	1771
About using the Desktop Agent when Lotus Notes is not configured for the current user	1772
About modifying Desktop Agent settings	1772
Changing schedule options for a DLO backup job	1773
Setting customized options on the Desktop Agent	1775
Moving the desktop user data folder	1777
Customizing connection policies	1777

About synchronizing desktop user data	1778
How synchronization works	1779
Synchronizing a folder across multiple desktops	1780
Changing or viewing a synchronized folder	1781
Removing a synchronized folder	1781
Resolving conflicts with synchronized files	1781
About the status of the Desktop Agent	1782
Starting a pending job from the Status view	1783
About suspending or canceling a job	1783
Viewing usage details	1783
Usage Details	1784
Restoring files by using the Desktop Agent	1786
Restore options	1787
Searching for desktop files and folders to restore	1788
About restoring Microsoft Outlook Personal Folder files	1789
About restoring deleted email messages	1789
About restoring files with alternate stream data	1790
About using Backup Exec Retrieve to restore files	1790
About monitoring job history in the Desktop Agent	1790
Viewing log files	1791
Searching for log files	1793
About grooming log files	1795
About using DLO with other products	1795
Troubleshooting the DLO Administration Console	1796
Troubleshooting the Desktop Agent	1800
Accessibility and DLO	1802

Appendix R

Symantec Backup Exec Intelligent Disaster Recovery Option	1805
About the Intelligent Disaster Recovery Option	1806
Requirements for using IDR	1807
About installing the IDR Option	1807
About using a trial version of the IDR Option	1808
About preparing computers for IDR	1808
About the the Intelligent Disaster Recovery Configuration Wizard	1810
About manually editing the default data paths for the *.dr files	1811
About creating and updating recovery media	1813
About requirements for running the Intelligent Disaster Recovery Preparation Wizard	1814

About running the Intelligent Disaster Recovery Preparation Wizard	1815
About creating recovery media after a disaster	1816
Creating the Intelligent Disaster Recovery nonbootable CD image only	1822
Copying the disaster recovery files	1824
Preparing IDR media by using other media servers	1825
Media server logon credential options	1826
About preparing to recover from a disaster by using IDR	1827
About changing hardware in the computer to be recovered	1829
About using IDR to recover IBM computers	1830
About the Intelligent Disaster Recovery Wizard	1830
About encrypted backup sets and the Intelligent Disaster Recovery Wizard	1831
Recovering a computer by using the Intelligent Disaster Recovery Wizard	1831
Performing an automated restore by using the Intelligent Disaster Recovery Wizard	1832
Restoring from a locally attached media device	1835
Restoring from remote backup-to-disk folders	1837
Restoring from a remote media server	1839
Installing network drivers	1840
About altering hard drive partition sizes	1841
Performing a manual restore by using the Intelligent Disaster Recovery Wizard	1841
Microsoft SQL Server recovery notes	1844
Microsoft Exchange recovery notes	1844
SharePoint Portal Server recovery notes	1844
Citrix Metaframe recovery notes	1844
About using IDR with the Central Admin Server Option	1845
About using IDR with Veritas Storage Foundation for Windows	1845
Best Practices for IDR	1846
Appendix S	
Symantec Backup Exec NDMP Option	1847
About the NDMP Option	1847
Requirements for using the NDMP Option	1848
About installing the NDMP Option	1849
Adding an NDMP server to Backup Exec	1849
Add NDMP Server options	1850
Sharing the devices on an NDMP server between multiple media servers	1850
Backing up NDMP resources	1851

NDMP backup options for NetApp/IBM/Fujitsu	1852
NDMP backup options for EMC	1853
NDMP backup options for providers other than NetApp/IBM/Fujitsu and EMC	1854
About including and excluding directories and files for NDMP backup selections	1854
Including specific directories in an NDMP backup selection	1856
How to use patterns to exclude files and directories from an NDMP backup selection	1857
Excluding directories and files from a NetApp/IBM/Fujitsu backup selection	1859
Excluding directories and files from an EMC backup selection	1860
Excluding files and directories from backup selections for NDMP providers other than NetApp/IBM/Fujitsu and EMC	1861
How to duplicate backed up NDMP data	1862
Restoring NDMP data	1863
NDMP restore options for NetApp/IBM/Fujitsu	1864
NDMP restore options for EMC	1865
NDMP restore options for providers other than NetApp/IBM/Fujitsu and EMC	1866
About redirecting restored NDMP data	1867
Setting the default backup and restore options for NDMP	1867
NDMP default options for NetApp/IBM/Fujitsu backup and restore	1868
NDMP default options for EMC backup and restore	1869
NDMP default backup and restore options for providers other than NetApp/IBM/Fujitsu and EMC	1871
Viewing NDMP server properties	1872
NDMP server properties	1872
Appendix T Symantec Backup Exec Remote Agent for Linux or UNIX Servers	1875
About the Remote Agent for Linux or UNIX Servers	1876
Requirements for the Remote Agent for Linux or UNIX Servers	1876
About installing the Remote Agent for Linux or UNIX Servers	1877
Installing the Remote Agent for Linux or UNIX Servers	1877
About configuring the Remote Agent for Linux or UNIX Servers	1881
About publishing Linux, UNIX, and Macintosh computers to media servers	1882
Adding media servers to which the Remote Agent for Linux, UNIX, and Macintosh can publish information	1883

About excluding files and directories from backup jobs for Linux, UNIX, and Macintosh computers	1884
Editing configuration options for Linux, UNIX, and Macintosh computers	1884
Configuration options for Linux, UNIX, and Macintosh computers	1885
About backing up data by using the Remote Agent for Linux or UNIX Servers	1891
Backing up Linux, UNIX, and Macintosh computers	1892
Backup job options for Linux, UNIX, and Macintosh computers	1918
Requirements for backing up Novell Open Enterprise Server on SUSE Linux Enterprise Server	1895
Novell Open Enterprise Server components that are supported for backup	1896
Backing up Novell Open Enterprise Server (OES) components	1896
Restoring data to Linux, UNIX, and Macintosh computers	1897
About restoring Novell OES components	1898
Restore job options for Linux, UNIX, and Macintosh computers	1898
Edit the default backup and restore job options for Linux, UNIX, and Macintosh computers	1899
Default backup and restore job options for Linux, UNIX, and Macintosh computers	1899
Uninstalling the Remote Agent for Linux or UNIX Servers	1903
Manually uninstalling the Remote Agent for Linux or UNIX Servers	1904
Runtime scripts to remove when manually uninstalling the Remote Agent for Linux or UNIX Servers	1905
Starting the Remote Agent for Linux or UNIX Servers daemon	1907
Stopping the Remote Agent for Linux or UNIX Servers daemon	1907
Troubleshooting the Remote Agent for Linux or UNIX Servers	1908
Appendix U	
Symantec Backup Exec Remote Agent for Macintosh Systems	1911
About the Remote Agent for Macintosh Systems	1911
Requirements for the Remote Agent for Macintosh Systems	1912
About the Backup Exec admin group on Macintosh systems	1912
Creating the Backup Exec admin group manually on Macintosh systems	1913
About installing the Remote Agent for Macintosh Systems	1914

Installing the Remote Agent for Macintosh Systems	1914
About configuring the Remote Agent for Macintosh Systems	1917
About backing up data by using the Remote Agent for Macintosh Systems	1917
Backing up Macintosh systems	1918
Macintosh restore options	1918
Restoring Macintosh systems	1919
Editing the default backup and restore options for Macintosh systems	1919
Default backup and restore job options for Macintosh systems	1919
Uninstalling the Remote Agent for Macintosh Systems	1923
Starting the Remote Agent for Macintosh Systems	1924
Stopping the Remote Agent for Macintosh Systems	1924
Manually uninstalling the Remote Agent for Macintosh Systems	1925
Troubleshooting the Remote Agent for Macintosh Systems	1926

Appendix V

Symantec Backup Exec Remote Agent for NetWare Systems	1929
About the Remote Agent for NetWare Systems	1929
Requirements for installing the Remote Agent for NetWare Systems on a NetWare server	1930
About installing the Remote Agent for NetWare Systems	1931
Installing the Remote Agent for NetWare Systems	1931
About publishing NetWare servers to the NetWare agents list	1933
Adding BESTART to the Autoexec.ncf file on the NetWare server	1934
Unloading the Remote Agent for NetWare Systems	1934
About backing up NetWare servers	1934
About backing up the NetWare Directory Services (NDS)	1936
Backing up NetWare servers	1938
NetWare SMS backup options	1939
About restoring NetWare servers	1939
Restoring NetWare servers	1939
About default options for the Remote Agent for NetWare Systems	1940
Setting default options for the Remote Agent for NetWare Systems	1941
Specifying TCP dynamic port ranges on the media server	1943
Saving configuration information for the NetWare server	1943

Appendix W	Symantec Backup Exec Remote Agent for Windows Systems	1945
	About the Remote Agent for Windows Systems	1945
	Requirements for the Remote Agent for Windows Systems	1946
	Stopping and starting the Remote Agent for Windows Systems	1947
	About the Remote Agent Utility for Windows Systems	1948
	Starting the Remote Agent Utility	1949
	Viewing the activity status of the remote computer in the Remote Agent Utility	1949
	Status options for the Remote Agent Utility	1950
	Viewing the activity status of the remote computer from the system tray	1950
	Starting the Remote Agent Utility automatically on the remote computer	1951
	Setting the refresh interval on the remote computer	1951
	About publishing the Remote Agent for Windows Systems to media servers	1951
	Configuring database access	1955
	Database access options for the Remote Agent Utility	1956
	About the Remote Agent Utility Command Line Applet	1958
	Using the Remote Agent Utility Command Line Applet	1958
	Remote Agent Utility Command Line Applet switches	1959
Appendix X	Symantec Backup Exec Remote Media Agent for Linux Servers	1965
	About the Remote Media Agent for Linux Servers	1966
	How the Remote Media Agent for Linux Servers works	1966
	Requirements for the Remote Media Agent for Linux Servers	1967
	About installing the Remote Media Agent for Linux Servers	1968
	Installing the Remote Media Agent for Linux Servers	1968
	About the Backup Exec operators group for the Remote Media Agent for Linux Servers	1971
	Creating the Backup Exec operators group manually for the Remote Media Agent for Linux Servers	1971
	Adding a Linux server as a Remote Media Agent	1972
	Add Remote Media Agent options	1973
	Changing the port for communications between the media server and the Remote Media Agent	1975
	Editing properties for the Remote Media Agent for Linux Servers	1975
	Remote Media Agent properties	1976

Sharing a Remote Media Agent between multiple media servers	1977
About creating device pools for devices attached to the Remote Media Agent for Linux Servers	1977
Deleting a Remote Media Agent for Linux Servers from a media server	1978
Backing up data by using the Remote Media Agent for Linux Servers	1978
Restoring data by using the Remote Media Agent for Linux Servers	1978
About the Tape Library Simulator Utility	1979
Creating a simulated tape library	1980
Viewing simulated tape libraries properties	1981
Deleting a simulated tape library	1983
Managing simulated tape libraries from the command line	1984
Uninstalling the Remote Media Agent for Linux Servers	1985
Finding simulated tape library files	1986
Troubleshooting the Remote Media Agent for Linux Servers	1987

Appendix Y

Symantec Backup Exec SAN Shared Storage Option	1991
About the SAN Shared Storage Option	1991
Requirements for the SAN Shared Storage Option	1993
About installing the SAN Shared Storage Option	1994
About devices in the SAN Shared Storage Option	1995
About media rotation in the SAN Shared Storage Option	1996
How to catalog media in the SAN Shared Storage Option	1997
About sharing media in the SAN Shared Storage Option	1997
About scheduling and viewing jobs in the SAN Shared Storage Option	1997
About sharing robotic libraries between Backup Exec for NetWare Servers and Backup Exec	1998
About robotic library sharing prerequisites	1999
Configuring partitions on Windows media servers for robotic library sharing	2000
Configuring partitions on NetWare media servers for robotic library sharing	2000
About device operations with the SAN Shared Storage Option	2003
About renaming robotic libraries and drives in the SAN Shared Storage Option	2003
How to use drive pools with the SAN Shared Storage Option	2004

About viewing media in the SAN Shared Storage Option	2005
How to monitor drives in the SAN Shared Storage Option	2005
About designating a new primary database server and setting up servers in the SAN Shared Storage Option	2005
Tips for maintaining the Backup Exec database servers and the shared ADAMM database in the SAN Shared Storage Option	2006
Creating a standby primary database server in the SAN Shared Storage Option	2007
About starting and stopping Backup Exec Services on multiple servers in the SAN Shared Storage Option	2009
About reconfiguration of the SAN Shared Storage Option environment	2009
Troubleshooting failed components in the SAN Shared Storage Option	2010
Troubleshooting offline devices in the SAN Shared Storage Option	2010
Finding hardware errors for the SAN Shared Storage Option	2012
Resetting the SAN in the SAN Shared Storage Option	2013
Bringing devices online after an unsafe device removal event in the SAN Shared Storage Option	2014
Best practices for the SAN Shared Storage Option	2014
Appendix Z	
Symantec Backup Exec Storage Provisioning Option	2017
About the Storage Provisioning Option	2018
Requirements for the Storage Provisioning Option	2019
Requirements for the Storage Provisioning Option in a CASO environment	2019
About installing the Storage Provisioning Option	2020
Viewing storage array components in Backup Exec	2020
About using the Storage Array Configuration Wizard	2021
Configuring a storage array by using the Storage Array Configuration Wizard	2022
Viewing properties for storage arrays	2023
General properties for storage arrays	2023
Properties of physical disks on storage arrays	2026
About the All Virtual Disks device pool in the Storage Provisioning Option	2029
About virtual disks in the Storage Provisioning Option	2029
Editing default options for a virtual disk on a storage array	2030

Advanced properties for storage arrays	2031
Editing the default options for all virtual disks on storage arrays	2032
Default options for all virtual disks on storage arrays	2033
Configuring a virtual disk on a storage array	2034
Viewing properties for unconfigured virtual disks on a storage array	2035
Properties for unconfigured virtual disks on storage arrays	2035
Editing general properties of virtual disks on storage arrays	2037
General properties for virtual disks on storage arrays	2038
About hot spares in the Storage Provisioning Option	2042
Adding a hot spare by using the Storage Array Configuration Wizard	2042
Changing a hot spare by using the Storage Array Configuration Wizard	2043
Detecting a new storage array	2044
Renaming a virtual disk or storage array	2044
About identifying the physical disks of a virtual disk	2045
Identifying the physical disks of a virtual disk	2046
About predicting disk usage in the Storage Provisioning Option	2046
Configuring an alert for low disk space on storage arrays	2047
Default options for Storage Provisioning Alert	2047
Troubleshooting the Storage Provisioning Option	2048
Appendix AA Symantec Online Storage for Backup Exec	2051
About Symantec Online Storage for Backup Exec	2051
Best practices for using Symantec Online Storage for Backup Exec	2052
Setting up Symantec Online Storage for Backup Exec	2053
About signing up for Symantec Online Storage for Backup Exec	2053
About downloading the Symantec Online Storage for Backup Exec Protection Agent	2054
About Symantec Online Storage folders	2054
Creating a Symantec Online Storage folder	2054
Pausing a Symantec Online Storage folder	2056
Resuming a Symantec Online Storage folder	2056
Sharing an existing Symantec Online Storage folder	2057
About creating duplicate backup jobs for Symantec Online Storage for Backup Exec	2057

	Creating duplicate backup jobs for Symantec Online Storage for Backup Exec	2058
	About managing Symantec Online Storage for Backup Exec jobs	2060
	Erasing Symantec Online Storage for Backup Exec files	2060
	Deleting Symantec Online Storage folders	2061
	About restoring Symantec Online Storage for Backup Exec jobs	2062
Appendix AB	Accessibility and Backup Exec	2063
	About accessibility and Backup Exec	2063
	About keyboard shortcuts in Backup Exec	2064
	Keyboard shortcuts unique to Backup Exec	2064
	Keyboard shortcuts unique to Backup Exec Utility	2066
	Keyboard shortcuts unique to Backup Exec Desktop and Laptop Option Administration Console	2067
	Keyboard shortcuts unique to Backup Exec Desktop and Laptop Option Desktop Agent	2069
	General keyboard navigation within the Backup Exec user interface	2070
	Keyboard navigation within dialog boxes in Backup Exec	2071
	List box navigation in Backup Exec	2072
	Tabbed dialog box navigation in Backup Exec	2072
	About setting accessibility options	2072
Glossary	2075
Index	2083

Introducing Backup Exec

This chapter includes the following topics:

- [About Backup Exec](#)
- [How Backup Exec works](#)
- [What's new in Backup Exec](#)
- [What's new in Backup Exec agents and options](#)
- [Backup Exec agents and options](#)
- [About the Administration Console](#)
- [About the Home view](#)

About Backup Exec

Symantec Backup Exec 2010 is a high-performance data management solution for Windows® servers networks. With its client/server design, Backup Exec provides fast, reliable backup and restore capabilities for servers and workstations across the network.

Backup Exec is available in the following configurations that can accommodate multi-platform networks of all sizes.

Table 1-1 Backup Exec configurations for multiple platforms

Backup Exec Edition	Description
Symantec Backup Exec™ 2010	

Table 1-1 Backup Exec configurations for multiple platforms (*continued*)

Backup Exec Edition	Description
	<p>Supports a wide variety of both tape and disk devices in almost any type of storage configuration such as the following:</p> <ul style="list-style-type: none"> ■ Fiber Channel ■ iSCSI ■ NAS ■ SAN, LAN, and WAN ■ Disk-based deduplication appliances <p>Backup Exec 2010 protects physical and virtual environments such as Windows, Linux, Solaris, MAC OS, VMware, and NetWare systems. Optional Backup Exec agents are available to protect remote systems, applications, and databases. You can add separate Backup Exec options to provide advanced features such as data deduplication, archiving, and centralized management.</p> <p>Each license of Backup Exec 2010 includes the following options:</p> <ul style="list-style-type: none"> ■ Intelligent Disaster Recovery Option <p>This option provides disaster recovery capabilities for systems without having to reinstall the operating system. Backup Exec System Recovery 2010 is sold separately for faster and advanced system recovery capabilities. These capabilities include dissimilar hardware recovery support and recovery to a virtual environment such as VMware, Microsoft Hyper-V, and Citrix Xen. See the following URL: www.backupexec.com/besr</p> ■ Advanced Open File Option <p>This option is now included and enabled by default to provide automatic open file protection using the Microsoft Volume Shadow Copy Services (VSS) snapshot infrastructures.</p> ■ Backup Exec Desktop and Laptop Option <p>This option provides continuous protection of the user data files that are sent to a file share on your network that Backup Exec 2010 can protect. Licenses for five desktop and laptop computers are included in the Backup Exec 2010 license.</p> ■ Support for stand-alone tape drives and backup-to-disk folders

Table 1-1 Backup Exec configurations for multiple platforms (*continued*)

Backup Exec Edition	Description
	<p>Backup Exec 2010 provides support for a large number of tape and disk-based backup devices.</p> <p>You can find a list of compatible devices at the following URL: http://entsupport.symantec.com/umi/V-269-2</p> <ul style="list-style-type: none"> ■ Support for one drive in every physical robotic library and support for every single-drive virtual tape library <p>To enable support for each additional drive in a physical robotic library, you can purchase the Library Expansion Option. You can also purchase the Virtual Tape Library Unlimited Drive Option to enable unlimited drive support in a virtual tape library.</p> <ul style="list-style-type: none"> ■ Backup Exec Remote Agents for Windows XP Professional OS workstations. <p>See “Backup Exec agents and options” on page 80.</p>

Table 1-1 Backup Exec configurations for multiple platforms (*continued*)

Backup Exec Edition	Description
<p>Small Business Server Edition (SBSE)</p>	<p>Installs on and protects supported versions of Microsoft Small Business Server for Windows.</p> <p>You can find a list of supported operating systems, platforms, and applications at the following URL:</p> <p>http://entsupport.symantec.com/umi/V-269-1</p> <p>The Small Business Server Edition includes the following features:</p> <ul style="list-style-type: none"> ■ Exchange Agent. ■ SQL Agent to protect Microsoft SQL on Small Business Server Premium Editions. ■ SharePoint Agent to protect Windows SharePoint Services. ■ Remote Agent for Windows Systems to protect a second server with the Small Business Server Premium Edition. ■ Desktop and Laptop Option to protect supported versions of Windows desktops and laptops. ■ Backup Exec System Recovery 2010 Small Business Server Edition to take a disk-based snapshot backup of a complete system while Windows is running. Ensures a complete system recovery for the entire Small Business Server system to one of the following: <ul style="list-style-type: none"> ■ The original hardware. ■ To different hardware. ■ To a VMware or Hyper-V virtual environment. <p>The Granular Recovery Option is also included.</p> ■ Backup Exec System Recovery 2010 Server Edition (when used on Microsoft SBS Premium Edition only). Included to provide complete system recovery for the additional Windows server that is included with Microsoft Windows Small Business Server (SBS) Premium Edition. Includes the Granular Recovery Option. ■ Backup Exec System Recovery 2010 Desktop Edition. Required to use the Granular Recovery Option to facilitate the restore of individual Exchange mail messages and SharePoint documents. <p>With the exceptions of the Central Admin Server Option and SAN Shared Storage Option, you can purchase additional Backup Exec agents and options for use with SBSE.</p>

Table 1-1 Backup Exec configurations for multiple platforms (*continued*)

Backup Exec Edition	Description
QuickStart Edition (QSE) (OEM release only)	<p data-bbox="628 326 1193 378">Installs on and protects any supported version of Microsoft Windows Server.</p> <p data-bbox="628 401 1193 453">You can find a list of supported operating systems, platforms, and applications at the following URL:</p> <p data-bbox="628 470 1193 496">http://entsupport.symantec.com/umi/V-269-1</p> <p data-bbox="628 513 1193 539">QuickStart Edition supports the following:</p> <ul data-bbox="628 557 1193 852" style="list-style-type: none"> <li data-bbox="628 557 1193 635">■ Backup Exec Desktop and Laptop Option. This option provides continuous protection of user data files for five desktop and laptop computers. <li data-bbox="628 652 1193 817">■ Single-drive robotic library or virtual tape library. To enable support for each additional drive in a physical robotic library, you can purchase the Library Expansion Option. You can also purchase the Virtual Tape Library Unlimited Drive Option to enable unlimited drive support in a virtual tape library. <li data-bbox="628 835 1193 852">■ Stand-alone tape drives and backup-to-disk drives. <p data-bbox="628 878 1193 956">You must purchase an upgrade to Symantec Backup Exec 2010 or the Small Business Server Edition of Backup Exec to use additional Backup Exec agents and options.</p> <p data-bbox="628 965 1193 1017">OEM-specific versions of QuickStart may support additional options.</p>

See “What’s new in Backup Exec” on page 70.

See “What’s new in Backup Exec agents and options” on page 75.

See “Backup Exec agents and options” on page 80.

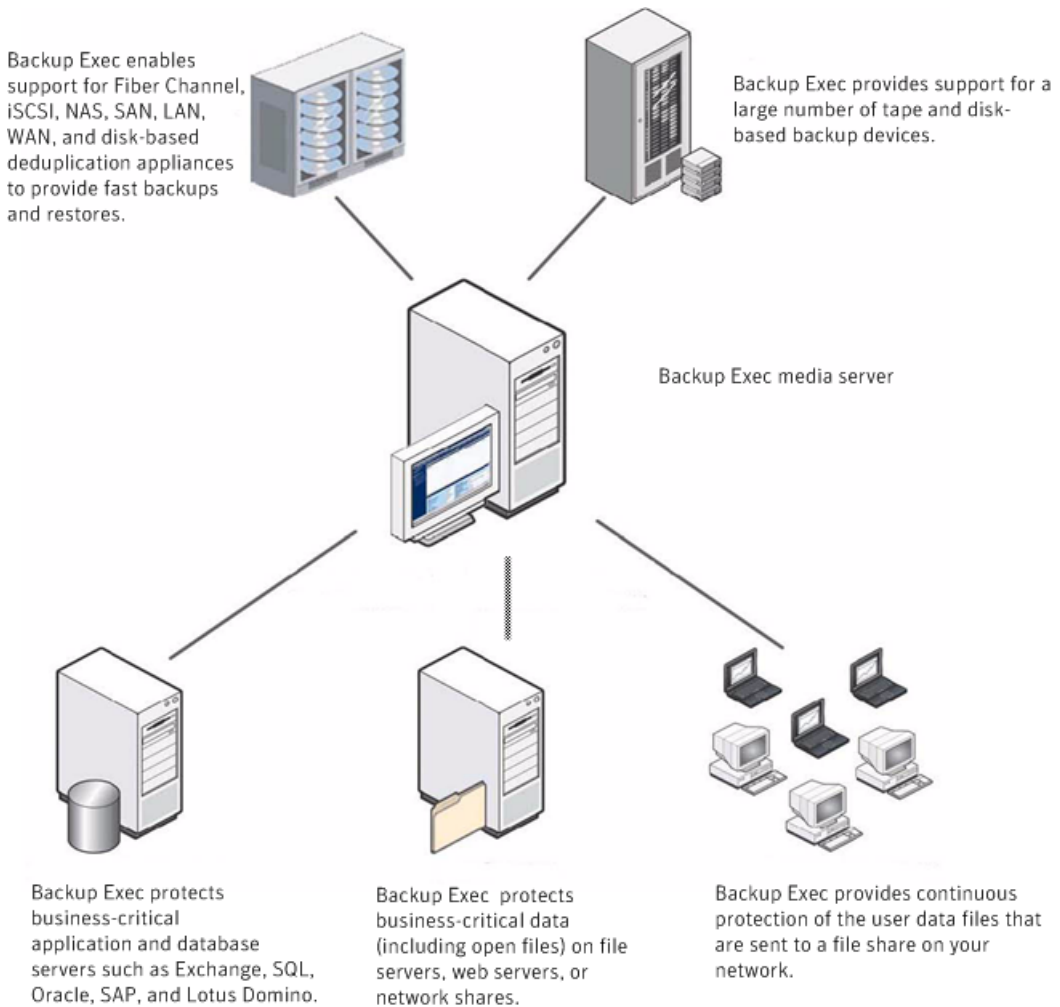
See “How Backup Exec works” on page 68.

How Backup Exec works

You use the Backup Exec Administration Console to submit backup, restore, and utility operations. Administrators can run the Administration Console from the media server (a Windows server on which Backup Exec is installed) or from a remote computer. After jobs are created, the Backup Exec media server processes the jobs or delegates the jobs for processing, depending on your environment.

Most interaction with Backup Exec, such as submitting jobs, viewing results, and performing device and media operations, is done through the Administration Console.

Figure 1-1 Backup and restore functionality for the entire network



Through the Administration Console, you configure the job defaults that you want Backup Exec to use for most jobs. However, you can override these default options while setting up a specific job such as a weekly backup of selected workstations, called resources. You can create a once-only job, such as the restore of a file to a server. Or, you can schedule recurring jobs, such as daily backup jobs. You can use policies to manage the recurring jobs that make up your backup strategy.

Wizards guide you through most Backup Exec operations, including the creation of a backup or restore job, setting up media rotation jobs, and setting media overwrite protection.

You can monitor a job's progress through the **Job Monitor**, or use Backup Exec's Calendar to quickly view all jobs scheduled to run for the day, week, or month.

The media server contains the media and device databases that organize and allocate the storage devices that are accessible to the media server. These databases also help prevent media from being accidentally overwritten. Through Backup Exec's device management functions, you can logically group storage devices together in device pools to share the backup workload. Through the media management function, you can organize, track, and troubleshoot all of the media in your library.

After a job has been processed, the job's results are stored in a job history database. A record of the data that was backed up is kept in Backup Exec's catalog. The job history is a report of what happened during the processing of the job (statistics, errors, and so on), and the catalog file is the record from which restore selections are made.

What's new in Backup Exec

This release of Backup Exec includes the following new features and capabilities:

Table 1-2 New features and capabilities in Backup Exec

New feature	Description
Support for Microsoft Windows Server 2008 R2	Lets you do the following: <ul style="list-style-type: none">■ Back up and restore data using the new Express (system state) writers■ Back up and restore operating system boot files from unnamed partitions■ Back up and restore Cluster Shared Volumes (CSV)■ Back up from and restore to native VHD files

Table 1-2 New features and capabilities in Backup Exec (*continued*)

New feature	Description
Support for the Server Core installation option of Windows Server 2008 R2	Lets you install the Backup Exec Remote Agent for Windows Systems on the Server Core for backup and restore operations. The Remote Agent also installs the Remote Agent Utility Command Line Applet. This applet lets you monitor Backup Exec operations on the remote computer.
Support for Microsoft Windows 7	Lets you install the Backup Exec Remote Agent for Microsoft Windows 7 computers for backup and restore operations. The Backup Exec Desktop and Laptop Option provides automated file protection for Windows 7 computers. Backup Exec also supports BitLocker drive encryption.
Enhanced Backup Exec License Assessment Tool	<p>Supports the license key scans on all installations of Backup Exec System Recovery and Backup Exec 2010 on your network.</p> <p>The License Assessment Tool report now provides the following new information:</p> <ul style="list-style-type: none"> ■ The versions of Backup Exec that are installed so that you can plan to upgrade your environment. ■ A web link on the report that provides upgrade assistance so that you can read about the new features in current releases.

Table 1-2 New features and capabilities in Backup Exec (*continued*)

New feature	Description
Support for new platforms for the Remote Agent for Linux or UNIX Servers	<p>Supports the following platforms:</p> <ul style="list-style-type: none"> ■ Oracle Enterprise Linux 5.2 ■ Ubuntu 8.10 ■ XenServer 5 ■ Debian 4.0, 5.0 ■ SUSE Linux Enterprise 11
Support for a new platform for the Remote Media Agent for Linux	Supports SUSE Linux Enterprise 11.
Enhancement for Library Expansion Option	Supports each additional drive that you add after the first drive that you add in each robotic library. When you install Backup Exec, support for the first drive in every robotic library is included. The Library Expansion Option enables support for each additional drive in a robotic library.
Home view on the Backup Exec Administration Console	Lets you add or delete items to customize the display of your important Backup Exec features. In one view, you can display summaries of jobs, alerts, and devices, and the technical support sites that you want fast access to. You can add as many or as few items as you want.
Installation DVD	Provides all of the Backup Exec installation files on a single DVD.

Table 1-2 New features and capabilities in Backup Exec (*continued*)

New feature	Description
Share Your Ideas web link	Lets you suggest new ideas for Symantec Backup Exec by clicking this link that is located at the top of the Backup Exec Administration Console . After you have submitted your suggestions, other community members can vote or comment on the idea. The ideas with the most votes move to the top of the list. Symantec product managers review these ideas for possible features in future releases.
DirectCopy	Enables data to be copied from a virtual device directly to a physical device. The Backup Exec media server records information about the data in the catalog. Because the information about the copied data is in the catalog, you can restore data from either the virtual device or the physical device.

Table 1-2 New features and capabilities in Backup Exec (*continued*)

New feature	Description
Verify backup sets template	<p>Lets you run verify operations independent of backup and duplicate backup jobs. The verify backup sets template lets you schedule a verify operation to run at any time after a backup. For example, you can schedule the verify operation to run outside of your backup window if your network resources are scarce. The verify backup sets template also greatly enhances the benefits you receive from Backup Exec's Deduplication Option by letting you verify backup sets locally.</p> <p>Backup Exec includes an example policy that is preconfigured with a verify backup sets template.</p>
Typical installation	<p>Provides a new, simpler wizard for installation. This installation method is designed for small or uncomplicated environments. Backup Exec makes many of the installation decisions for you, based on common installation scenarios. If you prefer to set all of your installation options, you can use the Custom installation option, which provides the same installation wizard that was available in Backup Exec 2010.</p>

Table 1-2 New features and capabilities in Backup Exec (*continued*)

New feature	Description
Backup Exec Renewal Assistant	<p>Lets you access maintenance contract expiration information from within Backup Exec.</p> <p>The Renewal Assistant uses the contract expiration information to automatically set the Backup Exec alerts that remind you to renew the maintenance contracts before they expire. Reminder alerts are set at 30-day, 60-day, and 90-day intervals, based on the expiration date of the maintenance contract.</p>
Symantec RSS Reader	Lets you view and add Backup Exec and Symantec RSS feeds in the Home view on the Backup Exec Administration Console.
Enhanced Backup Wizard	Provides a more efficient design that makes it easier for you to back up your computers. The wizard helps you create a backup job that uses the recommended best practices.

See [“What's new in Backup Exec agents and options”](#) on page 75.

See [“Backup Exec agents and options”](#) on page 80.

What's new in Backup Exec agents and options

This release of Backup Exec includes the following new features and capabilities in the agents and options:

Table 1-3 New features and capabilities in Backup Exec agents and options

Agent or option	New feature
Agent for VMware Virtual Infrastructure	<p>Includes support for the following:</p> <ul style="list-style-type: none"> ■ VMware vSphere v4.1, which includes ESX/ESXi 4.1, vCenter 4.1, and vStorage APIs for Data Protection ■ Differential and incremental backups of virtual machines that are configured with hardware version 7 ■ Non-staged backups, which provide improved performance without the need for VCB proxy servers ■ SAN-based restores of virtual machines <p>Also, includes the following new features:</p> <ul style="list-style-type: none"> ■ Single-pass backups of VSS-aware applications that are installed on virtual machines, with the ability to recover individual application items ■ Dynamic inclusion, which automatically protects any virtual machines that you added since the last backup ■ Ability to automatically exclude from jobs any virtual machines that are turned off ■ Ability to redirect a virtual machine to a different virtual machine folder or resource pool ■ Expanded Job History now lists the total number of virtual machines in each backup ■ Multiple alternate backup transport types to ensure that backups complete successfully
Agent for Microsoft Hyper-V	<p>Includes support for the following:</p> <ul style="list-style-type: none"> ■ Hyper-V 2008 R2 ■ Clustered Hyper-V virtual machines, with automatic discovery of Highly Available virtual machines ■ Cluster shared volumes ■ Live Migration ■ Single-pass backups of VSS-aware applications that are installed on virtual machines, with the ability to recover individual application items <p>This agent was formerly known as the Agent for Microsoft Virtual Servers.</p>

Table 1-3 New features and capabilities in Backup Exec agents and options
(continued)

Agent or option	New feature
Agent for Microsoft Exchange Server	<p>Includes the following new features:</p> <ul style="list-style-type: none"> ■ Support for Exchange Server 2010 ■ Support for Exchange Server 2010 Database Availability Groups ■ Support for Exchange Server 2010 in Hyper-V and ESX 4.0 environments through the Remote Agent for Windows Systems ■ CPS support Exchange Server 2010 stand-alone servers
Agent for Enterprise Vault	<p>Includes the following new features:</p> <ul style="list-style-type: none"> ■ Support for Enterprise Vault 8.0 ■ Support for Enterprise Vault Compliance Accelerator and Discovery Accelerator ■ Supports Backup Exec Migrator for Enterprise Vault. Backup Exec Migrator enables the migration of archived Enterprise Vault data from Enterprise Vault servers to the tertiary storage systems that Backup Exec media servers manage.
Exchange Mailbox Archiving Option	<p>Lets you archive Microsoft Exchange Server emails that have been backed up. After data is archived, it is deleted from its source location, which reduces the amount of data on the Exchange Server.</p> <p>The Archiving Option uses Enterprise Vault technology to move archive data to disk-based vault stores. Only data that is already backed up is archived so that there is little affect on the Exchange server. End users can retrieve current and previous versions of files by browsing a web interface called Backup Exec Retrieve.</p> <p>See “About the Archiving Option” on page 1414.</p>
File System Archiving Option	<p>Lets you archive Windows NTFS data that has been backed up. After data is archived, it is deleted from its source location, which reduces the amount of data on the file server.</p> <p>The Archiving Option uses Enterprise Vault technology to move archive data to disk-based vault stores. Only data that is already backed up is archived so that there is little affect on the file system server. End users can retrieve current and previous versions of files by browsing a web interface called Backup Exec Retrieve.</p> <p>See “About the Archiving Option” on page 1414.</p>

Table 1-3 New features and capabilities in Backup Exec agents and options
(continued)

Agent or option	New feature
Virtual Tape Library Unlimited Drive Option	<p>Supports all additional drives after you add the first drive in each virtual tape library (VTL). You can purchase the Virtual Tape Library Unlimited Drive Option to support all of the additional drives in each virtual tape library. You do not have to purchase separate instances of the Virtual Tape Library Unlimited Drive Option for each virtual drive.</p> <p>Additional VTL enhancements include the following:</p> <ul style="list-style-type: none"> ■ VTL recognition capabilities ■ VTL device-specific menus to ensure proper operation ■ Support for synthetic full backups when used with the Advanced Disk-based Backup Option.
Deduplication Option	<p>Includes the following new features:</p> <ul style="list-style-type: none"> ■ Ability to reduce the amount of disk storage that is required for backups by storing only unique data. ■ Ability to reduce backup network usage by sending only unique data across the network. ■ Support for non-staged restores of granular data from a duplicate backup set that is copied between deduplication devices that support non-staged granular restores. <p>Also, includes the following revised terminology:</p> <ul style="list-style-type: none"> ■ Direct Access is replaced with client-side deduplication ■ Remote Agent for Direct Access is replaced with Remote Agent deduplication ■ Server-side deduplication is replaced with media server deduplication ■ Source-side deduplication is replaced with client-side deduplication ■ Target-side deduplication is replaced with appliance deduplication
Active Directory Recovery Agent	Supports Microsoft Windows Server 2008 Active Directory Domain Services objects.
Agent for Lotus Domino Server	Supports Lotus Domino version 8.5, including support for the Domino Attachment and Object Service (DAOS).

Table 1-3 New features and capabilities in Backup Exec agents and options
(continued)

Agent or option	New feature
Symantec Online Storage for Backup Exec	Provides more efficient online backups. Symantec Online Storage for Backup Exec now compares your backup selections to any existing backup data from previous duplicate backup jobs. Any data that is unchanged from previous duplicate backup jobs is skipped. If only a portion of a file has changed, only that portion is backed up. This enhancement can reduce the amount of time and bandwidth that is required to run recurring backup jobs.
NDMP Option	Enables you to add new NDMP providers from multiple companies.
Agent for Microsoft SharePoint	<p>Includes the following:</p> <ul style="list-style-type: none"> ■ Support for SharePoint 2010. ■ Support for SharePoint Foundation 2010. ■ Support for Backup Exec Granular Recovery Technology (GRT) to the same extent as SharePoint Server 2007. ■ Support for Backup Exec Granular Recovery Technology (GRT) of new SharePoint 2010 content types, including external lists and document sets. <p>Also, includes the following new features:</p> <ul style="list-style-type: none"> ■ Redirection of SharePoint databases to alternate SQL instances. ■ Ability to overwrite existing databases. ■ Dynamic inclusion of components such as Web applications and service applications that are added to existing SharePoint farms.
Desktop and Laptop Option	<p>Supports Microsoft Windows Server 2008 R2 on the following components:</p> <ul style="list-style-type: none"> ■ DLO Administration Console ■ DLO Administration service ■ Maintenance service <p>Supports Microsoft Windows 7 on the following components:</p> <ul style="list-style-type: none"> ■ Desktop Agent ■ Change Log Service ■ DLO Administration Console (to support remote administration)

Backup Exec agents and options

Several Backup Exec options are available to provide protection for your network.

Options are categorized as follows:

- Media server components
See [“About Backup Exec media server components”](#) on page 80.
- Server protection agents
See [“About Backup Exec server protection agents”](#) on page 81.
- Application protection agents
See [“About Backup Exec application protection agents”](#) on page 82.
- Virtual machine agents
See [“About Backup Exec's virtual machine agents”](#) on page 85.
- Client protection agents
See [“About Backup Exec client protection agents”](#) on page 86.
- Media server storage options
See [“About Backup Exec media server storage options”](#) on page 87.

About Backup Exec media server components

The following media server components allow greater control of backups and disaster recovery:

Table 1-4 Backup Exec media server components

Item	Description
Symantec Backup Exec Advanced Open File Option	Ensures that all files on your network are protected even if they are being used. Whether used alone or in combination with specific database agents, this option handles open files at the volume level and is seamlessly integrated with Backup Exec. You do not need to know which files are open ahead of time; just set a scheduled backup to use this option. See “About the Advanced Open File Option” on page 937.

Table 1-4 Backup Exec media server components (*continued*)

Item	Description
Intelligent Disaster Recovery (IDR)	<p>Provides a recovery solution for both local and remote Windows computers. This option eliminates the need to manually re-install the entire operating system after a computer failure. IDR lets you use diskettes, CD-R/CD-RW, or bootable tape to restore from your last complete backup set to get back online fast.</p> <p>See “About the the Intelligent Disaster Recovery Configuration Wizard” on page 1810.</p>

About Backup Exec server protection agents

The following options provide protection for remote Microsoft Windows servers, Novell NetWare servers, Linux and UNIX servers, and Macintosh systems on the network:

Table 1-5 Backup Exec server protection agents

Item	Description
Symantec Backup Exec Remote Agent for Windows	<p>Provides backup and restore of remote Windows computers.</p> <p>See “About the Remote Agent for Windows Systems” on page 1945.</p>
Symantec Backup Exec Remote Agent for Linux or UNIX Servers	<p>Lets Windows Server network administrators perform backup and restore operations on Linux and UNIX servers that are connected to the network. This agent must be running on these servers before backup or restore operations can be performed.</p> <p>See “Backing up Linux, UNIX, and Macintosh computers” on page 1892.</p>

Table 1-5 Backup Exec server protection agents (*continued*)

Item	Description
Symantec Backup Exec Remote Media Agent for Linux Servers	<p>Lets you back up data to and restore data from the following devices:</p> <ul style="list-style-type: none"> ■ Storage devices that are directly attached to a Linux server. ■ A folder on a hard disk on the Linux server. <p>See “About the Remote Media Agent for Linux Servers” on page 1966.</p>
Symantec Backup Exec Remote Agent for Macintosh Systems	<p>Enables Windows Server network administrators to perform backup and restore operations on Macintosh systems that are connected to the network.</p> <p>See “Backing up Macintosh systems” on page 1918.</p>
Symantec Backup Exec Remote Agent for NetWare Systems	<p>Provides backup and restore of remote NetWare resources.</p> <p>See “About backing up NetWare servers” on page 1934.</p>

About Backup Exec application protection agents

The following application protection agents provide non-disruptive protection for corporate email messaging, knowledge base, and mission-critical database applications:

Table 1-6 Backup Exec application protection agents

Item	Description
Symantec Backup Exec Active Directory Recovery Agent	<p>Lets you restore the objects and attributes from the following Microsoft applications without having to perform an authoritative or non-authoritative full restore:</p> <ul style="list-style-type: none"> ■ Active Directory ■ Active Directory Application Mode ■ Active Directory Lightweight Directory Services <p>See “How the Active Directory Recovery Agent works” on page 882.</p>
Symantec Backup Exec Agent for Microsoft Exchange Server	<p>Provides backups for your Exchange Server data. You can restore individual mailboxes, mail messages, and public folders from the Information Store backups that have Backup Exec’s Granular Recovery Technology (GRT) option enabled. Use Backup Exec Continuous Protection Server (CPS) to provide complete recovery to any point in time of the Information Store, including the latest complete transaction log.</p> <p>See “About the Backup Exec Exchange Agent” on page 1094.</p>
Symantec Backup Exec Agent for Microsoft SQL Server	<p>Protects active databases and verifies all of your SQL data automatically. You can customize your data protection needs down to the filegroup level. For fast point-in-time backups, you can use this option to run transaction log backups with truncation. Redirected restores allow you to easily restore SQL data to other SQL servers on the network.</p> <p>See “About backup strategies for SQL” on page 1258.</p>

Table 1-6 Backup Exec application protection agents (*continued*)

Item	Description
Symantec Backup Exec Agent for Microsoft SharePoint	<p>Protects all of the files and attributes associated with a SharePoint installation. You can use the GRT option to restore individual workspaces and documents from a backup of the entire farm. You can restore data to the original Information Store or redirect it to another Information Store without affecting other workspaces.</p> <p>See “About using the SharePoint Agent with SharePoint Portal Server 2003 and Windows SharePoint Services 2.0” on page 1235.</p>
Symantec Backup Exec Agent for Oracle® on Windows and Linux Servers	<p>Provides the ability to initiate backup and restore operations from Backup Exec or from the RMAN console as a Database Administrator (DBA). Provides data protection of both individual table spaces as well as complete Oracle databases. You can also include archived redo files and control files without taking them offline.</p> <p>See “About the Backup Exec Oracle Agent” on page 1315.</p>
Symantec Backup Exec Agent for Lotus Domino Server	<p>Provides seamless online backup protection for Lotus Domino servers. The Lotus Agent uses Lotus Domino APIs to support transactional logging, which protects the entire Lotus Domino server.</p> <p>See “About the Agent for Lotus Domino Server” on page 1064.</p>
Symantec Backup Exec Agent for Enterprise Vault	<p>Provides a disaster recovery solution for Enterprise Vault archived data. Recovery of the archived data is not dependent on the archive source, such as Exchange Server or a specific file system.</p>

Table 1-6 Backup Exec application protection agents (*continued*)

Item	Description
Symantec Backup Exec Agent for DB2 on Windows Servers	<p>Allows backup and restore on Microsoft Windows servers. Backup and restore jobs can be started from the Backup Exec Administration Console or from a DB2 command line processor.</p> <p>See “About the Backup Exec DB2 Agent” on page 953.</p>
Symantec Backup Exec Agent for SAP Applications	<p>Provides superior data protection by allowing backups of critical data while the application is still online and in use. The SAP Agent, certified by SAP, is a reliable solution that provides both local and remote protection of the latest versions of SAP™ databases using the SAP (BC-BRI BACKINT) interface while enabling you to efficiently manage your data.</p> <p>See “About backing up and restoring with the SAP Agent” on page 1371.</p>

About Backup Exec's virtual machine agents

The following agents allow for protection and recovery of virtual machines:

Table 1-7 Backup Exec virtual machine agents

Agent	Description
Symantec Backup Exec Agent for VMware Virtual Infrastructure	<p>Lets you back up and restore the online virtual machines that use VMware ESX Server or vCenter Server (formerly VirtualCenter). You can restore a virtual machine to its original location, or redirect it to another virtual server.</p> <p>See “About the Agent for VMware” on page 1386.</p>

Table 1-7 Backup Exec virtual machine agents (*continued*)

Agent	Description
Symantec Backup Exec Agent for Microsoft Hyper-V	<p>Lets you do the following:</p> <ul style="list-style-type: none"> ■ Back up and restore the configuration settings for the virtual server host, which is the physical computer that runs the virtual server software. ■ Back up and restore all virtual machines, which are the virtual computers that reside on the virtual server host. ■ Back up and restore selected online and offline virtual machines. ■ Redirect restores of the virtual machines to a different virtual server host or virtual machine. <p>See “About the Agent for Microsoft Hyper-V” on page 1169.</p>

About Backup Exec client protection agents

The following options provide protection for remote Microsoft Windows servers and Macintosh systems on the network, as well as automated protection of desktop and laptop systems.

Table 1-8 Backup Exec client protection agents

Item	Description
Symantec Backup Exec Desktop and Laptop Option	<p>Lets you protect all business data. It provides continuous backup protection whether users are in the office or on the road. Users can synchronize files between their desktop and laptop.</p>
Symantec Backup Exec Remote Agent for Windows Systems	<p>Provides backup and restore of remote Windows systems.</p> <p>See “About the Remote Agent for Windows Systems” on page 1945.</p>
Symantec Backup Exec Remote Agent for Macintosh Systems	<p>Enables Windows Server network administrators to perform backup and restore operations on Macintosh systems that are connected to the network.</p> <p>See “Backing up Macintosh systems” on page 1918.</p>

About Backup Exec media server storage options

The following options let you extend Backup Exec's capabilities to use larger or more efficient media storage devices or share storage resources over a SAN.

Table 1-9 Backup Exec media server storage options

Item	Description
Deduplication Option	<p>Provides the following features to support a data reduction strategy:</p> <ul style="list-style-type: none">■ Reduces the amount of disk storage that is required for backups by storing only unique data.■ Reduces backup network usage by sending only unique data across the network. <p>See “About the Deduplication Option” on page 1568.</p>
Microsoft Exchange Mailbox Archiving Option	<p>Lets you archive Microsoft Exchange Server emails that have been backed up. After data is archived, it is deleted from its source location, which reduces the amount of data on the Exchange Server.</p> <p>The Archiving Option uses Enterprise Vault technology to move archive data to disk-based vault stores. Only data that is already backed up is archived so that there is little affect on the Exchange server. End users can retrieve current and previous version of files by browsing a web interface called Backup Exec Retrieve.</p> <p>See “About the Archiving Option” on page 1414.</p>

Table 1-9 Backup Exec media server storage options (*continued*)

Item	Description
File System Archiving Option	<p>Lets you archive Windows NTFS data that has been backed up. After data is archived, it is deleted from its source location, which reduces the amount of data on the file server.</p> <p>The Archiving Option uses Enterprise Vault technology to move archive data to disk-based vault stores. Only data that is already backed up is archived so that there is little affect on the file system server. End users can retrieve current and previous versions of files by browsing a web interface called Backup Exec Retrieve.</p> <p>See “About the Archiving Option” on page 1414.</p>
Symantec Backup Exec NDMP Option	<p>Enables Backup Exec to use the Network Data Management Protocol (NDMP) to initialize and control backups and restores on supported devices.</p> <p>See “About installing the NDMP Option” on page 1849.</p>
Symantec Backup Exec Library Expansion Option	<p>Enables support for each additional drive in a robotic library. When you install Backup Exec, support for the first drive in every robotic library is included.</p> <p>See “About the Library Expansion Option ” on page 453.</p>
Symantec Backup Exec Virtual Tape Library Unlimited Drive Option	<p>Enables support for all additional drives after the first drive in each virtual tape library. When you install Backup Exec, support for every single-drive virtual tape library is included.</p> <p>See “About the Virtual Tape Library Unlimited Drive Option ” on page 452.</p>

Table 1-9 Backup Exec media server storage options (*continued*)

Item	Description
Symantec Backup Exec Central Admin Server Option	<p>Maximizes your Backup Exec investment by providing centralized administration and load balanced job processing functionality for existing or newly configured media servers.</p> <p>See “How CASO works” on page 1504.</p>
Symantec Backup Exec SAN Shared Storage Option	<p>Lets Backup Exec operate in a Storage Area Network (SAN), providing a high performance LAN-free backup solution. SAN Shared Storage Option lets multiple distributed media servers share common, centralized storage devices connected over a SAN. This configuration provides greater efficiency and fault tolerance. In addition to increasing performance and backup speeds in the SAN environment, the SAN Shared Storage Option load balances backup activity across multiple Backup Exec media servers and centralizes management tasks while lowering the total cost of hardware ownership.</p> <p>See “About installing the SAN Shared Storage Option” on page 1994.</p>

Table 1-9 Backup Exec media server storage options (*continued*)

Item	Description
Symantec Backup Exec Advanced Disk-based Backup Option (ADBO)	<p>Provides the following features:</p> <ul style="list-style-type: none"> ■ Synthetic backup assembles, or synthesizes, data from one previous full or incremental backup and subsequent incremental backups. This feature eliminates the need to run full backups. The synthesis is performed on the Backup Exec media server without accessing the remote computer. The overall backup window and the network bandwidth requirements are reduced. ■ True image restore enables Backup Exec to restore the contents of directories to what they were at the time of a full backup or incremental backup. You choose restore selections from a view of the directories as they existed at the time of the backup. Files that were deleted before then are not restored. Only the correct versions of files are restored from the appropriate full or incremental backups. Previous versions are not restored and then overwritten. ■ Offhost backup moves the backup operation away from the remote computer to a Backup Exec media server in a fiber-connected SAN environment. When the backup is moved to the media server, the remote computer is free for other operations. Offhost backup for Exchange Server backups that have the Granular Recovery Technology (GRT) option enabled are also supported. <p>See “What’s new in Backup Exec agents and options” on page 75.</p> <p>See “About the synthetic backup feature” on page 899.</p> <p>See “About true image restore” on page 912.</p> <p>See “About offhost backup” on page 919.</p>


Table 1-9 Backup Exec media server storage options (*continued*)

Item	Description
Symantec Backup Exec Storage Provisioning Option	<p>Lets you configure, manage, and monitor a storage array that is attached to the media server. A wizard guides you through the configuration of the storage array. The wizard creates the virtual disks that Backup Exec uses as job destination devices on the storage array. The Storage Provisioning Option monitors disk usage trends to send alerts when low disk space occurs on the storage arrays. Disk usage trends also provide information about whether the current disk space is sufficient, and when you should add disk space.</p> <p>See “About the Storage Provisioning Option” on page 2018.</p>

About the Administration Console

From the Administration Console, you can access the Backup Exec features.

Figure 1-2 Administration Console



The Administration Console screen includes the following components:

Table 1-10 Administration Console components

Item	Description
Menu bar	Backup Exec's menu bar appears across the top of your screen. To display a menu, click the menu name or use the keyboard shortcut. You can launch Backup Exec operations by clicking options from a menu. Some options may be unavailable until you select an item from the console screen. For example, you cannot select Rename from the Edit menu unless you have first selected an item to rename from either the Devices view or the Media view.

Table 1-10 Administration Console components (*continued*)

Item	Description
Navigation bar	<p>The navigation bar appears under the menu bar and enables you to navigate to Backup Exec's views.</p> <p>Views that can be accessed through the navigation bar include the following:</p> <ul style="list-style-type: none"> ■ Home. Use this view to quickly access the Backup Exec features that you use frequently. You can customize the Home view by adding or deleting items. ■ Media Servers. This view displays only if you have installed the Central Admin Server Option (CASO). Use this view to monitor and manage media servers in a CASO-enabled Backup Exec environment. ■ Job Setup. Use this view to perform tasks for new backup, restore, and media rotation jobs, as well as to launch some utility jobs. ■ Job Monitor. Use this view to monitor jobs and job history. Through this view, you can also access the Outlook-like job calendar. ■ Alerts. Use this to view and respond to alerts, view alert history, apply alert filters, and set up notification recipients to receive e-mail or pager communications when alerts occur. ■ Reports. Use this to view, print, save, and schedule reports about your media server, its operations, and its device and media usage. Also, you can use this to create a custom report. You can view a report in Backup Exec in a PDF or HTML format. You can also save and print reports in PDF, XML, HTML, Microsoft Excel (XLS), and Comma Separated Value (CSV) formats. ■ Devices. Use this view to configure devices and to perform device operations and access device properties pages. ■ Media. Use this view to manage your media, create media sets, and create media locations. ■ Backup. Click Backup or click the arrow next to Backup to create a backup job. ■ Restore. Click Restore or click the arrow next to Restore to create a restore job.
Share Your Ideas	<p>Share Your Ideas is a link that you can use to suggest new ideas for Symantec Backup Exec. After you have submitted your suggestions, other community members can vote or comment on the idea. The ideas that get the most votes move to the top of the list. Symantec product managers review these ideas for possible features in future releases.</p>

Table 1-10 Administration Console components (*continued*)

Item	Description
Search Knowledge Base	Type your question or keywords in the Search Knowledge Base search box, and then click the magnifying glass icon. A browser window displays the results of the knowledge base search. You must have an active Internet connection to access the Symantec Knowledge Base.
Task pane	The task pane displays on the left side of the Administration Console by default, but can be hidden by selecting View, and then selecting Task Pane. Through the task pane, you can initiate actions such as creating a new backup job or responding to an alert. The contents of the task pane are dynamic, changing according to the view selected from the navigation bar. Some options may be unavailable until an item is selected from the console screen or a prerequisite task is performed. For example, you cannot select Rename from the Devices task pane unless you have first selected an item that can be renamed, such as a user-created drive pool.
Selection pane	The Selection pane is where you select items to work with, such as files to back up or restore.
Results pane	The Results pane is the large pane on the right side of the screen that usually contains a list or tree view of items that correspond to items that are selected in the Selection pane. For example, if you select a Backup-to-Disk folder in the Selection pane, the Backup-to-Disk files that are contained in the folder display in the Results pane. This pane can be divided to display a preview pane.
Preview pane	The preview pane displays on the bottom right of the Administration Console. It displays information about the item selected in the list or tree view. This pane can be hidden by selecting View, and then selecting Preview Pane.
Status bar	The status bar appears on the bottom of the Administration Console and provides information about the media server, jobs running or scheduled to run on the server, alerts, and services running.
Columns	You can change the location of columns by dragging and dropping them. In addition, you can right-click a column to select the columns you would like to make visible, configure column settings, or sort the columns. You can also change the order of the entries in a column by clicking the column heading. For example, names of reports display in alphabetical order by default. To display report names in reverse alphabetical order, click the Name column heading on the Reports view.

About the Home view

The **Home** view on the Backup Exec Administration Console is a central place from which you can quickly access the Backup Exec features you use frequently. You can customize the Home view by adding or deleting items. **Home** view items contain Backup Exec data and links to features. You can hide or display **Help and Technical Support**, **Summary**, and **Detail** items.

See [“Configuring the Home view”](#) on page 95.

See [“Restoring the Home view's default configuration”](#) on page 95.

See [“Editing items on the Home view”](#) on page 96.

See [“Help and Technical Support items”](#) on page 96.

See [“Summary items”](#) on page 98.

See [“Detail items”](#) on page 99.

Configuring the Home view

You can customize the **Home** view by adding or deleting items. **Home** view items contain Backup Exec data and links to features. You can select to hide or display items to create shortcuts for the Backup Exec features that you use frequently.

See [“About the Home view”](#) on page 95.

You can quickly restore the Home view to its default configuration at any time.

See [“Restoring the Home view's default configuration”](#) on page 95.

To configure the Home view

- 1 On the navigation bar, click **Home**.
- 2 In the task pane, under **Layout**, select the number and type of columns that you want to display on the **Home** view.
- 3 In the task pane, under **Help and Technical Support Items**, **Summary Items**, and **Detail Items**, select the items that you want to display on the **Home** view.
- 4 Drag the items to the column and position in which you want them to display to further customize the **Home** view.

Restoring the Home view's default configuration

You can customize the **Home** view by adding or deleting items to create shortcuts for the Backup Exec features that you use frequently.

See [“Configuring the Home view”](#) on page 95.

You can quickly restore the **Home** view to its default configuration at any time.

To restore the Home view's default configuration

- 1 On the navigation bar, click **Home**.
- 2 In the task pane, under **Layout**, select **Default Layout and Content**.

Editing items on the Home view

You can edit Detail items to control what information displays on them. **Home** view items that are editable have a pencil icon in their title bar.

See [“About the Home view”](#) on page 95.

See [“Detail items”](#) on page 99.

To edit items on the Home view

- 1 On the navigation bar, click **Home**.
- 2 Click the pencil icon to edit the item.
- 3 Complete the appropriate options.
- 4 Click **OK**.

Help and Technical Support items

You can customize the Backup Exec **Home** view by selecting the items that display.

See [“About the Home view”](#) on page 95.

The **Help and Technical Support** items help you configure Backup Exec and resolve usage issues.

Table 1-11 Help and Technical Support items

Item	Description
Getting Started	Provides a series of steps that you can follow to configure logon accounts, devices, media sets, and the Intelligent Disaster Recovery Option. The steps to configure the Intelligent Disaster Recovery Option display only if you have a license for it.

Table 1-11 Help and Technical Support items (continued)

Item	Description
Technical Support	Provides the following support options to help you understand product features and functionality or troubleshoot issues: <ul style="list-style-type: none">■ Backup Exec Tech Center■ Backup Exec Technical Support■ Use MySupport to manage new or existing support cases■ Symantec Remote Assistance■ Best Practices■ Register for software alerts■ Get software patches and updates
Documentation	Provides the following documentation options to help you understand product features and functionality or troubleshoot issues: <ul style="list-style-type: none">■ View Readme■ View Administrator's Guide (PDF)■ View Administrator's Guide Addendum (PDF)
Installation Tasks	Lets you access the Installation Wizard, which you can use to install additional agents and options to other servers.
Job Creation Tasks	Lets you create backup jobs, policies, and restore jobs by using wizards.
Advanced Configuration Tasks	Lets you perform the following advanced configuration tasks: <ul style="list-style-type: none">■ Set job defaults and preferences■ Configure alerts and notifications■ IDR Preparation Wizard (if installed)■ Configure the Symantec Volume Snapshot Provider

Table 1-11 Help and Technical Support items *(continued)*

Item	Description
Device and Media Tasks	<p>Lets you perform the following device and media tasks:</p> <ul style="list-style-type: none"> ■ Configure media sets ■ Configure devices ■ Configure device pools
Symantec RSS Reader	<p>Lets you view and add Backup Exec and Symantec RSS feeds.</p>
Licensing and Maintenance Contract Information	<p>Lets you enter serial numbers so that Backup Exec can search for expiration dates for your maintenance contracts. After you enter the serial numbers, you can set alerts to remind you to renew each contract before the maintenance contract expires.</p> <p>Entering the serial numbers also ensures your access to the latest upgrades and to technical support for your products.</p>

Summary items

You can customize the Backup Exec **Home** view by selecting the items that display.

See [“About the Home view”](#) on page 95.

The **Summary** items provide concise overviews on the status of your alerts, jobs, devices, and media.

Table 1-12 Summary items

Item	Description
Active Alert Summary	<p>Provides a summary view of any active alerts.</p>
Job History Summary	<p>Provides a summary view of job history. You can customize the amount of time for which you display information about completed jobs. The job information includes the number of jobs completed, the amount of data that was backed up, and the number of media that was used. It also details the job statuses.</p>

Table 1-12 Summary items (continued)

Item	Description
Current Job Summary	Provides a summary view of current jobs. The summary displays the number of active, scheduled, and on-hold jobs.
Device Summary	Provides a summary view of device information. The device information includes the number of devices and their current statuses.
Media Summary	Provides a summary view of media information. The media information displays the number of overwritable and appendable media that are available. You can also view or change the default media overwrite protection level.

Detail items

You can customize the Backup Exec **Home** view by selecting the items that display. See [“About the Home view”](#) on page 95.

Additionally, you can edit the information that displays on Detail items.

See [“Editing items on the Home view”](#) on page 96.

The **Detail** items provide itemized overviews on the status of your alerts and jobs.

Table 1-13 Detail items

Item	Description
Active Alerts	Lets you view all active alerts. You can display any or all of the following types of alerts: <ul style="list-style-type: none"> ■ Attention Required ■ Error ■ Warning ■ Information
Job History	Lets you view job history for the specified period of time. You can select the window of time for which you want to view completed jobs.

Table 1-13 **Detail items** (*continued*)

Item	Description
Current Jobs	Lets you view all current jobs. You can choose to display any or all of the following types of current jobs: <ul style="list-style-type: none">■ Active Jobs■ Scheduled Jobs■ Jobs On Hold

Installing Backup Exec

This chapter includes the following topics:

- [About installing Backup Exec](#)
- [Before you install](#)
- [System requirements](#)
- [About the differences between a typical installation and a custom installation](#)
- [Installing a typical installation of Backup Exec](#)
- [Installing a custom installation of Backup Exec](#)
- [Installing additional Backup Exec options to the local media server](#)
- [Special considerations for installing Backup Exec to remote computers](#)
- [Push-installing Backup Exec to remote computers](#)
- [About installing Backup Exec options to remote computers](#)
- [Push-installing the Remote Agent and Advanced Open File Option to remote computers](#)
- [Push-installing the Desktop Agent and DLO Maintenance Service from the media server to remote computers](#)
- [About installing the Remote Agent for Windows Systems](#)
- [Installing the Remote Administrator](#)
- [Installing Backup Exec using the command line \(silent mode\)](#)
- [Installing a trial version of Backup Exec agents and options](#)
- [About the installation log](#)

- [Repairing Backup Exec](#)
- [Starting and stopping Backup Exec services](#)
- [Uninstalling Backup Exec](#)
- [Uninstalling Backup Exec options from the local media server](#)
- [About updating Backup Exec with LiveUpdate](#)
- [Viewing license information](#)
- [Adding licenses](#)
- [Finding installed licenses in your environment](#)
- [About Backup Exec maintenance contract information](#)
- [About upgrading from previous versions of Backup Exec](#)
- [Post-installation tasks](#)

About installing Backup Exec

Several methods are available for installing Backup Exec.

You can do the following:

- Use the installation wizard, which guides you through the installation process.
- Use the command line, which is called silent mode installation. The silent mode installation uses the Setup.exe program on the Backup Exec installation media.

You can install Backup Exec and its options on a local computer, a remote computer, or both. Additionally, you can install the Remote Administrator , which lets you administrate the media server from a remote Windows server or workstation.

Backup Exec may install the additional products:

- Symantec LiveUpdate
- Microsoft XML Core Services (MSXML) 6.0
- Microsoft Report Viewer Redistributable 2005
- Microsoft.NET Framework 3.5 SP1
- Microsoft Windows Imaging Component
- Microsoft SQL Express 2005 SP3

See [“Installing a custom installation of Backup Exec”](#) on page 119.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

See [“Push-installing Backup Exec to remote computers”](#) on page 126.

See [“Push-installing the Remote Agent and Advanced Open File Option to remote computers”](#) on page 134.

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 153.

Before you install

Before you install Backup Exec, you should do the following:

- Run the Backup Exec Environment Check on the computer on which you want to install Backup Exec. The Environment Check analyzes the computer to make sure that the installation process can complete. If Backup Exec finds configuration issues that can be fixed during the installation, or that may prevent the installation, warnings appear. Although the Environment Check runs automatically during installation, you may want to run it manually before you install Backup Exec or before you back up data with Backup Exec. See [“Checking your environment before installing”](#) on page 105.
- Install the storage device hardware (controller, drives, robotic libraries) on the media server. Refer to the documentation that is included with your storage device hardware for installation instructions. Use the appropriate Windows hardware setup functions to configure your controller and storage devices. Refer to your Microsoft Windows documentation for more information.
- Check your Windows security settings to make sure that they work properly with the Backup Exec service account. See [“About the Backup Exec service account”](#) on page 106.
- If the drive on which you want to install Backup Exec is encrypted or compressed, and you would like to use a default SQL Express database, verify that an unencrypted and uncompressed drive is available for SQL Express installation.
- Check the computer name of the computer on which you want to install Backup Exec. It should only use standard ANSI characters. You may receive errors if you install Backup Exec on a computer with a name that uses non-standard characters.
- If you want to install Backup Exec to a non-English version of Windows, download the SQL Express SP3 setup file from the Microsoft Web site before you install Backup Exec if all of the following are true:

- You want to use a local Backup Exec SQL Express instance.
- You have non-English SQL Server instances on the computer on which you want to install Backup Exec.
- Exit all other programs.

About the Environment Check

The Symantec Backup Exec Environment Check is a utility that runs on a computer automatically during installation and that reports the following:

- If the computer meets the minimum requirements for installation, such as the operating system, disk and physical memory, sufficient logon account privileges.
 See “[System requirements](#)” on page 114.
- If third-party software that uses Backup Exec ports is configured correctly.
- If required components are installed, and if they are versions that are compatible with Backup Exec.
- If previous versions of Backup Exec and Backup Exec options are installed.
- If storage device hardware and associated drivers are properly installed and recognized by the Windows operating system.
- If the computer meets the minimum requirements for installation of the Desktop and Laptop Option.

One of the following results are reported for each item:

Table 2-1 Environment Check results

Result	Description
Passed	There are no incompatibilities to prevent the Backup Exec installation. For hardware, this result indicates that the hardware configuration is recognized by Backup Exec.
Warning	An incompatibility with Backup Exec exists, but can be resolved during the Backup Exec installation.
Failed	An incompatibility with Backup Exec exists, and it will cause the installation to fail. Action is required before you can successfully install Backup Exec.

Although the Environment Check runs automatically during installation, you may want to run it manually before installing Backup Exec or before backing up data with Backup Exec.

To remove the name of a computer from the list of computers on which the Environment Check runs

- Select the computer from the list.
- Click **Remove**.

6 If you want to save the results of the Environment Check, check **Save Results To**.

To change the location where the Environment Check results are saved, click **Change Path** to browse to a new location.

7 Click **Finish**.

About the Backup Exec service account

All Backup Exec services on the media server run in the context of a user account that is configured for the Backup Exec system services. You can create this account during the Backup Exec installation, or you can use an existing user account. To create a service account for Backup Exec during installation, enter the name and password of an Administrator account for the Backup Exec services to use.

Note: The Backup Exec service account and the Backup Exec System Logon Account are set to the same user name when Backup Exec is installed. If you need to change the user name for the service account is no longer used, then you should also change the Backup Exec System Logon Account to use new credentials.

See [“Changing service account information”](#) on page 107.

If this computer is in a domain, enter a Domain Administrators account, or an equivalent account that is part of the Domain Admins group. In the Domain list, select or enter the Domain name.

If this computer is in a workgroup, enter an Administrators account, or an equivalent account that is part of the Administrators group on the computer. In the Domain list, select or enter the computer name.

The account that you designate for Backup Exec services, whether it is a new account or an existing user account, is assigned the following rights:

- Authenticate as any user and gain access to resources under any user identity.
- Create a token object, which can then be used to access any local resources.
- Log on as a service.
- Administrative rights (provides complete and unrestricted rights to the computer).

- Backup operator rights (provides rights to restore files and directories).
- Manage auditing and security log.

See [“Required user rights for backup jobs”](#) on page 331.

Due to security implementations in Microsoft Small Business Server, the service account must be Administrator.

You cannot install Backup Exec with an account that has a blank password on Windows Server 2003/2008 or XP computers unless Windows is configured to allow it. If you try to do so, the following error message appears when Backup Exec services are created:

The given password is not correct for account [server]\[username].

You can, however, configure Windows to allow for blank passwords. For more information, see your Windows documentation.

Changing service account information

On the media server, all Backup Exec services run in the context of a user account that is configured for the Backup Exec system services.

Note: The Backup Exec service account and the Backup Exec System Logon Account are set to the same user name when Backup Exec is installed. If you need to change the user name for the service account is no longer used, then you should also change the Backup Exec System Logon Account to use new credentials.

See [“About the Backup Exec service account”](#) on page 106.

To change service account information

- 1 On the **Tools** menu, click **Backup Exec Services**.
- 2 Click **Services credentials**.
- 3 Click **Change service account information**.
- 4 Enter the user name, domain, and password for the new service account.

See [“Service Account Information options”](#) on page 107.

Service Account Information options

On the media server, all Backup Exec services run in the context of a user account that is configured for the Backup Exec system services.

See [“Changing service account information”](#) on page 107.

Table 2-2 Service Account Information options

Item	Description
Change service account information	Enables you to change the user name, domain, and password for the service account.
User name	Indicates the user name for the service account.
Domain name	Indicates the name of the domain for the service account.
New password	Indicates the password for the service account.
Confirm password	Confirms the password that you typed in the New password field.
Change startup options	Lets you change the startup options for the service account.
Automatic	Indicates that the service account starts automatically at system startup.
Manual	Indicates that the service account does not start automatically at system startup. You must start it manually.
Disabled	Indicates that the service account is disabled at system startup.
Grant Backup Exec system service rights to the service account	Lets the service account have the system service rights.

About changing Windows security

You can set up Windows security with the Backup Exec service account to protect your data.

Depending on how the Windows network is configured, change security properties for the following scenarios:

- Servers in one domain.
- Servers and selected workstations in one domain.
- Servers in more than one domain.
- Servers and workstations in more than one domain.

You can change Windows security to give the Backup Exec service account administrative rights in the appropriate domains and workstations. You must grant the Backup Exec service account administrative rights to give Backup Exec access to the administrative shares (for example, C\$) and the ability to protect the Windows registry.

Use the Active Directory Users and Computers tool and Domain and Trusts tool in the Active Directory administrative tools group to change Windows security properties.

See [“Changing Windows security to back up servers \(only\) in one domain”](#) on page 109.

See [“Changing Windows security to back up servers and selected workstations in one domain”](#) on page 109.

See [“Changing Windows security to back up servers in more than one domain”](#) on page 110.

See [“Changing Windows security to back up servers and workstations in more than one domain”](#) on page 111.

Changing Windows security to back up servers (only) in one domain

You can change Windows security to give the Backup Exec service account administrative rights in the appropriate domains and workstations. You must grant the Backup Exec service account administrative rights to give Backup Exec access to the administrative shares (for example, C\$) and the ability to protect the Windows registry.

To change Windows security to back up servers (only) in one domain

- ◆ When prompted for a user name, add the name of an existing or new service account (for example, Administrator) as a member of the local Administrators group for the Domain. It is highly recommended that you also enter a password.

See [“About changing Windows security”](#) on page 108.

Changing Windows security to back up servers and selected workstations in one domain

You can change Windows security to give the Backup Exec service account administrative rights in the appropriate domains and workstations. You must grant the Backup Exec service account administrative rights to give Backup Exec access to the administrative shares (for example, C\$) and the ability to protect the Windows registry.

Table 2-3 How to change Windows security to back up servers and selected workstations in one domain

Step	Description
Step 1	Add the name of an existing or new service account (for example, Administrator) as a member of the Global Domain Admins group.
Step 2	Ensure that on each workstation in the domain you want to back up, the Global Domain Admins group is a member of the workstation's local Administrators group.

See [“About changing Windows security”](#) on page 108.

Changing Windows security to back up servers in more than one domain

You can change Windows security to give the Backup Exec service account administrative rights in the appropriate domains and workstations. You must grant the Backup Exec service account administrative rights to give Backup Exec access to the administrative shares (for example, C\$) and the ability to protect the Windows registry.

Table 2-4 How to change Windows security to back up servers in more than one domain

Step	Description
Step 1	<p>Establish a One Way Trust Relationship between the Host Domain (the domain in which the media server resides) and the Target Domains (the domains that are to be backed up).</p> <p>Do the following in the order listed:</p> <ul style="list-style-type: none"> ■ In the Host Domain, permit the Target Domains to trust the Host Domain. ■ In each Target Domain, trust the Host Domain.
Step 2	In each Target Domain, add the Host Domain's name of an existing or new service account (for example, Administrator) in the local Administrators group.

See [“About changing Windows security”](#) on page 108.

Changing Windows security to back up servers and workstations in more than one domain

You can change Windows security to give the Backup Exec service account administrative rights in the appropriate domains and workstations. You must grant the Backup Exec service account administrative rights to give Backup Exec access to the administrative shares (for example, C\$) and the ability to protect the Windows registry.

Table 2-5 How to change Windows security to back up servers and workstations in more than one domain

Step	Description
Step 1	<p>Establish a One Way Trust Relationship between the Host Domain (the domain in which the media server resides) and the Target Domains (the domains that are to be backed up).</p> <p>Do the following in the order listed:</p> <ul style="list-style-type: none">■ In the Host Domain, permit the Target Domains to trust the Host Domain.■ In each Target Domain, trust the Host Domain.
Step 2	<p>In each Target Domain, add the Host Domain's name of an existing or new service account (for example, Administrator) in the local Administrators group.</p>
Step 3	<p>On each workstation to back up, add the Host Domain's name of an existing or new service account (for example, Administrator) in the Local Administrators group.</p>

See [“About changing Windows security”](#) on page 108.

About Microsoft SQL Server 2005 Express Edition components installed with Backup Exec

The Backup Exec installation program installs Microsoft SQL Server 2005 Express Edition components that are required to run Backup Exec.

Backup Exec prompts you to do one of the following:

- Install the required Microsoft SQL Express components with Backup Exec and create a default Backup Exec instance.
- Select a Microsoft SQL Server 2005 (SP3) or SQL Server 2008 instance that already exists on the network on which you want to run Backup Exec.

During the installation process and upgrade process, Backup Exec stops and starts the SQL service several times. Other user-created databases that use the SQL Server instance are unavailable during the process. To avoid such conflicts, you should install Backup Exec into its own SQL instance.

If you choose to install Backup Exec into an existing SQL 2005 instance, make sure that SQL 2005 Service Pack 3 or later is installed before you continue with the installation.

Caution: Backup Exec may not function properly if you install it into an existing SQL instance that uses case-sensitive collation. Symantec recommends that you avoid installing Backup Exec to a SQL instance that uses case-sensitive collation.

When Backup Exec is installed into an existing instance, the automated master database restore feature is not available. To recover the Master database, you must replace it with the Master database copy that Backup Exec automatically creates and updates when the Master database is backed up.

You cannot install multiple Backup Exec databases on the same SQL Server instance.

Note: If you are installing a managed media server, it is recommended that you select a local Microsoft SQL Server 2005 (SP3) instance or later on which to install the Backup Exec database for this managed media server. Do not select the same SQL Server instance that is used by the central administration server.

See [“System requirements”](#) on page 114.

About Backup Exec’s standard features

When you enter a Backup Exec license key, you can select any of the following additional features that are available for installation.

Table 2-6 Backup Exec's standard features

Feature	Description
Tape Device Drivers	Installs the Symantec tape device drivers for all supported tape devices that are attached to the server. If there are no tape devices attached to your media server, uncheck this option.
Online Documentation	Installs the Backup Exec Administrator's Guide in a pdf file format.
Enable Robotic Library Support	Enables support for tape libraries, or optical robotic libraries and library storage systems. Backup Exec includes support for one drive in every robotic library. Each additional drive in a library requires a Library Expansion Option license.
Copy Server Configurations	Enables you to copy jobs, selection lists, and job templates between media servers. This option is recommended for environments that contain multiple Backup Exec media servers. This option is required for the Central Admin Server Option.
Managed Media Server	Installs the managed media server component of the Central Admin Server Option. You can install managed media servers after you install a central administration server.
Advanced Open File Option	Ensures that all files on a Windows computer are backed up even if they are open and in use. This option is free with each license of Backup Exec, Backup Exec Remote Agent for Windows Systems, and Backup Exec application agents and options. To control specific snapshot settings, install the Advanced Open File Option along with the Advanced Disk-based Backup Option.
Intelligent Disaster Recovery	Provides a recovery solution for both local and remote Windows computers.
Virtual Tape Library Support	Provides support for every single-drive Virtual Tape Library (VTL). You must purchase the Virtual Tape Library Unlimited Drive Option to support additional drives in each VTL. If you select this option, the Enable Robotic Library Support option is selected automatically. You cannot uncheck Enable Robotic Library Support unless you uncheck Virtual Tape Library Support .

All other options and agents require the purchase of additional licenses. Installing a trial version enables many options that must be purchased separately and are not included as part of Backup Exec.

If you have a licensed version of Backup Exec, you can use a trial version of most options and agents for a specified period of time.

See “[Installing a trial version of Backup Exec agents and options](#)” on page 166.

System requirements

The following are the minimum system requirements to run this version of Backup Exec:

Table 2-7 Minimum system requirements

Item	Requirements
Operating system	<p>You can find a list of compatible operating systems, platforms, and applications at the following URL: http://entsupport.symantec.com/umi/V-269-1</p> <p>You cannot install a Backup Exec media server on a computer that runs the Windows Server Core installation option of Windows Server 2008. You can only install the Backup Exec Remote Agent for Windows Systems on Server Core computers.</p> <p>You cannot install SQL Express or SQL Server 2005 on a Windows Server 2008 computer that is configured in a Read Only Domain Controller (RODC) role. The Read Only Domain Controller role does not let you use the local accounts which are required for SQL Express and SQL Server 2005. When you install Backup Exec on an RODC computer you must select a remote SQL instance for the Backup Exec Database.</p>
Additional application support	<p>You can use Backup Exec with Microsoft Windows Microsoft Operations Manager (MOM) 2005.</p>
Internet browser	<p>Internet Explorer 6.0 or later. Service Pack 1 is required for SQL Server 2005 Express.</p>
Processor	<p>Intel Pentium, Xeon, AMD, or compatible.</p>

Table 2-7 Minimum system requirements (*continued*)

Item	Requirements
Memory	<p>Required: 512 MB RAM</p> <p>Recommended: 1 GB RAM (or more for better performance)</p> <p>Note: RAM requirements may vary depending on the operations performed, the options installed, and the specific computer configuration.</p> <p>For the Central Admin Server Option: 512 MB RAM required, 1 GB recommended.</p> <p>Virtual Memory Recommendations: 20 MB above the Windows recommended size for total paging file size (total for all disk volumes). Refer to your Microsoft Windows help documentation for instructions on how to view or set the paging file size.</p>
Installation disk space	<p>1.44 GB (Typical installation)</p> <p>2.32 GB (Includes all options)</p> <p>Note: Disk space requirements may vary depending on the operations performed, the options installed, and the specific system configuration. Backup Exec database and catalogs require additional space. An additional 330 MB is required for SQL Express.</p>
Other Hardware	<p>The following hardware is recommended:</p> <ul style="list-style-type: none"> ■ Network interface card or a virtual network adapter card. ■ CD/DVD drive. ■ (Recommended) A mouse. ■ (Optional for pager notification) Modem that Microsoft Windows supports. ■ (Optional for printer notification) Printer that Microsoft Windows supports.
Storage Hardware	<p>You can use storage media drives, robotic libraries, removable storage devices, and non-removable hard drives.</p> <p>You can find a list of compatible devices at the following URL: http://entsupport.symantec.com/umi/V-269-2</p> <p>Support is available for the first drive in each robotic library when you purchase Backup Exec. To enable support for each additional robotic library drive, you must purchase the Backup Exec Library Expansion Option.</p>

See “Installing a typical installation of Backup Exec” on page 116.

See [“Installing a custom installation of Backup Exec”](#) on page 119.

About the differences between a typical installation and a custom installation

The Backup Exec installation program provides two methods of installation: typical and custom. A typical installation is a simpler installation method than a custom installation and is designed for small or uncomplicated environments. For example, if you use a local media server and a few Backup Exec agents or options, then a typical installation may be best for you. A custom installation is designed for large or complex environments. You can also use the custom installation method if you prefer to set all of your options. For example, if you use a remote media server or use the Central Admin Server Option or the SAN Shared Storage Option, you should perform a custom installation.

With a typical installation, Backup Exec makes the following decisions for you, based on common installation scenarios:

- Backup Exec is installed to a local media server.
- SQL Express is installed with the default instance.
- Agents and options are installed if you enter the license keys for them. If you do not enter any license keys, a trial version of Backup Exec is installed.

Note: If you enter a license key for an agent or option that is not compatible with the typical installation method, such as the Central Admin Server Option, then the installation program automatically switches to the custom installation method.

- LiveUpdate runs automatically.

See [“Installing a typical installation of Backup Exec”](#) on page 116.

See [“Installing a custom installation of Backup Exec”](#) on page 119.

Installing a typical installation of Backup Exec

A typical installation is designed for small or uncomplicated environments. In a typical installation, Backup Exec makes many decisions for you, based on common installation scenarios. If you prefer to set all of your options, you should use the custom installation option.

See [“Installing a custom installation of Backup Exec”](#) on page 119.

See “[About the differences between a typical installation and a custom installation](#)” on page 116.

Note: Before you install, make sure that your license keys for Backup Exec and any agents or options are available. If you do not have license keys, go to the following URL to activate the product:

<https://licensing.symantec.com>

License keys are required to install Backup Exec and its agents and options. However, you can install a trial version of Backup Exec without a license key.

To install a typical installation of Backup Exec

- 1 From the installation media browser, click **Install Products**, and then select **Backup Exec**.

If the Microsoft.NET Framework 3.5 SP1 is not already installed on this computer, Backup Exec installs it. The installation of the Microsoft.NET Framework may take some time.

- 2 On the **Welcome** panel, read the license agreement, and then click **I accept the terms of the license agreement**.

- 3 Click **Next**.

- 4 On the **Installation Type** panel, click **Typical Installation**, and then click **Next**.

For first-time installations and upgrade installations, the Backup Exec Environment Check runs automatically after you click **Next**.

- 5 Review the results of the Environment Check.

- 6 Do one of the following:

- If the Environment Check does not reveal any issues that may prevent a successful installation of Backup Exec, click **Next**.
- If the Environment Check reveals any issues that may prevent a successful installation of Backup Exec, click **Cancel** to exit the wizard. Correct the issues before you attempt to install Backup Exec again.

- 7 Select one of the following methods to enter license keys:

To enter license keys manually

Do the following in the order listed:

- Type the Backup Exec license key into the license key field.
- Click **Add**.
- Repeat for each license key for each option or agent that you want to install.

To import license keys from a file

Do the following in the order listed:

- Click **Import from file**.
- Select the besernum.xml file.

The default location for the besernum.xml file is %allusersprofile%\Application Data\Symantec\Backup Exec.

To install a trial version

Leave the license key field blank.

8 Click **Next**.

9 On the **Service Account** panel, provide a user name, password, and domain for an Administrator account that the Backup Exec system services can use.

10 If you want to change the directory where Backup Exec files are installed, click **Change**, and then select a new location.

11 Click **Next**.

12 On the **Remote Computers** panel, do one of the following:

To push-install the Remote Agent for Windows Systems to a single remote computer

Do the following in the order listed:

- Click **Add**, and then select **Add a Single Computer**.
- Enter the name of the remote computer, and then enter the credentials for the remote computer.
- Click **Next**.

To push-install the Remote Agent for Windows Systems to multiple remote computers

Do the following in the order listed:

- Click **Add**, and then select **Add Multiple Computers with the Same Settings**.
- Enter the name of a remote computer, and then click **Add to List**. Repeat this step for each remote computer that you want to install the Remote Agent on.
- Enter the credentials for the remote computers.
- Click **Next**.

To skip the installation of the Remote Agent for Windows Systems

Go to step 13.

13 Click **Next**.

14 Review the installation summary, and then click **Install**.

The installation process creates an installation log named Bkupinst.htm on the computer where Backup Exec is installed.

Installing a custom installation of Backup Exec

A custom installation is designed for large or complex environments or for customers who prefer to set all of their options. For example, if you use a remote media server or use the Central Admin Server Option or the SAN SSO Option, you should perform a custom installation. If you prefer to have Backup Exec set default options or if you have a small, uncomplicated environment, you can use the typical installation.

See [“Installing a typical installation of Backup Exec”](#) on page 116.

Note: If you install Backup Exec through terminal services and the installation media is on a shared drive (network share), you must install it using a UNC path. Installation by mapped drives is not supported in this situation.

Note: Before you install, make sure that your license keys for Backup Exec and any agents or options are available. If you don't have license keys, go to the following URL to activate the product:

<http://licensing.symantec.com>

License keys are required to install Backup Exec and its agents and options. However, you can install a trial version of Backup Exec without a license key.

To install a custom installation of Backup Exec

- 1 From the installation media browser, click **Install Products**, and then select **Backup Exec**.

If the Microsoft.NET Framework 3.5 SP1 is not already installed on this computer, Backup Exec installs it. The installation of the Microsoft.NET Framework may take some time.

- 2 On the **Welcome** panel, read the license agreement, and then click **I accept the terms of the license agreement**.
- 3 Click **Next**.
- 4 On the **Install Type** panel, click **Custom Installation**, and then click **Next**.
- 5 On the **Menu** panel, check **Local Install**, and then click **Install Backup Exec software and options**.
- 6 Click **Next**.

For first-time installations and upgrade installations, the Backup Exec Environment Check runs automatically after you click **Next**.

- 7 Review the results of the Environment Check.
- 8 Do one of the following:
 - If the Environment Check does not reveal any issues that may prevent a successful installation of Backup Exec, click **Next**.
 - If the Environment Check reveals any issues that may prevent a successful installation of Backup Exec, click **Cancel** to exit the wizard. Correct the issues before you attempt to install Backup Exec again.
- 9 Select one of the following methods to enter license keys:

- | | |
|------------------------------------|---|
| To enter license keys manually | Do the following in the order listed: <ul style="list-style-type: none"> ■ Type the Backup Exec license key into the license key field. ■ Click Add. ■ Repeat for each license key for each option or agent that you want to install. |
| To import license keys from a file | Do the following in the order listed: <ul style="list-style-type: none"> ■ Click Import from file. ■ Select the besernum.xml file. <p>The default location for the besernum.xml file is %allusersprofile%\Application Data\Symantec\Backup Exec.</p> |
| To install a trial version | Leave the license key field blank. |

10 Click Next.

The license keys that you entered are saved to the besernum.xml file, which is located in the %allusersprofile%\Application Data\Symantec\Backup Exec directory.

11 Select any additional options or agents that you want to install.

12 Click Next.

If you selected the File System Archiving Option or the Microsoft Exchange Mailbox Archiving Option, the Archiving Option Environment Check runs. The Archiving Option Environment Check verifies that the computer meets the minimum requirements for installing and configuring Enterprise Vault. If the computer does not meet the minimum requirements, you must uncheck the archiving options or fix the errors before you can continue with the installation.

13 Do one of the following:

- | | |
|---|--|
| To change the directory where the Backup Exec files are installed | Click Change to select a new directory. |
| To accept the default directory (recommended) | Proceed to step 14. |

Symantec recommends that you do not select a mount point as the destination directory because if you delete the mount point, Backup Exec is uninstalled.

14 Click **Next**.

15 Provide a user name, password, and domain for an Administrator account that the Backup Exec system services can use, and then click **Next**.

See “[About the Backup Exec service account](#)” on page 106.

16 On the **Choose SQL Server** panel, do one of the following to select a location to store the Backup Exec Database .

The **Choose SQL Server** panel does not appear for upgrades. You cannot change the database location during the upgrade process. If you want to change the database location after the upgrade, use BE Utility.

To create a local Backup Exec SQL Express instance

Do the following in the order listed:

- Click **Create a local Backup Exec SQL Express instance to store the database on**.
- To change the location of the Backup Exec SQL Express instance, click **Browse**.
- Select the location, and then click **OK**.

To use an existing SQL Server 2005 or SQL Server 2008 instance

Do the following in the order listed:

- Click **Use an existing instance of SQL Server 2005 (SP3a or later) or SQL Server 2008 on the network to store the database on**.
- Select the instance.

When Backup Exec is installed into an existing instance, the automated master database restore feature is not available. To recover the Master database, replace it with the Master database copy that Backup Exec automatically creates and updates when the Master database is backed up.

Caution: During the installation process and upgrade process, Backup Exec stops and starts the SQL service several times. Other user-created databases that use the SQL Server instance are unavailable during the process. To avoid such conflicts, you should install Backup Exec into its own SQL instance.

17 Click **Next**.

Backup Exec attempts to connect to the instance.

18 If the Symantec Backup Exec Database panel appears, perform the following steps to identify the location of the SQL Express SP3 setup file:

- Click **Browse**.
 - Navigate to the location where you downloaded the SQL Express SP3 setup file.
 - Click **OK**.
 - Click **Next**.
- 19** If you are prompted, select how the **Symantec Device Driver Installer** should install device drivers for the tape storage devices that are connected to the server, and then click **Next**.
- Symantec recommends that you select **Use Symantec device drivers for all tape devices**.
- 20** If you are prompted, enter information or choose settings for the additional options that you want to install, and then click **Next** after each selection.
- 21** Read the Backup Exec installation summary, and then click **Install**.
- The installation process takes several minutes to complete. During the process, the progress bar may not move for several minutes.
- 22** When the installation is complete, you can run LiveUpdate, view the readme, and create a shortcut to Backup Exec on the desktop.
- 23** Click **Finish** to close the installation wizard.
- 24** If Restart System appears, restart the computer in order for the configuration to take effect.

The installation process creates an installation log named Bkupinst.htm on the computer where Backup Exec is installed.

See [“About the installation log”](#) on page 166.

See [“Post-installation tasks”](#) on page 182.

See [“About Backup Exec’s standard features”](#) on page 112.

See [“About upgrading from previous versions of Backup Exec”](#) on page 180.

Installing additional Backup Exec options to the local media server

You can install agents and options when you install Backup Exec. However, if you have already installed Backup Exec and want to install additional options, review the documentation for those options to ensure that your system meets all minimum requirements. The Backup Exec services may be stopped while the additional

options are installed. If any active jobs are running, you are prompted to stop them, or to wait for the jobs to finish.

See [“Installing a custom installation of Backup Exec”](#) on page 119.

Note: If you install Backup Exec through Terminal Services and the installation media is on a shared drive (network share) you must install using a UNC path. Installation by mapped drives is not supported.

If you installed the trial version or the Not For Resale (NFR) edition of Backup Exec, you can install trial versions of the additional options. If you have a licensed version of Backup Exec, you can use a trial version of most options and agents for a specified period of time.

See [“Installing a trial version of Backup Exec agents and options”](#) on page 166.

Note: If the Central Admin Server Option is installed, and you want to install additional options on a managed media server, you can pause the managed media server. When a managed media server is paused, the central administration server does not delegate jobs to it. When the installation is complete, un-pause, or resume, the managed media server.

See [“Pausing a managed media server in CASO”](#) on page 1561.

To install additional Backup Exec options to the local media server

- 1 On the **Tools** menu, click **Install Options and License Keys on this Media Server**.
- 2 Select one of the following methods to enter license keys:

To manually enter license keys

Do the following in the order listed:

- Type a license key into the license key field.
- Click **Add**.
- Repeat for each license key for each option or agent that you want to install.

To import license keys from a file

Do the following in the order listed:

- Click **Import from file**.
- Select the besernum.xml file.

To install a trial version

Leave the license key field blank.

- 3 Click **Next**.

- 4 Select the additional options that you want to install, and then click **Next**.
- 5 If you are prompted, enter information or choose settings for the additional options that you want to install. Click **Next** after each selection.
- 6 Read the Backup Exec installation summary, and then click **Install**.

The Backup Exec services are stopped while the additional options are installed. If any active jobs are running, you are prompted to stop them, or to wait for the jobs to finish.

When the installation is complete, the services are restarted.

- 7 Click **Finish**.

Special considerations for installing Backup Exec to remote computers

There are special considerations that you should be familiar with before you install Backup Exec to remote computers.

Table 2-8 Special considerations for installing Backup Exec to remote computers

Item	Consideration
Windows XP SP2/Server 2003 SP1	<p>To push-install Backup Exec to a Windows XP SP2/Server 2003 computer, you must enable File and Printer Sharing on the Windows Firewall Exceptions list for the following ports:</p> <ul style="list-style-type: none"> ■ 135 (RPC) ■ 445 (TCP) ■ 103X (mostly 1037) ■ 441 (RPC) <p>For more information about the Windows Firewall Exceptions list, refer to your Microsoft Windows documentation.</p> <p>During the installation process, Backup Exec sets the Remote Launch and Remote Access security permissions for the Administrator's group.</p> <p>You should enable the "Allow remote administration exception" group policy for the computer to which you push the installation.</p>

Table 2-8 Special considerations for installing Backup Exec to remote computers (*continued*)

Item	Consideration
Windows Server 2008	To push-install Backup Exec to a computer that runs Windows Server 2008, you must enable certain items on the destination computer's Windows Firewall Exceptions list. You must enable the following items: <ul style="list-style-type: none">■ File and Printer Sharing■ Windows Management Instrumentation (WMI) For more information refer to your Microsoft Windows documentation.
Symantec Endpoint Protection (SEP) 11.0 or later	To push-install Backup Exec to a computer that runs Symantec Endpoint Protection (SEP) version 11.0 or later, you must configure SEP to share files and printers. The file and printer sharing feature is turned off by default.

See [“Push-installing Backup Exec to remote computers”](#) on page 126.

Push-installing Backup Exec to remote computers

If you install Backup Exec through Terminal Services and the installation media is on a shared drive (network share) you must use a UNC path. Installation by mapped drives is not supported.

You can set up multiple server installations. Backup Exec processes up to five remote computer installations concurrently.

Before, you install Backup Exec to remote computers, you should review the special considerations.

See [“Special considerations for installing Backup Exec to remote computers”](#) on page 125.

Note: You can also use Microsoft's Add or Remove Programs utility to install Backup Exec to a remote computer. See your Microsoft documentation for more information.

The installation process creates an installation log named Bkupinst.htm on the computer where Backup Exec is installed.

See [“About the installation log”](#) on page 166.

To push-install Backup Exec to remote computers

1 Do one of the following:

To push-install Backup Exec to remote computers from the installation media

Do the following steps in the order listed:

- From the installation media browser, click **Installation**, and then click **Backup Exec**.
- On the **Welcome** panel, click **Next**.
- Select **I accept the terms of the license agreement**, and then click **Next**.
- Select **Custom installation**.
- Uncheck **Local Installation**, and then check **Remote Installation**.
- Click **Next**.
- On the **Remote Computers** panel, click **Add**.
- To install Backup Exec on one remote computer, select **Add a Single Computer**, or to install Backup Exec on multiple computers using the same settings, select **Add Multiple Computers with the Same Settings**.

To push-install Backup Exec to remote computers from the Backup Exec media server

On the **Tools** menu, click **Install Agents and Media Servers on Other Servers**.

- 2 Select **Symantec Backup Exec**, and then click **Next**.
- 3 Type the fully qualified name, IP address, or computer name of the remote computer or click **Browse Remote Computers** to locate the remote computer.
- 4 Click **Add to List**, and then repeat steps 3 and 4 for each remote computer to which you want to push-install the programs.

If you are push-installing from the installation media and you selected **Add a Single Computer** in step 1, you can skip this step.

- 5 Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.

You must use Administrator credentials. These remote computer logon credentials are not the same as the Backup Exec service account credentials in step 12.

- 6 Click **Next**.

7 Select one of the following methods to enter license keys:

To enter license keys manually	Do the following in the order listed: <ul style="list-style-type: none">■ Type a license key into the license key field.■ Click Add.■ Repeat for each license key for each option or agent that you want to install.
To import license keys from a file	Do the following in the order listed: <ul style="list-style-type: none">■ Click Import from file.■ Select the besernum.xml file.
To install a trial version	Leave the license key field blank.

8 Click **Next**.

9 Select the agents and options that you want to install, and then click **Next**.

10 In the **Destination Folder** field, enter the location where you want to install Backup Exec.

11 Click **Next**.

12 Complete the service account credentials options as follows:

User Name	Type the user name for an Administrator account that the Backup Exec services can use. If the remote computer is in a domain, use a domain administrators account or an equivalent account that is part of the domain administrators group. If the remote computer is in a workgroup, use an Administrators account or an equivalent account that is part of the Administrators group on the computer.
Password	Type the password for an Administrator account that the Backup Exec services can use.
Domain	If the computer is in a domain, select the domain in which the computer is located. If the computer is in a workgroup, select the computer name.

13 Click **Next**.

- 14** Do one of the following to select a location on which to store the Backup Exec Database , and then click **Next**.

To create a local Backup Exec SQL Express instance

Do the following in the order listed:

- Click **Create a local Backup Exec SQL Express instance to store the database on**.
- To change the location of the database, type the new location in the **Destination Folder** field.

To use an existing SQL Server 2005 or SQL Server 2008 instance

Do the following in the order listed:

- Click **Use an existing instance of SQL Server 2005 (SP3a or later) or SQL Server 2008 on the network to store the database on**.
- Select the instance.

When Backup Exec is installed into an existing instance, the automated master database restore feature is not available. To recover the Master database, you must replace it with the Master database copy that Backup Exec automatically creates and updates when the Master database is backed up.

Caution: During the installation process and upgrade process, Backup Exec stops and starts the SQL service several times. Other user-created databases that use the SQL Server instance are unavailable during the process. To avoid such conflicts, you should install Backup Exec into its own SQL instance.

Backup Exec attempts to connect to the instance.

This step is skipped during upgrades.

- 15** Click **Next**.
- 16** Review the note about tape device drivers, and then click **Next**.
- 17** Click **Next**.
- 18** If you are prompted, enter information or choose settings for additional options being installed, and then click **Next** or **OK** after each selection.
- 19** After Backup Exec validates the remote computers, you can change the list in any of the following ways:

To manually add one remote computer	Click Add , and then click Add a Single Computer .
To manually add multiple remote computers	Click Add , and then click Add Multiple Computers with the Same Settings .
To add multiple remote computers by importing an existing list of computers	Click Import and Export , and then select one of the following options: <ul style="list-style-type: none">■ Select Import from File to enable Backup Exec to add the names of the remote computers from a selected list.■ Select Import Servers Published to this Media Server to enable Backup Exec to add the names of all the remote computers that are set up to publish to this media server. You must enter remote computer logon credentials for the list of remote computers.
To change the product that you selected to install or to change other properties you selected for this installation	Select the remote computer that you want to change, and then click Edit .
To delete a remote computer from the list	Select the remote computer that you want to delete, and then click Delete .
To save this list of remote computers and the associated remote computer logon credentials	Verify that Save the server list for future remote install sessions is checked. This option enables the names and the credentials of all of the remote computers to be added automatically the next time you install Backup Exec or options to these remote computers.
To save the list of remote computers to an XML file	Click Import and Export , and then click Export to File . You can select the location to save the Push_Export.xml file. This option is useful if you want to use the same list for multiple media servers. When you import the list, you must re-enter the remote computer logon credentials.
To fix the errors that were located during the validation	Right-click the name of the computer, and then click Fix Errors .

To enable Backup Exec to attempt to re-validate an invalid remote computer Right-click the name of the computer, and then click **Retry Validation**.

20 After all of the computers in the list are validated and the list is complete, click **Next**.

21 Read the Backup Exec installation review, and then click **Install**.

See [“About the installation log”](#) on page 166.

22 Click **Next**, and then click **Finish** to exit the wizard.

If you did not restart the remote computer, you may need to do it now in order for the configuration to take effect.

About installing Backup Exec options to remote computers

You can install the following options to remote computers:

- Remote Agent for Windows Systems
- Advanced Open File Option (AOFO)
- Desktop and Laptop Option (DLO) Desktop Agent
- Desktop and Laptop Maintenance Service

See [“Push-installing the Remote Agent and Advanced Open File Option to remote computers”](#) on page 134.

See [“Push-installing the Desktop Agent and DLO Maintenance Service from the media server to remote computers”](#) on page 137.

These features are push-installed to remote computers from a media server. Push installations save time by eliminating the need for local access at the target computer for the installation to be successful. You can install Backup Exec to as many as five remote computers concurrently.

There are special considerations that you should be familiar with before you install Backup Exec options on remote computers.

Table 2-9 Special considerations for installing Backup Exec options to remote computers

Item	Consideration
32-bit and 64-bit computers	If you try to push-install an option from a 32-bit computer to a 64-bit computer, you may be prompted to insert the 64-bit installation media.
Remote Agent for Windows Systems	<p>You cannot push-install the Remote Agent for Windows Systems when the remote computer is in the ForceGuest configuration and it is not in a domain. ForceGuest is an operating system configuration that limits incoming users to Guest-level access. Instead, use the installation media or the network to install the Remote Agent on the Windows computer.</p> <p>See “Installing Backup Exec using the command line (silent mode)” on page 153.</p> <p>You can also turn off ForceGuest. In Windows XP, ForceGuest is configured by the Use simple file sharing option. In Windows Vista, ForceGuest is configured by the Network Access: Sharing and security model for local accounts settings. Refer to your Microsoft Windows documentation for more information.</p> <p>Backup Exec installs a command line version of the Remote Agent on the computers that run the Server Core installation option of Windows Server 2008. The Remote Agent Utility command line applet is installed with the Remote Agent. This applet lets you monitor Backup Exec operations on the remote computer.</p> <p>See “Remote Agent Utility Command Line Applet switches” on page 1959.</p>
Terminal Services	If you install Backup Exec agents and options through Terminal Services and the installation media is on a shared drive (network share) you must install using a UNC path. Installation via mapped drives is not supported.

Table 2-9 Special considerations for installing Backup Exec options to remote computers *(continued)*

Item	Consideration
Windows XP SP2/Server 2003 SP1	<p>To push-install Backup Exec options to a Windows XP SP2/Server 2003 SP1 computer, you must enable File and Printer Sharing on the Windows Firewall Exceptions list for the following ports:</p> <ul style="list-style-type: none"> ■ 135 (RPC) ■ 445 (TCP) ■ 103X (mostly 1037) ■ 441 (RPC) <p>For more information about the Windows Firewall Exceptions list, refer to your Microsoft Windows documentation.</p> <p>During the installation process, Backup Exec sets the Remote Launch and Remote Access security permissions for the Administrator's group.</p> <p>You should enable the "Allow remote administration exception" group policy for the computer to which you push the installation.</p>
Windows Vista/Server 2008	<p>To push-install Backup Exec options to a computer that runs Windows Vista/Server 2008, you must enable certain items on the destination computer's Windows Firewall Exceptions list. You must enable the following items:</p> <ul style="list-style-type: none"> ■ File and Printer Sharing ■ Windows Management Instrumentation (WMI) <p>For more information refer to your Microsoft Windows documentation.</p> <p>To push-install to a Windows Vista computer, the destination computer must be part of a domain.</p> <p>For more information, refer to the Microsoft Knowledge Base.</p>
Symantec Endpoint Protection 11.0 or later	<p>To push-install options to a computer that runs Symantec Endpoint Protection (SEP) version 11.0 or later, you must configure SEP to share files and printers. File and printer sharing is turned off by default.</p>

Push-installing the Remote Agent and Advanced Open File Option to remote computers

You can install the following options to remote computers:

- Remote Agent for Windows Systems
- Advanced Open File Option (AOFO)

Before you install Backup Exec options on remote computers, review the special considerations.

See [“About installing Backup Exec options to remote computers”](#) on page 131.

The installation process creates an installation log named Bkupinst.htm on the computer where Backup Exec is installed .

See [“About the installation log”](#) on page 166.

If there are problems installing the Backup Exec Remote Agent using this method, you can try to manually install the Remote Agent.

See [“Using a command prompt to install the Remote Agent on a remote computer”](#) on page 145.

To push-install the Remote Agent and AOFO to remote computers

1 Do one of the following:

To push-install Backup Exec options to remote computers from the installation media

Do the following steps in the order listed:

- From the installation media browser, click **Installation**, and then click **Backup Exec**.
- On the **Welcome** panel, select **I accept the terms of the license agreement**, and then click **Next**.
- Click **Custom installation**.
- Uncheck **Local Installation**, and then check **Remote Installation**.
- Click **Next**.
- On the **Remote Computers** panel, click **Add**.
- To install Backup Exec on one remote computer, select **Add a Single Computer**, or to install Backup Exec on multiple computers using the same settings, select **Add Multiple Computers with the Same Settings**.

Push-installing the Remote Agent and Advanced Open File Option to remote computers

To push-install Backup Exec options to remote computers from the Backup Exec media server On the **Tools** menu, click **Install Agents and Media Servers on Other Servers**.

- 2 Select **Remote Agent for Windows Systems**, and then click **Next**.
- 3 Type the fully qualified name of the remote computer or click **Browse** to locate the remote computer.
- 4 Click **Add to List**, and then repeat steps 3 and 4 for each remote computer to which you want to push-install the options.

If you are push-installing from the installation media and you selected **Add a Single Computer** in step 1, you can skip this step.

- 5 Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.

You must use Administrator credentials.

- 6 Click **Next**.
- 7 Select **Advanced Open File Option** if you want to install it with the Remote Agent for Windows Systems.
- 8 In the **Destination Folder** field, enter the path where you want to install the files.
- 9 Click **Next**.
- 10 Verify that the option to enable the remote agent to publish information to the media servers is selected.
- 11 Verify that the media servers to which you want to publish are listed. You can add, edit, or remove media servers.
- 12 Click **Next**.
- 13 After Backup Exec validates the remote computers, you can change the list in any of the following ways:

To manually add one remote computer Click **Add**, and then click **Add a Single Computer**.

To manually add multiple remote computers Click **Add**, and then click **Add Multiple Computers with the Same Settings**.

To add multiple remote computers by importing an existing list of computers	Click Import and Export , and then select one of the following options: <ul style="list-style-type: none">■ Select Import from File to enable Backup Exec to add the names of the remote computers from a selected list.■ Select Import Servers Published to this Media Server to enable Backup Exec to add the names of all the remote computers that are set up to publish to this media server. You must enter remote computer logon credentials for the list of remote computers.
To change the product that you selected to install or to change other properties you selected for this installation	Select the remote computer that you want to change, and then click Edit .
To delete a remote computer from the list	Select the remote computer that you want to delete, and then click Delete .
To save this list of remote computers and the associated remote computer logon credentials	Verify that Save the server list for future remote install sessions is checked. This option enables the names of all of the remote computers and their credentials to be added automatically the next time you want to install Backup Exec or options to these remote computers.
To save the list of remote computers to an XML file	Click Import and Export , and then click Export to File . You can select the location to save the XML file. This option is useful if you want to use the same list for multiple media servers. When you import the list, you must re-enter the remote computer logon credentials.
To fix the errors that were located during the validation	Right-click the name of the computer, and then click Fix Errors .
To enable Backup Exec to attempt to re-validate an invalid remote computer	Right-click the name of the computer, and then click Retry Validation .

14 After all of the computers in the list are validated and the list is complete, click **Next**.

15 Read the Backup Exec installation review, and then click **Install**.

See “[About the installation log](#)” on page 166.

16 Click **Next**, and then click **Finish** to exit the wizard.

If you did not restart the remote computer, you may need to do it now in order for the configuration to take effect.

Push-installing the Desktop Agent and DLO Maintenance Service from the media server to remote computers

You can install the following options to remote computers:

- Desktop Agent (This agent can be push-installed to a remote computer only if the Desktop and Laptop Option is already installed on the Backup Exec media server.)
- DLO Maintenance Service

To push-install the Desktop Agent and DLO Maintenance Service to remote computers

- 1 On the **Tools** menu, click **Install Agents and Media Servers on Other Servers**.
- 2 On the **Remote Computers** panel, click **Add**.
- 3 To install on one remote computer, select **Add a Single Computer**, or to install on multiple computers using the same settings, select **Add Multiple Computers with the Same Settings**.
- 4 Select **Desktop and Laptop Agent** or **DLO Maintenance Service**, and then click **Next**.
- 5 Type the fully qualified name of the remote computer or click **Browse Remote Computers** to locate the remote computer.
- 6 Click **Add to List**, and then repeat steps 5 and 6 for each remote computer to which you want to push-install the options.

This step does not apply if you selected **Add a Single Computer** in step 4.

- 7 Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.

You must use Administrator credentials.

- 8 Click **Next**.

9 After Backup Exec validates the remote computers, you can change the list in any of the following ways:

To manually add one remote computer	Click Add , and then click Add a Single Computer .
To manually add multiple remote computers	Click Add , and then click Add Multiple Computers with the Same Settings .
To add multiple remote computers by importing an existing list of computers	Click Import and Export , and then select Import from File to enable Backup Exec to add the names of the remote computers from a selected list. You must enter remote computer logon credentials for the list of remote computers.
To change the product that you selected to install or to change other properties you selected for this installation	Select the remote computer that you want to change, and then click Edit .
To delete a remote computer from the list	Select the remote computer that you want to delete, and then click Delete .
To save this list of remote computers and the associated remote computer logon credentials	Verify that Save the server list for future remote install sessions is checked. This option enables the names of all of the remote computers and their credentials to be added automatically the next time you want to install Backup Exec or options to these remote computers.
To save the list of remote computers to an XML file	Click Import and Export , and then click Export to File . You can select the location to save the XML file. This option is useful if you want to use the same list for multiple media servers. When you import the list, you must re-enter the remote computer logon credentials.
To fix the errors that were located during the validation	Right-click the name of the computer, and then click Fix Errors .
To enable Backup Exec to attempt to re-validate an invalid remote computer	Right-click the name of the computer, and then click Retry Validation .

- 10 After all of the computers in the list are validated and the list is complete, click **Next**.
- 11 Read the Backup Exec installation review, and then click **Install**.
See [“About the installation log”](#) on page 166.
- 12 Click **Next**, and then click **Finish** to exit the wizard.
If you did not restart the remote computer, you may need to do it now in order for the configuration to take effect.

About installing the Remote Agent for Windows Systems

You can install the Remote Agent for Windows Systems by using the following methods, depending on your environment:

- Install the Remote Agent from the Backup Exec installation media by taking the media to the computer and running the Backup Exec installation program.
See [“Installing additional Backup Exec options to the local media server”](#) on page 123.
- Push-install the Remote Agent and the Advanced Open File Option (AOFO) to one or more remote computers from the media server.
See [“Push-installing the Remote Agent and Advanced Open File Option to remote computers”](#) on page 134.
- Push-install the Remote Agent and AOFO to a remote computer that is displayed in the backup selections list.
See [“Installing the Remote Agent and the Advanced Open File Option to a remote computer in the backup selections list”](#) on page 140.
- Use a Microsoft Active Directory network to centrally manage the installation of the Remote Agent and the AOFO to computers in the network.
See [“How to install the Remote Agent and Advanced Open File Option in an Active Directory network”](#) on page 140.
- Install the Remote Agent and AOFO by using command script files.
See [“Using a command script to install the Remote Agent and AOFO ”](#) on page 148.

There are special considerations for installing the Remote Agent.

See [“About installing Backup Exec options to remote computers”](#) on page 131.

Installing the Remote Agent and the Advanced Open File Option to a remote computer in the backup selections list

As you are making backup selections, you can install the Remote Agent and the Advanced Open File Option (AOFO) on computers that you want to back up.

Note: AOFO can only be installed on 32-bit computers that run Windows 2000/XP.

To install the Backup Exec Remote Agent and the Advanced Open File Option to a remote computer in the backup selections list

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the **View by Resource** tab, do one of the following:

If the remote computer has been added to Favorite Resources Do the following in the order listed:

- Expand **Favorite Resources**.
- Expand **Windows Systems**.

If the remote computer has not been added to Favorite Resources Do the following in the order listed:

- Expand **Domains**.
- Expand **Microsoft Windows Network**.
- Expand the appropriate domain.

- 4 Right-click the computer that you want to install the Remote Agent and the AOFO to, and then click **Install Remote Agent/Advanced Open File Option**.
- 5 Use the installation wizard to complete the installation.

See “[Push-installing the Remote Agent and Advanced Open File Option to remote computers](#)” on page 134.

How to install the Remote Agent and Advanced Open File Option in an Active Directory network

You can centrally manage the installation of the Backup Exec Remote Agent for Windows Systems and the Advanced Open File Option (AOFO) to computers in an Active Directory network. You configure the installation once, and then use a Group Policy Object to assign that installation to computers in an Organizational Unit. The options are installed automatically whenever a computer in the Organizational Unit is started.

Note: Review your organization’s deployment plans before you implement a rollout of the Backup Exec Remote Agent and Advanced Open File Option to client computers. You should also review your Group Policy Desktop Management and Active Directory documentation.

Note: AOFO can only be installed on 32-bit computers that run Windows 2000/XP.

Table 2-10 Installing the Remote Agent and Advanced Open File Option in an Active Directory network

Action	Description
<p>Create a transform for the Remote Agent and/or AOFO.</p> <p>See “Creating a transform” on page 142.</p>	<p>A transform contains changes that you want to make to the Remote Agent’s Windows Installer package when a computer starts, such as installation path, which computers to publish to, and whether to install AOFO. You must create separate transforms for 32-bit computers and 64-bit computers.</p> <p>Requirements to create a transform are as follows:</p> <ul style="list-style-type: none"> ■ The computer on which you want to create the transform must have Microsoft Windows 2000 or later. ■ The computers on which you want to install the Remote Agent must be running MSI 3.1. ■ The computers on which you want to install the Remote Agent must be running MSXML 6.0. ■ Only assignment to computers is supported. Assignment to users is not supported.
<p>Create a distribution point (share) that contains the source file of the Remote Agent that you want to install.</p> <p>See “Creating a software distribution point (share)” on page 143.</p>	<p>You must copy the transform that you create, and the Backup Exec RAWS32 or RAWSX64 directory, to the distribution point.</p>

Table 2-10 Installing the Remote Agent and Advanced Open File Option in an Active Directory network (*continued*)

Action	Description
Configure a Group Policy Object to assign the transform and the RAWS32 or RAWSX64 directory in the distribution point to computers in an Active Directory Organizational Unit. See “Configuring a Group Policy Object” on page 144.	The software is installed automatically when the computers in the Organizational Unit are started.

Creating a transform

To install the Remote Agent and the Advanced Open File Option in an Active Directory network, you must create a transform.

See [“How to install the Remote Agent and Advanced Open File Option in an Active Directory network”](#) on page 140.

To create the transform

- 1 Do one of the following:
 - From the Backup Exec installation media browser, click **Installation**, and then click **Backup Exec Remote Agent**.
 - From a media server on which Backup Exec is installed, go to \Program Files\Symantec\Backup Exec\Agents\RAWS32 and double-click **Setup.exe**.
- 2 On the Welcome panel, click **Next**.
- 3 On the Install Type panel, click **Create a Transform to use Active Directory to install the Remote Agent**, and then click **Next**.
- 4 On the Install Option panel, do the following:
 - Select the options that you want to include in the transform.
 The configuration that you specify in the transform becomes the default setting for a client computer in the Active Directory network when it performs setup.
 - Enter the path where the Remote Agent will be installed on client computers
 To change the default path, click **Change**.
 The path should not be a removable drive or a network drive.
- 5 Click **Next**.

- 6 Do the following in the order listed:
 - Verify that the option **Enable the Remote Agent to publish the IP address and name of the remote computer and the version of the Remote Agent to media servers** is selected.
 - Click **Add** to enter the media server name or IP address of all of the media servers that you want the Remote Agent to publish to after the transform has been applied.
- 7 Click **Next**.
 The computer that the Remote Agent is installed on is displayed in the media server’s backup selection tree under **Favorite Resources**.
- 8 Enter a file name and a path where the transform will be created, and then click **Next**.
 To change the default path, click **Change**.
 Use a meaningful file name for the transform. For example, the name could include the names of the options in the transform and the platform you plan to apply the transform to, such as RemoteAgentDefaultPathNoPublishing.
- 9 To create the transform, click **Install**.
- 10 After the transform is created, set up a distribution point for the source files.
 See [“Creating a software distribution point \(share\)”](#) on page 143.

Creating a software distribution point (share)

To install the Remote Agent and the Advanced Open File Option in an Active Directory network, you must create a software distribution point after you create a transform.

See [“Creating a transform”](#) on page 142.

See [“How to install the Remote Agent and Advanced Open File Option in an Active Directory network”](#) on page 140.

Table 2-11 How to create a software distribution point (share)

Step	Description
Step 1	Create a shared folder, and then set permissions so that client computers that will run the installation have access to the shared folder.

Table 2-11 How to create a software distribution point (share) *(continued)*

Step	Description
Step 2	Copy the following directories from the media server to the shared folder: <ul style="list-style-type: none"> ■ RAWS32 or RAWSX64 ■ MSXML By default, these folders are located in \Program Files\Symantec\Backup Exec\Agents.
Step 3	Copy the transform from the path where it was created to the RAWS32 or RAWSX64 directory on the shared folder.
Step 4	Configure a Group Policy Object to deploy the source files. See “Configuring a Group Policy Object” on page 144.

Configuring a Group Policy Object

To install the Remote Agent and the Advanced Open File Option in an Active Directory network, you must configure a Group Policy Object after you create a software distribution point and create a transform.

See [“Creating a transform”](#) on page 142.

See [“Creating a software distribution point \(share\)”](#) on page 143.

See [“How to install the Remote Agent and Advanced Open File Option in an Active Directory network”](#) on page 140.

To configure a Group Policy Object to deploy the software

- 1 From the Active Directory snap-in that manages users and groups, click **Properties**, and create a new Group Policy Object or edit an existing one.
 Refer to your Microsoft Windows documentation for information on creating a Group Policy Object.
- 2 Under Computer Configuration, expand **Software Settings**.
- 3 Right-click **Software Installation**, click **New**, and then click **Package**.
- 4 On the File Open dialog box, browse to the software distribution point by using the Universal Naming Convention (UNC) name, for example, \\server name\share name, select the package file, and then click **Open**.

- 5 Select the package file **Symantec Backup Exec Remote Agent for Windows Systems.msi**, and then click **Open**.
- 6 When prompted, apply the **Advanced Option**.
- 7 After Active Directory checks the msi package, on the General Properties tab, make sure the correct versions of the options are being installed.
- 8 On the **Deployment** tab, set up the configuration for your environment.
Make sure the option **Make this 32-bit x86 application available to WIN64 machines** is not selected.
If you want the Remote Agent to be uninstalled if the computer is removed from the Organization Unit, select the option **Uninstall this application when it falls out of the scope of management**.
- 9 On the **Modifications** tab, click **Add**, browse to the share, and select the transform that you created.
- 10 Select **Open**, and make any other changes that are necessary, and then click **OK**.
- 11 Close all of the dialog boxes.
When a computer in the Organizational Unit that you specified is started, the transform is processed and the options that you specified are installed.
- 12 View the installation log that is created on the destination computers to verify the installation of the Remote Agent and/or the AOFO.

Using a command prompt to install the Remote Agent on a remote computer

You can install the Remote Agent by using a command prompt.

The installation process creates an installation log named RAWSinstr.htm.

See [“About the installation log”](#) on page 166.

To use a command prompt to install the Remote Agent on a remote computer

- 1 At a remote computer, map a drive letter to the Backup Exec media server Agents directory. By default, the Agents directory is located at the following path:

`\Program Files\Symantec\Backup Exec\Agents`

or you can copy the following folders to the same local directory:

To install to a 32-bit computer: RAWS32 and MSXML folders

To install to a 64-bit computer: RAWSX64 and MSXML folders

- 2 Open a command prompt and type the drive letter that you mapped in step 1 and the following path:

To install to a 32-bit computer: `\RAWS32 :`

To install to a 64-bit computer: `\RAWSX64 :`

3 Do one of the following:

- | | |
|--|--|
| To install the Remote Agent to a 32-bit computer without publishing enabled: | Run the following command:

<code>setup.exe /RANT32: /S: -boot</code> |
| To install the Remote Agent to a 32-bit computer with publishing enabled: | Run the following command:

<code>setup.exe /RANT32: /S: /ADVRT:
<media server name 1> <media
server name 2></code> |
| To install the Remote Agent to a 64-bit computer without publishing enabled: | Run the following command:

<code>setup.exe /RAWSX64: /S: -boot</code> |
| To install the Remote Agent to a 64-bit computer with publishing enabled: | Run the following command:

<code>setup.exe /RAWSX64: /S: /ADVRT:
<media server name 1> <media
server name 2></code> |

The Remote Agent is installed on the remote computer in the following directory:

- | | |
|---|---|
| If you installed the Remote Agent to a 32-bit computer: | <code>\Program Files\Symantec\Backup Exec\RAWS</code> |
| If you installed the Remote Agent to a 64-bit computer: | <code>\Program Files\Symantec\Backup Exec\RAWS</code> |

Using a command prompt to uninstall the Remote Agent from a remote computer

You can uninstall the Remote Agent by using a command prompt.

To use a command prompt to uninstall the Remote Agent from a remote computer

- 1 At the remote computer, map a drive letter to the Backup Exec media server Remote Agent directory using the following path:

To uninstall the Remote Agent from a 32-bit computer: `\Program Files\Symantec\Backup Exec\Agents\RAWS32`

To uninstall the Remote Agent from a 64-bit computer: `\Program Files\Symantec\Backup Exec\Agents\RAWSX64`

- 2 Open a command prompt, and then type the drive letter that you mapped in step 1.
- 3 Run the following command:

To uninstall the Remote Agent from a 32-bit computer: `setup.exe /RANT32: /S: -u`

The `/S:` parameter is used to run the operation in silent mode, without the benefit of a user interface. The `-u` parameter specifies an uninstall operation.

To uninstall the Remote Agent from a 64-bit computer: `setup.exe /RAWSX64: /S: -u`

See [“Using a command prompt to install the Remote Agent on a remote computer”](#) on page 145.

Using a command script to install the Remote Agent and AOFO

You can use command script files to install the Remote Agent and the Advanced Open File Option (AOFO). The command script files are included in the Remote Agent installation directory.

The installation process creates an installation log named `RAWSinst.htm`.

See [“About the installation log”](#) on page 166.

To use a command script to install the Remote Agent and AOFO

- 1 Map a drive letter to the Agents directory of a Backup Exec media server. By default, the Agents directory is located at the following path:

`\Program Files\Symantec\Backup Exec\Agents`

- 2 Do one of the following:

To install the Remote Agent on a 32-bit computer Double-click **setupaa** in the RAWS32 directory.

To install the Advanced Open File Option on a 32-bit computer Double-click **setupafo** in the RAWS32 directory. By default, the command script installs the option automatically on the remote server in the following directory:

`\Program Files\Symantec\Backup Exec\RAWS`

To install the Remote Agent on a 64-bit computer Double-click **setupaax64** in the RAWX64 directory.

- 3 If you installed the Advanced Open File Option, you must restart the remote computer.

Using a command script to uninstall the Remote Agent and AOFO

One command script file is available to uninstall both the Remote Agent and the AOFO. The uninstall command script removes both options together. You cannot remove the options separately using the command script.

To use a command script to uninstall the Remote Agent and Advanced Open File Option

- 1 Map a drive letter to the Backup Exec media server by using one of the following paths:

To a 32-bit computer \Program Files\Symantec\Backup
Exec\Agents\RAWS32

To a 64-bit computer \Program Files\Symantec\Backup
Exec\Agents\RAWSX64

- 2 Do one of the following:

For a 32-bit computer Double-click **Removeafofo**.

For a 64-bit computer Double-click **Uninstallafofo64**.

Note: This script applies only to the Remote Agent for Windows Systems 12.5.

Both the Remote Agent and the Advanced Open File Option are removed from the computer.

- 3 Restart the remote computer.

See [“Using a command script to install the Remote Agent and AOFO”](#) on page 148.

Installing the Remote Administrator

The Remote Administrator lets you administer the media server from a remote Windows server or workstation. To support the Remote Administrator, the Backup Exec system services must be running on the media server that you want to administer.

To install the Remote Administrator

- 1 From the installation media browser, click **Installation**.
- 2 Click **Backup Exec**.
- 3 On the **Welcome** panel, select **I accept the terms of the license agreement**, and then click **Next**.
- 4 On the **Installation Type** panel, select **Custom installation**, and then click **Next**.
- 5 To install the Administration Console as a Remote Administrator, click **Install Remote Administration Console only**, and then click **Next**.

- 6 To change the location where the files are installed, click **Change** to select another directory for the installation.
- 7 Click **Next**.
- 8 Review the installation summary, and then click **Install**.
- 9 Click **Finish**.

See [“Running the Remote Administrator”](#) on page 151.

Running the Remote Administrator

The Remote Administrator lets you administer the media server from a remote Windows server or workstation. To support the Remote Administrator, the media server requires that the Backup Exec system services must be running.

You may be prompted for a user name and password to browse some network shares even if you are logged into the Remote Administrator computer under an account that is valid for those shares. Provide a domain-qualified user name and password when prompted (for example, domain1\howard).

For workgroup accounts, when logging in between different workgroups, you can provide only a user ID when prompted, and leave the workgroup line blank.

See [“Installing the Remote Administrator”](#) on page 150.

To run the Remote Administrator

- 1 Click **Start**.
- 2 Point to Programs, and then click **Symantec Backup Exec**.

If you are connecting to a remote administration console from a media server, on the **Network** menu, click **Connect to Local Media Server** to break the connection. Click **Connect to Media Server** to connect to another media server.

- 3 Select the appropriate options.

See [“Connect to Media Server options”](#) on page 151.

The status of the local services appears at the bottom of this dialog box. If you try to connect to a server and the connection fails, this dialog box displays the services status for the server you attempted to connect to.

- 4 Click **OK**.

Connect to Media Server options

On this dialog box, you can enter the credentials that are required to administer a media server from a remote Windows server or workstation.

See [“Running the Remote Administrator”](#) on page 151.

Table 2-12 **Connect to Media Server options**

Item	Description
Server	<p>Indicates the name of the media server. You can select the name from the list or type the name of the server if you are running the Remote Administrator from a media server.</p> <p>Each server in the domain that has Backup Exec installed automatically appears in the list box.</p>
Low speed connection (RAS)	<p>Minimizes the amount of information initially retrieved from the media server to which you are connected. When this option is selected, views such as the device view and the media view do not expand automatically when the Administration Console is loaded. This option reduces the time that is required to connect to the remote media server. Information for each view is updated when the view is selected.</p> <p>This option is useful if you connect to the media server over a modem line.</p>
User name	<p>Indicates an administrator user name for the server to which you want to connect.</p> <p>You cannot log on to the remote administration console with a user name that has a blank password on Windows Server 2003/2008 and XP/Vista computers. You must configure Windows to allow blank passwords. Otherwise, the error message "Logon failure: user account restriction" appears. For more information, see your Windows documentation.</p>
Password	<p>Indicates the password for the user.</p>
Domain	<p>Indicates the domain to which the user belongs. You can select the domain from the list or type the domain name.</p>

- 4 Type `setup /TS:` and the appropriate switches. For example:

```
setup /TS: /USER:<user> /DOM:domain /PASS:password /SNO:license key /S:
```

See [“Command line switches for silent mode installation of Backup Exec”](#) on page 154.

If you use the command line switches without the /S: switch, the Backup Exec installation program launches with the command line parameters as defaults for the installation options. For example, if /S: had been left in the above example, the Backup Exec installation program launches with the user name, domain, password, and license key appearing on the installation dialog boxes.

- 5 Press **Enter**.

Command line switches for silent mode installation of Backup Exec

The command line switches used for silent mode installation of Backup Exec are described in the following table.

Note the following general rules for using these switches:

- Substitute values appropriate for your environment for values shown in italics; for example substitute your password for *password*.
- Enclose the value in quotation marks if it contains spaces, such as "Operations Weekly Backup".

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 153.

Table 2-13 Command line switches for silent mode installation of Backup Exec

Switch	Additional Switches	Description
/TS:		Installs Backup Exec using the options specified with the additional switches. The /USER:" <i>user</i> " /DOM:" <i>dm</i> " /PASS:" <i>pw</i> " is required.

Table 2-13 Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/USER:" <i>user</i> " /DOM:" <i>dm</i> " /PASS:" <i>pw</i> "	<p>Required. Specifies an existing user, domain, and password for the Backup Exec system service account. Silent mode installation will not create a user.</p> <p>Note: When using /PASS:, if a quote is needed as part of the password, specify it as \". For example, if the password is pass"word, type it as /PASS:pass\"word. If the characters \" are used as part of the password, you must precede each character with a \. For example, if the password is pass\"word, type it as /PASS:pass\\\"word.</p>
	/DEST:" <i>path</i> "	<p>Specifies the path where Backup Exec will be installed. Otherwise, the default path Program Files\Symantec\Backup Exec is used.</p>
	/DOCS:	Installs online documentation.

Table 2-13 Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/NOINSTALL:	Allows you to select all install options without actually installing the Backup Exec software. This option can be used in conjunction with the /CPF: switch.
	/SNO: <i>license key</i>	<p>Specifies one or more license keys to use for installing Backup Exec and additional options. License keys are not required to install the Remote Administrator. You may specify up to 99 license keys. If none are specified, then a trial copy of Backup Exec is installed.</p> <p>The following examples show how the /SNO switch can be used:</p> <p>/SNO:<i>s1</i></p> <p>/SNO:<i>s1 s2 s3 s4</i></p> <p>Note: If you install a license for an option or agent, you must also type a switch that specifies the option or agent. The switches that specify an option or agent are included in this table.</p>

Table 2-13 Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/TD:NEW or ALL	<p>/TD:NEW installs tape drivers only for drives that do not have drivers loaded.</p> <p>/TD:ALL installs tape drivers for all drives.</p> <p>Note: To install the Symantec tape drivers, the Windows driver signing policy must be set to Ignore. See your Microsoft Windows documentation for instructions on changing the driver signing policy.</p>
	/CPF:"path\filename.cpf"	Creates a file containing all of the installation parameters provided. Note that the file is not encrypted, which exposes parameters.
	/DBSERVER:<server\instance>	Installs the Backup Exec database to the specified SQL server.
	/DBINSTPATH: <SQL Express destination folder>	Installs the default instance of SQL Express in the specified folder.
	/NOUPDATE:	Skips the installation of Symantec LiveUpdate.

Table 2-13 Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/DISADVRT	Installs the Remote Agent without publishing it.
	/SQLXSETUP:<SQL Express Install Package>	Specifies the location of the language-specific install package for Microsoft SQL Server 2005 Express Edition.
	/LOADER:	Installs the Library Expansion Option.
	/IDR:	Installs the Intelligent Disaster Recovery Option.
	/AOFO:	Installs the Advanced Open File Option.
	/DLO:	Installs the Backup Exec Desktop and Laptop Option.
	/DLO5:	Installs the five-user version of the Backup Exec Desktop and Laptop Option.
	/MMS:<CAS server name>	Creates a managed media server for use with the Central Admin Server Option.

Table 2-13 Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/CASOPVLOCAL: <1 or 0>	/CASOPVLOCAL:<1> indicates that device and media data will be stored locally on the managed media server. Use this switch with /MMS:. /CASOPVLOCAL:<0> indicates that device and media data will be stored on the central administration server. Use this switch with /MMS:.
	/R3:	Installs the Backup Exec Agent for SAP Applications (SAP Agent).
	/SSO:	Installs the SAN Shared Storage Option with this server as the primary server.
	/SSO:server name	Installs the SAN Shared Storage Option with this server as the secondary and the <server name> as the primary.
	/SHAREPT:	Installs the Agent for Microsoft SharePoint.
	/EXCH:	Installs the Agent for Microsoft Exchange Server.

Table 2-13 Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/LOTUS:	Installs the Agent for Lotus Domino.
	/ORACLE:	Installs the Agent for Oracle on Windows or Linux Servers.
	/SQL:	Installs the Agent for Microsoft SQL Server.
	/EV:	Installs the Agent for Enterprise Vault.
	/NTA:	Installs the Remote Agent for Windows Systems.
	/ADBO:	Installs the Advanced Disk-based Backup Option.
	/CASO:	Installs the Central Admin Server Option.
	/ADR:	Installs the Active Directory Recovery Agent.
	/NDMP:	Installs the NDMP Option.
	/DB2:	Installs the Agent for DB2.
	/MAC:	Installs the Remote Agent for Macintosh Servers.
	/RAULUS:	Installs the Remote Agent for Linux or UNIX Servers.

Table 2-13 Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/VRTSRV:	Installs the Agent for Microsoft Virtual Server.
	/VMWARE:	Installs the Agent for VMware Virtual Infrastructure.
	/STORPROV:	Installs the Storage Provisioning Option.
	/DEDUPE:	Installs the Deduplication Option.
	/EXCHARCH:	Installs the Exchange Mailbox Archiving Option.
	/NTFS:	Installs the File System Archiving Option.
	/VTL:	Installs the Virtual Tape Library Unlimited Drive Option.
	/FIXEDSPO:	Installs the Storage Provisioning Option - Basic.
	/RMAL:	Installs the Remote Media Agent for Linux Servers.
	/COPYCONFIG:	Installs the Copy Server Configuration option.
-?		Provides help on all command line operations, usage, and special switches.

Installing the Remote Administrator using the command line

You can also use Silent Mode Installation to install the Remote Administrator. Options for the Remote Administrator are specified with the use of additional command switches.

To install the Remote Administrator using the command line

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change directories to one of the following:

For 32-bit computers `\be\winnt\install\be32`

For 64-bit computers `\be\winnt\install\bex64`

- 4 Type `setup /RA:` and the appropriate switches. For example:

```
setup /RA: /S:
```

The command line switches used for silent mode installation of the Remote Administrator are described in the following table.

Remember the following general rules for using these switches:

- Substitute values appropriate for your environment for values in italics; for example, substitute your password for *password*.
- Enclose the value in quotation marks if it contains spaces, such as "Program Files\Symantec\Backup Exec".

Table 2-14 Command line switches for Remote Administrator silent mode installation

Switch	Additional Switches	Description
/RA:		Installs Remote Administrator using the options specified with the additional switches.

Table 2-14 Command line switches for Remote Administrator silent mode installation (*continued*)

Switch	Additional Switches	Description
	<code>/DEST:"path"</code>	Specifies the path where Remote Administrator will be installed. Otherwise, the default path <code>Program Files\Symantec\Backup Exec</code> is used.
	<code>/DOCS:</code>	Installs online documentation.
	<code>/NOINSTALL:</code>	Allows you to select all install options without actually installing the Backup Exec software. This option can be used with the <code>/CPF:</code> switch.
	<code>/CPF:"path\filename.cpf"</code>	Creates a file containing all of the installation parameters provided. Note that the file is not encrypted, which exposes parameters such as the password
<code>-?</code>		Provides help on all command line operations, usage, and special switches.

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 153.

Uninstalling Backup Exec using the command line

If Backup Exec is already installed, you can use the setup.exe program to uninstall Backup Exec program files and Backup Exec data.

To uninstall Backup Exec using the command line

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change directories to one of the following:

For 32-bit computers `\be\winnt\install\be32`

For 64-bit computers `\be\winnt\install\bex64`

- 4 To remove the Backup Exec program files but keep all of the Backup Exec data, type:

```
SETUP /UNINSTALL:
```

To remove the Backup Exec program files and the Backup Exec data, type:

```
SETUP /REMOVEALL:
```

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 153.

Creating installation parameter files

If you use the command line switches without the /S: switch, the Backup Exec installation program launches with the command line parameters as defaults for the installation options. For example, suppose you type:

```
SETUP /TS: /USER:user /DOM:domain /PASS:password /SNO:license  
key
```

The Backup Exec installation program is launched. The screens that allow you to enter the logon credentials and the license key will appear with the information you provided on the command line.

You can also use the /CPF: command to create a parameter file that contains all of the command line options you provided. This parameter file can then be used to provide the options for installing either Backup Exec or the Remote Administrator. Note that the file is not encrypted, which exposes parameters such as the password.

To create installation parameter files

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change directories to one of the following:

For 32-bit computers \be\winnt\install\be32

For 64-bit computers \be\winnt\install\bex64

- 4 Type `setup /TS:` and the appropriate switches, including `/CPF:` and the full path name of the parameter file. For example, type:

```
setup /TS: /USER:user /DOM:domain /PASS:password /SNO:license
key /CPF:"A:\file name" /S:
```

Backup Exec will be installed on your server and a parameter file containing the user name, domain, password, and license key will be saved to a floppy diskette. You can use this parameter file to install to another computer.

See [“Using installation parameter files”](#) on page 165.

Using installation parameter files

You can use the `/CPF:` command to create a parameter file that contains all of the command line options you provided. This parameter file can then be used to provide the options for installing either Backup Exec or the Remote Administrator.

See [“Creating installation parameter files”](#) on page 164.

To use installation parameter files

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change directories to `\WINNT\INSTALL\BE`.
- 4 Type:`SETUP /PARAMS:"A:\file name" /S:`
- 5 If you want to overwrite a parameter, specify the new parameter. For example, to change the password, type:`SETUP /PARAMS:"A:\file name" /PASS:new password/S:`

Installing a trial version of Backup Exec agents and options

You can install a trial version of most Backup Exec agents and options at any time after the core product is licensed. Each agent and each option has its own independent trial period. When a trial period is about to expire, Backup Exec warns you with an alert.

You can view a list of agents and options that are available for a trial period. You can also view the amount of time that is left in each individual trial period.

See “[Viewing license information](#)” on page 174.

To install a trial version of Backup Exec agents and options

- 1 On the **Tools** menu, click **Install Options and License Keys on this Media Server**.
- 2 Click **Next**.
- 3 Select the agents or options you want to evaluate.
- 4 Click **Next**.
- 5 If you are prompted, enter information or choose settings for the additional options that you want to install. Click **Next** after each selection.
- 6 Read the Backup Exec installation review, and then click **Install**.

The Backup Exec services are stopped while the additional options are installed. If any active jobs are in progress, you are prompted to stop them, or to wait for the jobs to finish.

When the installation is complete, the services restart.

- 7 Click **Finish**.

About the installation log

Backup Exec creates an installation log file, named Bkupinst.htm, when you install Backup Exec and when you install patches. This log file can help you troubleshoot installation problems. The log file provides links to TechNotes for the most common errors. If you install the Remote Agent for Windows Systems, a log file called RAWSinst.htm is also created.

In addition, the text in the log file uses the following colors so you can identify warnings and errors:

Table 2-15 Installation log colors

This color	Indicates
Black	Normal operations
Orange	Warning messages
Red	Error messages

For most versions of Windows, Bkupinst.htm is saved in:

```
%allusersprofile%\Application Data\Symantec\Backup Exec\Logs.
```

The Application Data folder is a hidden folder. If you do not see the Application Data folder, refer to the Microsoft Windows documentation for instructions on how to display hidden folders.

Repairing Backup Exec

If you have missing or corrupted Backup Exec files or registry keys on the local media server, run the Repair option. The program stops all Backup Exec services, reinstalls corrupted files and registry keys, reinstalls tape devices (stand-alone drives and libraries), and restarts the services. The database is not reinstalled.

Any changes made to Backup Exec program files and registry keys will be reset to the original settings.

To repair Backup Exec

- 1 Close the Backup Exec application.
- 2 From the Windows Control Panel, double-click **Add or Remove Programs**.
- 3 At the **Add or Remove Programs** dialog box, select **Symantec Backup Exec**, and then click **Change**.
- 4 On the **Welcome** panel, click **Next**.
- 5 Select **Local installation** and **Repair**, and then click **Next**.
Make sure the option **Remote installation** is not selected.
- 6 If you are prompted to enter credentials for the Backup Exec service account, type the correct credentials, and then click **Next**.
- 7 Select **Install**.
- 8 Click **Finish**.

Starting and stopping Backup Exec services

You can start, stop, and restart Backup Exec services.

To start or stop Backup Exec services

- 1 On the **Tools** menu, click **Backup Exec Services**.
- 2 Select the appropriate options.

See “[Backup Exec Services Manager options](#)” on page 168.

Backup Exec Services Manager options

You can start, stop, and restart Backup Exec services.

See “[Starting and stopping Backup Exec services](#)” on page 168.

Table 2-16 Backup Exec Services Manager options

Item	Description
Server	Indicates the name of a server for which you want to start, stop, or restart services. You can type the name of a server or import a list of servers.
Add	Enables you to add the name of a server for which you want to start, stop, or restart services.
Import List	Enables you to import a list of servers for which you want to start, stop, or restart services.
Start all services	Starts all Backup Exec services for the selected server.
Stop all services	Stops all Backup Exec services for the selected server.
Restart all services	Stops all Backup Exec services and then restart the services for the selected server.
Services credentials	Changes service account information or changes startup options.
Refresh	Refreshes this dialog box.
Remove Servers	Removes a selected server or servers from the server list.

Uninstalling Backup Exec

Use Microsoft’s Add or Remove Programs option to remove Backup Exec from a computer. For additional information on Add or Remove Programs, refer to your Microsoft documentation.

Uninstalling Backup Exec also removes Symantec tape class drivers. If you reinstall Backup Exec and want to use Symantec tape class drivers, you must reinstall them.

Note: When using the Shared Storage Option, you must uninstall Backup Exec from the secondary servers before uninstalling it from the primary server.

To uninstall Backup Exec

- 1 Close Backup Exec.
- 2 From the Windows Control Panel, select **Add or Remove Programs**.
- 3 On the **Add or Remove Programs** dialog box, select **Symantec Backup Exec**, and then click **Remove**.
- 4 When you are prompted to confirm that you want to uninstall Backup Exec from your computer, click **Yes**.
- 5 If you want to remove all of the files that are associated with Backup Exec, click **Yes, please remove Backup Exec and all of its associated files from the system**.
- 6 Click **Next**.
If the uninstall program fails, click **View Installation Log File** for additional information.
- 7 If you are prompted, restart the computer.

See [“Uninstalling Backup Exec options from the local media server”](#) on page 169.

Uninstalling Backup Exec options from the local media server

The install wizard removes Backup Exec options from the local media server. All corresponding files, registry keys, and configurations are removed.

To uninstall Backup Exec options from the local media server

- 1 On the **Tools** menu, click **Install Options and License Keys on this Media Server**.
- 2 On the **License Keys** panel, select the license key of the option that you want to uninstall, and then click **Remove**.
- 3 Click **Next**.
- 4 On the **Configure Options** panel, verify that the option you want to remove is not checked, and then click **Next**.

- 5 If you are prompted to enter credentials for the Backup Exec service account, type the correct credentials, and then click **Next**.
- 6 Read the installation summary, and then click **Install** to start the process.
- 7 When the install wizard has completed, click **Finish**.

See [“Uninstalling Backup Exec”](#) on page 168.

About updating Backup Exec with LiveUpdate

Symantec LiveUpdate, which provides updates, upgrades, and new versions of Backup Exec, is installed automatically with Backup Exec. If you enable the automatic update option, you can configure LiveUpdate to poll the main Symantec Web server on a scheduled interval. The automatic update option only searches for Backup Exec updates. It does not show updates for other Symantec products that use LiveUpdate. Likewise, when LiveUpdate is scheduled to automatically update other Symantec products, it does not search for Backup Exec updates.

Note: By default, LiveUpdate checks for updates every Sunday night at 10pm. If there is an update, LiveUpdate notifies you with an alert.

In addition to scheduling LiveUpdate, you can also run it manually at any time. You can access LiveUpdate from several locations in Backup Exec. However, you cannot access it from the Windows Start menu.

Caution: During the installation and upgrade processes, Backup Exec stops and starts the SQL Server service several times. Other user-created databases that use the SQL Server instance are unavailable during the installation or upgrade process. To avoid such conflicts, you should install Backup Exec into its own SQL instance.

Backup Exec installs the latest version of LiveUpdate. If a previous version of LiveUpdate is detected on the computer, Backup Exec upgrades it.

You can view any hot fixes or service packs that are installed on the media server.

See [“Viewing installed updates”](#) on page 173.

When LiveUpdate installs updates on the Backup Exec media server, it also determines if computers on which the Remote Agent for Windows Systems have the latest updates. If you do not have the latest updates you receive an alert to install the updates. You can push-install or manually copy those updates to Backup Exec Remote Agents.

If LiveUpdate installs any files, the Bkupinst.htm installation log file is updated with information about those files.

You can use the LiveUpdate Administrator Utility with LiveUpdate. The LiveUpdate Administrator Utility allows an administrator to modify LiveUpdate so that network users can download program and virus definition updates from an internal server rather than going to the Symantec LiveUpdate server over the Internet.

Go to ftp://ftp.symantec.com/public/english_us_canada/liveupdate/luadmin.pdf

See “[About scheduling automatic updates using LiveUpdate](#)” on page 171.

See “[Running LiveUpdate manually](#)” on page 173.

See “[About the installation log](#)” on page 166.

About scheduling automatic updates using LiveUpdate

You can schedule LiveUpdate to check for updates as follows:

- Once on a specific date at a specific time
- Every day at a specific time
- Every week on a specific day of the week and at a specific time
- Every month on a specific day of the month and at a specific time

When you schedule automatic updates through Backup Exec, the settings apply only to updates for Backup Exec. Changes that you make to the LiveUpdate schedule for Backup Exec do not affect the schedule for any other software applications that use LiveUpdate.

At the scheduled time, LiveUpdate automatically connects to the appropriate Web site, and then determines if your files need to be updated. Depending on the options that you select, Backup Exec either downloads and installs the files in the proper location or sends an alert to notify you that updates are available.

Backup Exec sends the following LiveUpdate alerts:

Table 2-17 LiveUpdate alerts

Backup Exec sends this alert	When
LiveUpdate Informational Alert	An update is installed successfully.
LiveUpdate Warning Alert	An update is installed successfully. However, you must restart the computer.
LiveUpdate Error Alert	An update fails to install.

See “[Scheduling automatic updates using LiveUpdate](#)” on page 172.

Scheduling automatic updates using LiveUpdate

You can schedule LiveUpdate to check for updates for Backup Exec.

See [“About scheduling automatic updates using LiveUpdate”](#) on page 171.

To schedule automatic updates using LiveUpdate

- 1 On the **Tools** menu, click **Options**.
- 2 On the properties pane, under **Settings**, click **LiveUpdate**.
- 3 Complete the appropriate options.
See [“Application settings for LiveUpdate”](#) on page 172.
- 4 Click **OK**.

Application settings for LiveUpdate

You can schedule LiveUpdate to check for updates for Backup Exec.

See [“Scheduling automatic updates using LiveUpdate”](#) on page 172.

Table 2-18 Application settings for LiveUpdate

Item	Description
Enable scheduled automatic updates	Lets you schedule automatic updates, and then choose the frequency of the updates.
Automatically download and install all available updates	Enables Backup Exec to download and install all updates that are available without prompting you first.
Only notify me of available updates	Enable Backup Exec to alert you when updates are available. Updates are not downloaded or installed. This option is the default. If you select this option, you need to run LiveUpdate manually to download and install the available updates.
Once	Enables Backup Exec to check for new updates only on the date and time that you specify in the On and At fields.
Daily	Enables Backup Exec to check for new updates every day. In the At field, enter the time to check for new updates.
Weekly	Enables Backup Exec to check for new updates once a week. In the Every field, select the day of the week on which to check for updates. In the At field, enter the time to check for new updates.

Table 2-18 Application settings for LiveUpdate (*continued*)

Item	Description
Monthly	Enables Backup Exec to check for new updates once a month. In the Every field, select the day of the month on which to check for updates. In the At field, enter the time to check for new updates.
Interval	Lets you set the date and time that you want Backup Exec to check for new updates.

Running LiveUpdate manually

You can either set a schedule for LiveUpdate or run LiveUpdate manually at any time to check for updates. You can configure LiveUpdate to run in either Interactive mode or Express mode. Interactive mode gives you the flexibility to choose which updates you want to install. Express mode automatically installs all of the Backup Exec updates. For information about how to change the LiveUpdate mode, see the LiveUpdate documentation.

Note: By default, LiveUpdate is configured for Interactive mode. If you change it to Express mode you must cancel the LiveUpdate session and restart it before the change takes place.

To run LiveUpdate manually

- 1 On the **Tools** menu, click **LiveUpdate**.
- 2 Select **Click here to run LiveUpdate**.
- 3 Do one of the following:
 - If LiveUpdate is set for Express mode Click **Start**.
 - If LiveUpdate is set for Interactive mode Click **Next**.

See [“Scheduling automatic updates using LiveUpdate”](#) on page 172.

Viewing installed updates

You can view hot fixes and service packs that are installed on a media server. You must be logged on with administrator privileges.

If a hot fix is installed prior to a service pack, that hot fix no longer displays as installed since the service pack contains the hot fix.

A hot fix that is offered after the service pack is released is displayed with the prior service pack.

To view installed updates

- 1 On the **Help** menu, click **About**.
- 2 Click **Installed Updates**.

Viewing license information

You can view information about the Backup Exec options that are licensed and installed on a media server. You can also view a list of agents and options that are available for a trial, as well as how much time is left in each individual trial period.

See [“Adding licenses”](#) on page 175.

To view license information

- ◆ Do one of the following:

To view license information from the Help menu Do the following in the order listed:

- On the **Help** menu, click **About**.
- Click **License Information**.

To view license information from the media server properties Do the following in the order listed:

- On the navigation bar, click **Devices**.
- Select the media server from the tree view.
- Under **General Tasks** in the task pane, click **Properties**, and then click **License Information**.

See [“License information options”](#) on page 174.

License information options

You can view information about the Backup Exec options that are licensed and installed on a media server.

See [“Viewing license information”](#) on page 174.

Table 2-19 License Information options

Item	Description
Option	Lists the names of the available Backup Exec options.
Licensed	Displays Yes if the option is licensed on the media server. Displays No if the option is not licensed.
Installed	Displays Yes if the option is installed on the media server. Displays No if the option is not installed. If the option is installed, it may still require some additional configuration.
Trial	Lists the following statuses: <ul style="list-style-type: none"> ■ Available ■ Expired ■ N/A If the option is in a trial period, the remaining number of days in the trial period appears.
Maintenance Contract Information	Launches the Maintenance Contract Information dialog box. See " Backup Exec maintenance contract information " on page 179.

Adding licenses

You can add licenses to activate additional agents or options at any time. If the trial period runs out on an agent or option, you need to enter a license key to continue to use it.

For information about how to obtain a license key, contact your reseller, or go to the following URL:

<https://licensing.symantec.com>

To add licenses

- 1 On the **Tools** menu, click **Install Options and License Keys on this Media Server**.
- 2 Select one of the following methods to enter license keys:

To manually enter license keys

Do the following in the order listed:

- Type a license key into the license key field.
- Click **Add**.
- Repeat for each license key for each option or agent that you want to install.

To import license keys from a file

Do the following in the order listed:

- Click **Import from file**.
- Select the besernum.xml file.

3 Click **Next**.

4 Verify that the additional options are selected for installation, and then click **Next**.

5 If you are prompted, enter information or choose settings for the additional options that you want to install. Click **Next** after each selection.

6 Read the Backup Exec installation review, and then click **Install**.

The Backup Exec services stop while the agents or options are installed. If any active jobs are in progress, you are prompted to stop them, or to wait for the jobs to finish.

When the installation is complete, the services are restarted.

7 Click **Finish**.

See “[Viewing license information](#)” on page 174.

Finding installed licenses in your environment

The License Assessment Tool lets you run a license key scan on the computers on which the following are installed:

- Backup Exec 2010
- Backup Exec System Recovery

Both of these products are Backup Exec installations.

You can also use the License Assessment Tool to view maintenance contract expiration dates so that you can renew contracts before they expire.

The License Assessment Tool reports maintenance contract information when the following conditions are met:

- The media server in the license key scan must run Backup Exec 2010 R2.

- The media server must have maintenance contract information available. To obtain this information, on the Backup Exec Administration Console, click **Tools > Licensing and Maintenance Contract Information > Renewal Assistant**.
- The media server must have a network connection.

See “[About Backup Exec maintenance contract information](#)” on page 177.

On each Backup Exec installation for which you run a license key scan, the License Assessment Tool reviews the resources that are backed up. Resources are files such as Windows shares, or application databases such as Microsoft SQL Server. A report compares the number of resources that are backed up with the number of license keys that are installed.

Note: Scans for time periods and date ranges do not apply to the Backup Exec Archiving Option. Only the resources that are backed up by Remote Agents are scanned if you select a time period or date range.

The License Assessment Tool report provides the following information:

- The number of additional licenses that are recommended for a Backup Exec installation.
- The versions of Backup Exec that are installed so that you can consider purchasing upgrades.

Running the License Assessment Tool does not ensure license compliance. For more information about licensing, contact your reseller, or go to the following URL:

<https://licensing.symantec.com>

To find installed licenses in your environment

- 1 On the **Tools** menu, click **Licensing and Maintenance Contract Information > Backup Exec License Assessment Tool**.
- 2 Follow the on-screen prompts.

About Backup Exec maintenance contract information

After you purchase maintenance contracts for your Backup Exec products, Symantec automatically updates the Symantec Licensing Portal Web site with your maintenance contract information. Maintenance contract information includes the contract serial number and the contract expiration date.

You can use the **Renewal Assistant** to access the maintenance contract expiration information from within Backup Exec. The **Renewal Assistant** then uses the contract expiration information to automatically set Backup Exec alerts that remind you to renew the maintenance contracts before they expire. Reminder alerts are set at 30-day, 60-day, and 90-day intervals, based on the expiration date of the maintenance contract.

To retrieve the contract expiration dates, you enter your maintenance contract serial numbers in the **Renewal Assistant**. The **Renewal Assistant** connects to the Symantec Licensing Portal Web site, at which point you may be prompted to enter customer and technical contact information. The **Renewal Assistant** then retrieves the maintenance contract information for each contract that you have purchased.

Note: The **Renewal Assistant** requires an active Internet connection on the media server where you run the **Renewal Assistant**. If the Internet is unavailable, you can use the **Renewal Assistant** to set a one-time manual alert for the contracts that appear on your purchase certificate.

All Backup Exec products and maintenance contracts are assigned Symantec serial numbers. The serial numbers appear on the printed certificate that you receive with your order. To determine the correct serial numbers to enter in the **Renewal Assistant**, look for the Maintenance/Subscription columns on your certificate. Each maintenance contract specifies a start date and an end date; these dates appear in the Maintenance/Subscription columns.

Maintenance contract serial numbers consist of a letter and a series of numbers. They are not the same as Backup Exec license keys, nor do they use the same format as license keys.

Table 2-20 Backup Exec serial number formats and license key formats

Serial number format	License key format
M0000000000	00-A00-0000-000000

See [“Running the Backup Exec Renewal Assistant”](#) on page 178.

See [“Viewing Backup Exec maintenance contract information”](#) on page 179.

Running the Backup Exec Renewal Assistant

You can run the **Renewal Assistant** at any time to enter maintenance contract serial numbers.

See [“About Backup Exec maintenance contract information”](#) on page 177.

To run the Backup Exec Renewal Assistant

- 1 On the **Tools** menu, click **Licensing and Maintenance Contract Information > Renewal Assistant**.
- 2 Follow the on-screen directions.

Viewing Backup Exec maintenance contract information

You can view maintenance contract information after you enter contract serial numbers in the **Renewal Assistant**. Backup Exec stores maintenance contract details locally on the media server after it retrieves the information from the Symantec Licensing Portal Web site.

See [“About Backup Exec maintenance contract information”](#) on page 177.

See [“Running the Backup Exec Renewal Assistant”](#) on page 178.

To view Backup Exec maintenance contract information

- 1 On the **Tools** menu, click **Licensing and Maintenance Contract Information > Maintenance Contract Information**.
See [“Backup Exec maintenance contract information”](#) on page 179.
- 2 Click **OK**.

Backup Exec maintenance contract information

You can view the maintenance contract information that Backup Exec has retrieved from the Symantec Licensing Portal Web site.

Note: A maintenance contract serial number must be entered in the **Renewal Assistant** before the contract information appears.

See [“About Backup Exec maintenance contract information”](#) on page 177.

See [“Running the Backup Exec Renewal Assistant”](#) on page 178.

See [“Viewing Backup Exec maintenance contract information”](#) on page 179.

Table 2-21 Maintenance contract information details

Item	Description
Feature	<p>Shows the name of the feature that the maintenance contract covers.</p> <p>A green check icon indicates that the maintenance contract for the feature is valid.</p> <p>A red X icon indicates that the maintenance contract for the feature has expired.</p>
Serial Number	Shows the serial number of the maintenance contract.
Expiration Date	Shows the expiration date of the maintenance contract.
Renewal Assistant	Starts the Renewal Assistant .

About upgrading from previous versions of Backup Exec

You can use the Backup Exec installation media to upgrade from Backup Exec version 11d and later to the current version. No separate upgrade utility is necessary. The current version of Backup Exec replaces any previous versions. Separate installations of different versions cannot exist on the same computer. Most settings and all catalogs and all data directories from previous versions of Backup Exec are kept, unless you choose to remove them. Backup Exec 2010 R2 can read and restore data from any previous version of Backup Exec or Backup Exec for NetWare, except where Symantec has made end-of-life decisions.

When you upgrade, Backup Exec automatically converts your existing definitions, configurations, and jobs to the current version. If any aspect of the data migration fails, then Backup Exec displays an error message stating that the upgrade did not complete successfully. Backup Exec provides you with the choice to either return to the previous version of Backup Exec or to continue with the upgrade. If you choose to continue with the upgrade, then devices, media, and jobs are not migrated from the previous version of Backup Exec.

Backup Exec 2010 R2 provides backward compatibility as follows:

- Backup Exec 2010 R2 can communicate with Backup Exec 12 Remote Agent for Windows Systems and later.
- Backup Exec 2010 R2 can communicate with Backup Exec for NetWare 9.2 Remote Agent.

- Backup Exec 2010 R2 Central Admin Server Option server can communicate with Backup Exec 12.5 (with the appropriate patch) or Backup Exec 2010 (with the appropriate patch) for the purpose of rolling upgrades.
- Backup Exec 2010 R2 Command Line Applet is compatible with Backup Exec 12 and later.

A Backup Exec 2010 R2 Remote Administration Console can manage media servers on which Backup Exec version 11d and later is installed. However, if the media server uses a previous version of Backup Exec, you cannot use any of the new features of the current version. If you want to use the features of the current version, you must use the current version on both the Remote Administration Console and the media server. A Remote Administration Console that uses a previous version of Backup Exec cannot be used with a media server on which the current version is installed. For example, a Backup Exec 2010 Remote Administration Console cannot manage a Backup Exec 2010 R2 media server. However, a Backup Exec 2010 R2 Remote Administration Console is backward-compatible with a Backup Exec 2010 Remote Administration Console.

Before you upgrade Backup Exec, do the following:

- Delete the job histories and the catalogs that you no longer need to shorten the upgrade window.
- Run a database maintenance job.
- Upgrade any existing instances of SQL Server 2000 to either SQL Server 2005 with SP3 or to SQL Server 2008.

You cannot change the configuration of your media servers during an installation. For example, you cannot change a central administration server to a managed media server. If you want to change the configuration of your media servers, do it either before or after you upgrade to the current version. You cannot change the database location during the upgrade process. If you want to change the database location after the upgrade, use BEUtility.

To upgrade the options that are installed on remote computers, you must reinstall them. Options that are push-installed are not upgraded until you reinstall them. The Remote Agent for Windows Systems and the Advanced Open File Option are push-installed.

Note: If you upgrade from a previous version of Backup Exec that uses a non-English version of Windows, you must download the SQL Express SP3 setup file for that language from the Microsoft Web site.

See [“Installing a typical installation of Backup Exec”](#) on page 116.

See [“Installing a custom installation of Backup Exec”](#) on page 119.

Post-installation tasks

For best results before starting Backup Exec, do the following:

- Make sure that your storage devices are connected and configured properly. See [“About storage devices”](#) on page 441.
- Decide if your backup will be to a tape device or a disk device. You can configure both devices when you prepare your Backup Exec environment.

Note the following:

- If you’re backing up to a tape device, verify that the device is supported. You can install drivers for the devices when you configure your Backup Exec environment.
- If you’re backing up to a disk device using the Backup-to-Disk feature, decide where you can create a backup folder. You should create it on a disk that won’t be included in the backup jobs and that has enough free space to contain the backup job. See [“About backup-to-disk folders”](#) on page 496.
- Understand how Backup Exec provides overwrite protection for your media. See [“About media overwrite protection”](#) on page 223.
- Understand the default media set and its four-week overwrite protection period. See [“About media in Backup Exec”](#) on page 219.
- Learn about creating new media sets with weekly, monthly, or quarterly retention periods. See [“About the default media set”](#) on page 226.
- Decide which resource credential you want your Backup Exec logon account to use when browsing and making backup selections. You can use an existing Backup Exec logon account, or create a new one. See [“Creating a Backup Exec logon account”](#) on page 187.
- Decide the format that you want to display all reports, either HTML or Adobe Portable Document Format (PDF). The default setting is HTML. See [“Editing application settings for reports”](#) on page 719.

Configuring Backup Exec settings and options

This chapter includes the following topics:

- [About configuring Backup Exec](#)
- [About configuring logon accounts](#)
- [About Backup Exec defaults](#)
- [About job priority](#)
- [Changing the default device and media set for jobs](#)
- [Changing application settings for Backup Exec preferences](#)
- [Copying configuration settings to another media server](#)
- [Copying logon account information](#)
- [About audit logs](#)
- [About database maintenance](#)
- [Viewing the location of Backup Exec databases](#)
- [About the Symantec RSS Reader](#)
- [Hiding columns](#)
- [Showing a hidden column](#)
- [Rearranging columns](#)
- [Sorting column information](#)

- [Viewing properties](#)

About configuring Backup Exec

During installation and prior to first use, you configure several Backup Exec features, such as a default logon account, database information, and audit log settings. Configuring Backup Exec lets you standardize Backup Exec before jobs are created and run.

You can perform the following initial operations:

- Configure logon accounts.
See [“About configuring logon accounts”](#) on page 184.
- Set the default Backup Exec logon account.
See [“About the default Backup Exec logon account”](#) on page 185.
- Create new Backup Exec system logon accounts.
See [“About the Backup Exec System Logon Account”](#) on page 189.
- Copy configuration settings and logon information to another media server.
See [“Copying configuration settings to another media server”](#) on page 199.
- Configure audit logs.
See [“About audit logs”](#) on page 205.
- Configure database maintenance.
See [“Configuring database maintenance”](#) on page 209.
- Configure and organize columns in Backup Exec.
See [“Hiding columns”](#) on page 215.

About configuring logon accounts

A Backup Exec logon account stores the credentials of a user account that you use to access a resource, such as a Windows computer. Backup Exec logon accounts enable Backup Exec to manage user names and passwords and can be used to browse resources or process jobs. Using Backup Exec logon accounts enables you to apply credential changes to the jobs that use them.

Backup Exec logon accounts are used to browse local and remote resources. Whenever the Backup Exec logon credentials are passed between the media server and the remote resource, the credentials are encrypted.

Backup Exec logon accounts can also be associated with selection list entries at the device level such as shares, databases, etc. If you need to edit the credentials,

you can edit the Backup Exec logon account and the changes will be applied to the selected resources that use the Backup Exec logon account.

Backup Exec logon accounts are not user accounts. When you create a Backup Exec logon account, an entry for the account is entered into the Backup Exec database; no operating system accounts are created. If your user account credentials change, you must update the Backup Exec logon account with the new information. Backup Exec does not maintain a connection with the user account.

You can view, create, delete, edit, and replace Backup Exec logon accounts.

The following types of logon accounts are included in Backup Exec:

Table 3-1 Types of logon accounts

Type of logon account	Description
Default Backup Exec logon account	Used to browse local and remote resources, make backup job selections, and restore data. See “About the default Backup Exec logon account” on page 185.
Backup Exec system logon account	Used to access most or all of your resources. It contains the Backup Exec Services credentials. See “About the Backup Exec System Logon Account” on page 189.
Backup Exec logon account	Used to manage Backup Exec user names and passwords, browse local and remote resources, process jobs, and apply credential changes to the jobs that use them. See “Creating a Backup Exec logon account” on page 187.

See [“Creating a Backup Exec logon account”](#) on page 187.

About the default Backup Exec logon account

The default Backup Exec logon account enables you to browse, make selections, or restore data. The first time you start Backup Exec, you must specify a default Backup Exec logon account using the Logon Account Wizard. You can select an existing Backup Exec logon account or create a new one.

You can create multiple Backup Exec logon accounts; however, each Backup Exec user can only have one default Backup Exec logon account.

Your default Backup Exec logon account enables you to perform the following:

- Browse resources. Your default Backup Exec logon account enables you to browse local and remote resources when you create backup jobs. To browse

resources, each user must have a default Backup Exec logon account that is associated with their user account. The Backup Exec logon account does not have to be the same user name as the user that is used to log on to Backup Exec.

For example, you are logged on to a media server named MEDIASERVER as the local Windows administrator. When you start Backup Exec, you are prompted to create a default Backup Exec logon account for the local administrator because one does not exist. You can create a Backup Exec logon account for the local administrator that has the credentials for a domain administrator. The Backup Exec logon account will have the following properties:

User name: DOMAIN\Administrator

Description: MEDIASERVER\Administrator Default Account

Owner: MEDIASERVER\Administrator

When you change your default Backup Exec logon account, you can use your new default Backup Exec logon account to browse resources immediately; you do not have to restart your system in order for the changes take effect.

- Make backup selections. You can select a different Backup Exec logon account when you make selections for backup. If your default logon account does not have rights, the Logon Account Selection dialog box appears and enables you to create or select a different Backup Exec logon account. You can also change the Backup Exec logon account when making backup selections using the Connect As command in the context menu.

See [“How to use Backup Exec logon accounts for SQL resources”](#) on page 1256.

See [“Requirements for accessing Exchange mailboxes”](#) on page 1101.

- Restore. You can assign Backup Exec logon accounts to resources when you create restore jobs. The default Backup Exec logon account is used unless you choose a different Backup Exec logon account when you create the restore job, in Resource Credentials in the Restore Job Properties.

See [“Changing your default Backup Exec logon account”](#) on page 193.

See [“About Backup Exec restricted logon accounts”](#) on page 186.

About Backup Exec restricted logon accounts

Backup Exec logon accounts can be common or restricted. When you create a Backup Exec logon account, you can designate it as a restricted account. To use a restricted logon account, you must be the owner of the logon account or you must know the password for the logon account. The person who created the logon account is the owner. If you authorize only a few people to back up or restore data, you can make the logon account a restricted logon account.

The main reasons to restrict a logon account are as follows:

- To help you limit access to the resources available for backup.
- To help you limit the computers to which you can restore.

When you use a restricted logon account to select the resources for a job, the logon account information is saved with the selection list. Anyone who tries to edit the job must provide the password to the restricted logon account. Backup Exec loads the selections for that job only when the password for the restricted logon account is provided.

See [“Creating a Backup Exec logon account”](#) on page 187.

See [“Editing a Backup Exec logon account”](#) on page 190.

Creating a Backup Exec logon account

You can create Backup Exec logon accounts using the Logon Account Wizard, which guides you through the creation of a Backup Exec logon account, or by using the Logon Account Management dialog. You can enter Backup Exec logon account property information when you create the Backup Exec logon account; however, Backup Exec assigns the Backup Exec logon account owner to the user name you used to log on to Backup Exec. The owner of the Backup Exec logon account cannot be modified.

See [“Editing a Backup Exec logon account”](#) on page 190.

See [“Replacing a Backup Exec logon account”](#) on page 192.

See [“Changing your default Backup Exec logon account”](#) on page 193.

To create a Backup Exec logon account

- ◆ Do one of the following:

To create a new logon account by using the Logon Account wizard

On the **Tools** menu, click **Wizards > Logon Account Wizard**.

The wizard guides you through the setup process.

To create a new logon account manually Do the following in the order listed:

- 1 On the **Network** menu, click **Logon Accounts**.
- 2 Click **New**.
- 3 Enter the appropriate options.

See [“Add Logon Credentials options”](#) on page 188.

Add Logon Credentials options

You can enter Backup Exec logon account property information when you create the Backup Exec logon account.

See [“Creating a Backup Exec logon account”](#) on page 187.

Table 3-2 New Logon Account options

Item	Description
User name	Indicates the fully qualified user name for the Backup Exec logon account. For example, DOMAIN\Administrator. The user name is provided when you attempt to connect to a resource. The user name is not case sensitive for the resources that are accessed.
Password	Indicates the password for the account. The password you enter is encrypted for security. You can leave this field blank if this Backup Exec logon account does not need a password.
Confirm password	Verifies the password. The password must match the password you typed in the Password field.
Account Name	Indicates the unique name for the Backup Exec logon account. The user name is automatically added if you do not enter information into the field.
Notes	Indicates how the Backup Exec logon account will be used.

Table 3-2 New Logon Account options (*continued*)

Item	Description
This is a restricted logon account	Enables the Backup Exec logon account to be used only by the owner of the logon account and those who know the password. If this is not selected, the Backup Exec logon account will be a common account. Common accounts are shared accounts that can be accessed by all users. See “About Backup Exec restricted logon accounts” on page 186.
This is my default logon account	Makes this account your default Backup Exec logon account, which is used to browse, make selections, or restore data on your local and remote resources.

About the Backup Exec System Logon Account

The Backup Exec System Logon Account (SLA) is created when you install Backup Exec. When the SLA is created, the user name and password match the credentials provided during install for the Backup Exec Services credentials. The owner of the SLA is the user that installed Backup Exec and is a common account by default. Common accounts are shared accounts that can be accessed by all users.

The Backup Exec System Logon Account may have access to most or all of your resources since it contains the Backup Exec Services credentials. If you want to make Backup Exec more secure, you can change the SLA to be a restricted account. You can also delete it after making another logon account the default. However, if you delete the SLA, the jobs in which it is used may fail. If the SLA is deleted, you can re-create it using the Logon Account Management dialog box.

The SLA is used for the following tasks and jobs:

- Jobs migrated from a previous version of Backup Exec
- Duplicate backup data jobs
- Command Line Applet (bemcmd.exe)
- Backup Exec Agent for SAP Applications

See [“Creating a new Backup Exec System Logon Account”](#) on page 193.

See [“Creating a Backup Exec logon account”](#) on page 187.

See [“Editing a Backup Exec logon account”](#) on page 190.

See [“Replacing a Backup Exec logon account”](#) on page 192.

See [“Deleting a Backup Exec logon account”](#) on page 192.

See [“Changing your default Backup Exec logon account”](#) on page 193.

See [“Copying configuration settings to another media server”](#) on page 199.

Editing a Backup Exec logon account

When you edit a Backup Exec logon account, the changes are automatically applied to all the resources that use the Backup Exec logon account. Changes made to a Backup Exec logon account are applied immediately. You do not have to restart your system for the changes to take effect.

You can edit the following properties for a Backup Exec logon account:

- Type (restricted or common)
- Account name
- Password
- User name
- Notes

See [“Changing your default Backup Exec logon account”](#) on page 193.

To edit a Backup Exec logon account

- 1 On the **Network** menu, click **Logon Accounts**.
- 2 Select the Backup Exec logon account you want to change, and then click **Edit**.

If you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can edit the account.

- 3 Modify the Backup Exec logon account properties as needed.

See [“Edit Logon Credentials options”](#) on page 190.

- 4 On the **Edit Logon Credentials** dialog box, click **OK**.

Edit Logon Credentials options

You can change the properties of an existing logon account.

See [“Editing a Backup Exec logon account”](#) on page 190.

Table 3-3 Edit Logon Credentials options

Item	Description
User name	Indicates the fully qualified user name for the Backup Exec logon account. For example, DOMAIN\Administrator. The user name is provided when you attempt to connect to a resource. The user name you enter is not case sensitive for the resources that are accessed.
Change Password	Enables you to change the password for the account. The password you enter is encrypted for security.
Account Name	Indicates the unique name for the Backup Exec logon account. The user name is automatically added if you do not enter information into the field.
Notes	Indicates how the Backup Exec logon account will be used.
This is a restricted logon account	Enables the Backup Exec logon account to be used only by the owner of the logon account and those who know the password. If this is not selected, the Backup Exec logon account will be a common account. Common accounts are shared accounts that can be accessed by all users. See “About Backup Exec restricted logon accounts” on page 186.
This is my default logon account	Makes this account your default Backup Exec logon account used to browse, make selections, or restore data on your local and remote resources.

Changing a Backup Exec logon account password

You can change a Backup Exec logon account password using the following steps. Changes made to a Backup Exec logon account password are applied immediately.

See [“About configuring logon accounts”](#) on page 184.

To change a Backup Exec logon account password

- 1 On the **Network** menu, click **Logon Accounts**.
- 2 Select the Backup Exec logon account to change, and then click **Edit**.
If you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can edit the account.
- 3 Click **Change password**.
- 4 In the **Password** field, type a new password.
- 5 In the **Confirm** field, re-type the password, and then click **OK**.
- 6 On the **Edit Logon Credentials** dialog box, click **OK**.
- 7 On the **Logon Account Management** dialog box, click **OK**.

Replacing a Backup Exec logon account

You can replace a Backup Exec logon account within all existing jobs and selections lists. The resources and selections lists in existing jobs that use the Backup Exec logon account will be updated to use the new Backup Exec logon account. If the new Backup Exec logon account is restricted, you must provide the password.

See [“About configuring logon accounts”](#) on page 184.

To replace a Backup Exec logon account

- 1 On the **Network** menu, click **Logon Accounts**.
- 2 Select the Backup Exec logon account you want to replace, and then click **Replace**.
- 3 Select the Backup Exec logon account with which you want to replace the selected Backup Exec logon account.

If the Backup Exec logon account is restricted and you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can edit the account.

- 4 Click **OK**.

Deleting a Backup Exec logon account

You cannot delete a Backup Exec logon account in the following situations:

- It is being referenced by a job.
- It is owned by a user who is logged on to the media server.

- It is set as the default Backup Exec logon account of a user who is logged on to the media server.

You can delete a Backup Exec logon account when the owner is logged off and all users who have it set as their default logon account are logged off.

See [“About configuring logon accounts”](#) on page 184.

To delete a Backup Exec logon account

- 1 On the **Network** menu, click **Logon Accounts**.
- 2 Select the Backup Exec logon account you want to delete, and then click **Delete**.
- 3 Click **Yes** to confirm the deletion.

Changing your default Backup Exec logon account

You can change your default Backup Exec logon account that enables you to browse, make selections, or restore data.

See [“About the default Backup Exec logon account”](#) on page 185.

To change your default Backup Exec logon account

- 1 On the **Network** menu, click **Logon Accounts**.
- 2 Select the Backup Exec logon account you want to use as your default Backup Exec logon account, and then do one of the following:
 - Click **Set as Default**.
 - Click **Edit**, select **This is my default logon account**, and then click **OK**.

Creating a new Backup Exec System Logon Account

The Backup Exec System Logon Account enables you to perform several operations. It is also used with Backup Exec Agent for SAP Applications and Command Line Applet. If you delete the Backup Exec System Logon Account, you should create a new one that enables you to perform the specified operations and use the agent and applet.

See [“About the Backup Exec System Logon Account”](#) on page 189.

To create a new Backup Exec System Logon Account

- 1** On the **Network** menu, click **Logon Accounts**.
- 2** Click **System Account**.
- 3** Select the appropriate options, and then click **OK** to create the system logon account.

See “[Edit Logon Credentials options](#)” on page 190.

About Backup Exec defaults

When you start Backup Exec for the first time, defaults are already configured. You can adjust the defaults to meet the needs of your environment. Default settings are available for various types of jobs, such as backup, restore, and test run. You also can set defaults for catalogs, media management, bar code rules, and database maintenance.

The defaults that will probably affect you the most are the backup job defaults. You can change many of these defaults after devising a media rotation strategy and creating additional media sets and drive pools. In the short term though, you can run Backup Exec and backup and restore jobs safely by using only the defaults set during installation.

See the following sections for more information about default options:

Table 3-4 Backup Exec default options

Item	Description
Jobs	<p>See “Setting default backup options” on page 391.</p> <p>See “Setting defaults for restore jobs” on page 638.</p> <p>See “Setting test run default options” on page 387.</p> <p>See “Changing application settings for the catalogs” on page 601.</p> <p>See “Creating separate selection lists for each computer or resource” on page 309.</p> <p>See “Setting priority and availability windows for selection lists” on page 307.</p>
Customizing Backup Exec	<p>See “Setting default pre/post commands” on page 400.</p> <p>See “Setting default backup network and security options” on page 404.</p>
Configuring Backup Exec	<p>See “Configuring database maintenance” on page 209.</p>

Table 3-4 Backup Exec default options (*continued*)

Item	Description
Reports	See “Editing application settings for reports” on page 719.
Administrating Backup Exec	See “Configuring default schedule options” on page 366. See “Using checkpoint restart on Microsoft Cluster Server failover” on page 822. See “Setting thresholds to recover jobs” on page 596. See “About scheduling automatic updates using LiveUpdate” on page 171.
Device and Media	See “Changing application settings for Backup Exec preferences” on page 196. See “Media locations and vaults” on page 250. See “Bar code rules in mixed media libraries” on page 245.
Options	See “Setting offhost backup options for a backup job” on page 925. See “Setting default options for the Advanced Open File Option” on page 943. See “Setting default backup and restore options for SQL” on page 1265. See “How to prepare for disaster recovery of Exchange Server” on page 1165. See “Configuring default Lotus Domino options” on page 1069. See “Setting default options for the Remote Agent for NetWare Systems ” on page 1941. See “Setting default options for SharePoint Server 2003 and later” on page 1196.

About job priority

You can set the priority of access to the devices for Backup Exec jobs.

You can choose from the following levels of priority:

- **Highest**
- **High**
- **Medium**
- **Low**

■ **Lowest**

This option is most useful if there are limited devices in your environment, but you want certain jobs to have priority access to the devices. A ready job that has a higher priority runs before a ready job that has a lower priority. A ready job that has a higher priority also runs before a ready job that has an earlier scheduled start time.

If multiple jobs are ready to run but must wait for a device to become available, then Backup Exec determines which jobs to run first. To make this determination, Backup Exec reviews the job priority and the scheduled start time of the job.

Changing the default device and media set for jobs

You can set the default device and media set to use for each job that you create. You can change the defaults for each job individually.

To change the default device and media set for jobs

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Options - Set Application Defaults** dialog box, under **Job Defaults**, click **Device and Media**.
- 3 In the **Device** field, select the device that you want to use as the default device for jobs.
- 4 If you have installed the Deduplication Option, select the default deduplication method to use for deduplication backup jobs.
- 5 In the **Media set** field, select the media set that you want to be used as the default media set for jobs.

See [“About the Deduplication Option”](#) on page 1568.

Changing application settings for Backup Exec preferences

You can change the settings for how you prefer Backup Exec to display various screens, indicators, and alerts.

To change application settings for Backup Exec preferences

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Settings**, click **Preferences**.
- 3 Select the appropriate options.

See “[Application settings for Backup Exec preferences](#)” on page 197.

Application settings for Backup Exec preferences

You can change the settings for how you prefer Backup Exec to display various screens, indicators, and alerts.

See “[Changing application settings for Backup Exec preferences](#)” on page 196.

Table 3-5 Application settings for Backup Exec preferences

Item	Description
Show splash screen at startup	Displays the splash screen when you start Backup Exec. If this option is cleared, the Backup Exec Administration Console is the first thing to display on startup.
Include robotic libraries in inventory job when Backup Exec services start up	Enables Backup Exec to inventory all of the slots in a robotic library when Backup Exec’s services are starting. Depending on the number of storage devices that are attached to your system, this process may take a few minutes.
Display the job summary before creating a job	Enables Backup Exec to display a summary of the job options you selected before submitting the job to the job queue.
Create jobs after a new policy is created	Enables Backup Exec to automatically display the Create or Delete Policy Jobs dialog box after you create a policy. You can use the Create or Delete Policy Jobs dialog box to create jobs by associating selection lists with policies.
Display progress indicators for backup jobs. This requires additional time to pre-scan devices.	<p>Displays the percentage complete number while a backup job processes. These indicators appear in the Job Activity tab, and they allow you to monitor the progress of the active job. Backups might take a little longer to complete when this option is selected because the target resources must be scanned to determine the amount of data to be backed up.</p> <p>Due to the time required to scan the target resources, selecting this option when backing up remote resources is not recommended.</p>

Table 3-5 Application settings for Backup Exec preferences (*continued*)

Item	Description
Enable percentage bars when available	<p>Displays a shaded percentage complete bar in the Percent Complete column for active jobs. The percentage complete bar displays in addition to the percentage complete number.</p> <p>If the color depth on your computer is set to 256 or less, this option may appear as unavailable.</p>
Enable ScreenTips	<p>Enables or disables ScreenTips, which provide brief explanations of selected items on the Administration Console. When ScreenTips are enabled, you can view them by holding the mouse pointer on an item. Only selected items have ScreenTips.</p>
Participate in the Backup Exec product improvement program	<p>Enables you to participate in the Symantec Backup Exec product improvement program.</p> <p>General Backup Exec usage and statistical information is periodically collected and sent anonymously to Symantec. Symantec uses the information to help improve the Backup Exec customer experience.</p> <p>Note: Although usage and statistical information is collected, Symantec never collects specific user information.</p>
Automatically display new alerts	<p>Enables alerts to automatically appear on the desktop when they are sent. Alerts that require a response always appear on the Backup Exec console.</p> <p>If you do not choose this option, you are required to view and respond to alerts through the Alerts view.</p>
Play sound after receiving alert	<p>Enables Backup Exec to send an audible tone when an alert is generated. Information about alerts can be found in the Alerts view.</p>
Shade alternate rows in Backup Exec views	<p>Enables or disables highlighting on every other row in various lists, such as the current jobs and job history lists on the Job Monitor. Highlighting facilitates viewing of long lists. This option is selected by default.</p> <p>If the color depth on your computer is set to 256 or less, this option may appear as unavailable.</p>
Set contrast of shading	<p>Sets the darkness of the shading in the rows, if you selected the Shade alternate rows in Backup Exec views option.</p>

Table 3-5 Application settings for Backup Exec preferences (*continued*)

Item	Description
Send an alert on this date as a reminder to renew your maintenance contracts	<p>Lets you select the date on which you want to receive a reminder to renew your maintenance contracts for Backup Exec features that are installed.</p> <p>Symantec recommends that you keep your maintenance contracts current for the Backup Exec products that are installed. Maintenance contracts ensure that you have access to the latest upgrades and to technical support. If your maintenance contracts expire, your entitlement to technical support also expires, and upgrades are no longer automatically sent.</p>

Copying configuration settings to another media server

You can copy configuration settings and logon information from one media server to another. This copy ability allows you to quickly set up a group of media servers with the same configuration or logon settings.

See [“Copying logon account information”](#) on page 204.

To copy configuration settings and logon information to other media servers, the Copy Server Configurations feature must be installed.

See [“Push-installing the Remote Agent and Advanced Open File Option to remote computers”](#) on page 134.

To copy configuration settings to another media server

- 1 On the Tools menu, click **Copy Settings to Media Servers**.
- 2 Select the appropriate settings to copy.
See [“Copy Settings options”](#) on page 203.
- 3 Do one of the following:

If the media server to which you want to copy the setting appears in the **Destination media servers** list Select the name of the media server.

If the media server to which you want to copy the settings does not appear in the **Destination media servers** list

Do one of the following:

- Click **Add** to add a media server to the list. After you add the media server, you can select it as a destination. See [“Add Server options in a CASO environment”](#) on page 202. See [“Add Server options in a non-CASO environment”](#) on page 201.
- Click **Import List** to add multiple media servers from a list. After you add the list, you can select any of the media servers on the list as a destination.

- 4 Click **OK**.

Adding multiple destination media servers by importing a list

You can copy some settings from one media server to another media server. If the media server to which you want to copy the settings does not appear in the Destination media servers list on the Copy Settings dialog box, you can add it by importing a list. After you add a media server to the **Destination media servers** list, you can select it as a destination.

See [“Copying configuration settings to another media server”](#) on page 199.

To add multiple destination media servers by importing a list

- 1 On the **Tools** menu, click **Copy Settings to Media Servers**.
- 2 On the **Copy Settings** dialog box, click **Import List**.
- 3 Browse to select the list, and then click **Open**.
- 4 Click **OK**.

Adding a destination media server in a non-CASO environment

You can copy some settings from one media server to another media server. If the media server to which you want to copy the settings does not appear in the Destination media servers list on the Copy Settings dialog box, you can add it. After you add a media server to the **Destination media servers** list, you can select it as a destination.

See [“Copying configuration settings to another media server”](#) on page 199.

To add a destination media server in a non-CASO environment

- 1 On the **Tools** menu, click **Copy Settings to Media Servers**.
- 2 Click **Add**.
- 3 Enter the name of the media server.
- 4 If necessary, click **Change Logon Account** and select or enter the correct logon account information

See [“About configuring logon accounts”](#) on page 184.

Changing a logon account does not permanently change the logon account for each selected media server.

- 5 Click **OK**.

Add Server options in a non-CASO environment

You can use the **Add Server options** dialog box to add media servers to which you want to copy settings.

See [“Copying configuration settings to another media server”](#) on page 199.

See [“Adding a destination media server in a CASO environment”](#) on page 201.

Table 3-6 Add Server options in a non-CASO environment

Item	Description
Media server name	Indicates the name of the media server that you want to add to the Destination media servers list on the Copy Setting dialog box.
Logon account to use to connect to the media server	Shows the name of the logon account that is used to access the media servers you selected.
Change Logon Account	Lets you change the logon account that Backup Exec uses to access the media servers you selected.

Adding a destination media server in a CASO environment

You can copy some settings from one media server to another media server. If the media server to which you want to copy the settings does not appear in the Destination media servers list on the Copy Settings dialog box, you can add it. After you add a media server to the **Destination media servers** list, you can select it as a destination.

See [“Copying configuration settings to another media server”](#) on page 199.

To add a destination media server in a CASO environment

- 1 On the **Tools** menu, click **Copy Settings to Media Servers**.
- 2 Click **Add**.
- 3 Select the appropriate options.
See [“Add Server options in a CASO environment”](#) on page 202.
- 4 If necessary, click **Change Logon Account** and select or enter the correct logon account information to be used to complete the copy operation.
See [“About configuring logon accounts”](#) on page 184.
Changing a logon account's credentials for a copy operation does not permanently change the logon account.
- 5 Click **OK**.

Add Server options in a CASO environment

You can use the **Add Server options** dialog box to add media servers to which you want to copy settings.

See [“Copying configuration settings to another media server”](#) on page 199.

See [“Adding a destination media server in a CASO environment”](#) on page 201.

Table 3-7 Add Server options in a CASO environment

Item	Description
Add an individual media server	Lets you select a single media server to add to the Destination media servers list on the Copy Settings dialog box. After you add the name of the media server, you can then copy settings to it.
Media server name	Indicates the name of the media server that you want to add to the Destination media servers list on the Copy Setting dialog box.
Add all managed media servers	Lets you add all of the managed media servers in your environment to the Destination media servers list on the Copy Settings dialog box. This option lets you copy settings to any managed media server.

Table 3-7 Add Server options in a CASO environment (*continued*)

Item	Description
Logon account to use to connect to the media servers	Shows the name of the logon account that is used to access the media servers you selected.
Change Logon Account	Lets you change the logon account that Backup Exec uses to access the media servers you selected.

Copy Settings options

On the **Copy Settings** dialog box, you can select the type of settings to copy to another media server.

See [“Copying configuration settings to another media server”](#) on page 199.

Table 3-8 Copy Settings options

Item	Description
Default job options	Lets you copy default job options from this media server to another media server.
Default schedule	Lets you copy the default schedule settings from this media server to another media server.
Error-handling rules	Lets you copy error-handling rules from this media server to another media server.
Alert configuration	Lets you copy the alert configuration from this media server to another media server.
Add	Lets you add a media server to the Destination media servers list. After you add a media server to the list, you can copy settings to it.
Edit	Lets you change the logon account that is used to connect to the selected media server.
Remove	Lets you remove the selected media server from the Destination media servers list.

Table 3-8 Copy Settings options (*continued*)

Item	Description
Import List	Lets you import a list of media servers to the Destination media servers list. After you add media servers to the list, you can copy settings to them.

Copying logon account information

You can copy logon account information to a different media server.

To copy logon accounts information

- 1 From the **Network** menu, click **Logon Accounts**.
- 2 Select the logon account information you want to copy, and then click **Copy to Servers**.
- 3 In the **Server Name** field, enter the name of the media server you want to copy the logon account information to, and then click **Add**.
- 4 Click **OK**.

Copy Logon Account options

You can copy logon account information to a different media server.

See [“Copying logon account information”](#) on page 204.

Table 3-9 Copy Logon Account options

Item	Description
Server Name	Indicates the name of the media server you want to copy the logon account information to, and then click Add .
Add	Adds the media server from the Server Name field to the list of media servers.
Remove	Removes a media server from the list.
Import List	Imports a list of media servers to be added to the media servers in the list. The list should include only the media server name, with one per line.
Logon Account	Specifies the logon account to use when connecting to the media servers in the list.

Table 3-9 Copy Logon Account options (*continued*)

Item	Description
If an account with this description already exists on the destination server, overwrite it	Overwrites logon accounts for an existing job having the same name. This option only appears if you are copying a job to another media server.

About audit logs

Use audit logs to examine and review information about operations that have been performed in Backup Exec. The audit log displays the date and time of the activity, who performed it, what the activity was, and a description of the activity.

You can view information about activities that occur for all or any of the following:

- Alerts
- Audit logs
- Devices and media
- Encryption keys
- Error-handling rules
- Jobs
- Logon accounts
- Policies and job templates
- Selection lists
- Server configuration

You can delete the audit logs as part of the Backup Exec database maintenance, and you can save the audit log to a file. Changes made to the audit log, such as when database maintenance occurs, can also be displayed in the audit log.

See [“Configuring the audit log”](#) on page 206.

See [“Viewing the audit log”](#) on page 206.

See [“Removing entries from the audit log”](#) on page 208.

See [“Saving the audit log to a file”](#) on page 208.

Configuring the audit log

Configure the audit log to display information about specific operations that are performed on items in Backup Exec.

See [“About audit logs”](#) on page 205.

See [“Viewing the audit log”](#) on page 206.

To configure the audit log

- 1 On the **Tools** menu, click **Audit Log**.
- 2 Click **Configure Logging**.
- 3 On the **Audit Log Configuration** dialog box, select the check box of the category that you want to display in the audit log.

Expand the category by clicking the plus sign to the left of the category. Select the operations that you want to display for the category.

Clear the check box of any item or operation that you do not want to display.

- 4 Click **OK**.

Viewing the audit log

You can view audit logs to see when changes were made in Backup Exec and which users made the changes.

See [“Configuring the audit log”](#) on page 206.

To view the audit log

- 1 On the **Tools** menu, click **Audit Log**.
- 2 In **Select category to view**, select the category for which you want to view audit information.

See [“Audit Logs options”](#) on page 207.

- 3 Use the scroll bar at the bottom of the Audit log window to view the whole entry, or double-click the entry to display the same information in an easy-to-read Audit Log Record.

Audit Log Record options

You can view audit logs to see when changes were made in Backup Exec and which users made the changes.

See [“About audit logs”](#) on page 205.

Table 3-10 Audit Log Record options

Item	Description
Date/Time	Shows the date and time that this change was made in Backup Exec.
User Name	Shows the domain and the user name of the user that made the change.
Category	Shows the category to which the log belongs.
Message	Shows the action that was recorded by Backup Exec for the operation that was performed.

Audit Logs options

You can view audit logs to see when changes were made in Backup Exec and which users made the changes.

See [“Viewing the audit log”](#) on page 206.

See [“Removing entries from the audit log”](#) on page 208.

See [“Saving the audit log to a file”](#) on page 208.

Table 3-11 Audit Logs options

Item	Description
Select category to view	Lets you select the category for which you want to view audit logs.
Date/Time	Shows the date and time that this change was made in Backup Exec. Click the column head to sort the information by date.
User Name	Shows the domain and the user name of the user that made the change. Click the column head to sort the information alphabetically.
Category	Shows the category to which the log belongs. Click the column head to sort the information alphabetically.
Message	Shows the action that was recorded by Backup Exec for the operation that was performed. Click the column head to sort the information alphabetically.

Table 3-11 **Audit Logs** options (*continued*)

Item	Description
Refresh	Updates the audit log with new entries.
Clear Category Log	Removes all entries from an audit log category.
Save Log to File	Indicates where to save audit log entries. You can save the audit log as a text (.txt) file.
Properties	Provides information about the selected entry.
Configure Logging	Lets you select the categories and options to include in the audit log.

Removing entries from the audit log

You can remove the entries for all categories or for a selected category.

See “[About audit logs](#)” on page 205.

To remove entries from the audit log

- 1 On the **Tools** menu, click **Audit Log**.
- 2 In **Select category to view**, select the category for which you want to view audit information.
- 3 Click **Clear Category Log** to remove all entries from an audit log category.

If you select specific categories to view, only the logs generated for the selected categories are cleared when you click **Clear Category Log**.

Saving the audit log to a file

You can save the audit log as a text (.txt) file.

See “[About audit logs](#)” on page 205.

To save the audit log to a file

- 1 On the **Tools** menu, click **Audit Log**.
- 2 Click **Save Log to File** to specify a file name and location to save the audit log entries.

About database maintenance

The Database Maintenance option enables you to manage the Backup Exec database and the Desktop and Laptop Option (DLO) database. Each database maintenance operation is performed independently on each database. The Backup Exec database maintains a record of files and data you have configured such as templates and catalogs.

Database maintenance enables you to perform the following:

- Optimize database size.
- Delete expired data.
- Save the contents of the database files.
- Perform a database consistency check.

Informational alerts are generated at the beginning and the end of the database maintenance process each time database maintenance is performed. The alerts provide details about the type of maintenance that was performed on each database and the amount of time the maintenance took to complete. If the database maintenance process fails, the alert indicates where the failure occurred and the reason for the failure.

See [“Configuring database maintenance”](#) on page 209.

Configuring database maintenance

The Database Maintenance option enables you to manage the Backup Exec database and the Desktop and Laptop Option (DLO) database. Each database maintenance operation is performed independently on each database. The Backup Exec database maintains a record of files and data you have configured such as templates and catalogs.

You do not have to select all the options; however, each one performs a different process that enables you to protect and maintain your database. Selecting all the options enables you to recover the database quickly and maintain optimal performance.

See [“About database maintenance”](#) on page 209.

To configure database maintenance

- 1 On the **Tools** menu, click **Options**.
- 2 Under **Settings**, click **Database Maintenance**.
- 3 Select the appropriate options, and then click **OK**.

See [“Application settings for database maintenance”](#) on page 210.

Application settings for database maintenance

You can manage the Backup Exec Database and the Desktop and Laptop Option (DLO) database.

See “[Configuring database maintenance](#)” on page 209.

Table 3-12 Application settings for database maintenance

Item	Description
Enable Backup Exec database maintenance	Activates the database maintenance process.
Last time maintenance was performed	Indicates the date and time the last database maintenance was performed.
Perform database maintenance at	Indicates the time you want to perform database maintenance. All the maintenance will occur once a day at the time you specify.
Delete aged data	<p>Activates the deletion of expired job history, job logs, alert history, and reports from the Backup Exec Database after the specified number of days have passed.</p> <p>For the Desktop and Laptop Option (DLO) database, only the Alert History setting applies. DLO does not have job history, job logs, or reports.</p>
Keep job history for data on media that have current overwrite protection periods	Keeps all job history data for any media to which an overwrite protection policy is currently assigned. After a media’s overwrite protection policy expires, the media’s job history data can be deleted.
Keep job history for specified number of days	Indicates the number of days to keep job history data in the database before it is deleted. Job history data includes summary statistics for a job and details about media, devices, and backup sets that were used to process the job.
Job logs	Indicates the number of days to keep job logs in the database before they are deleted. Job logs include detailed information about the job.
Alert history	Indicates the number of days to keep alert history data in the database before it is deleted. Alert history data includes property and response information for the alert.

Table 3-12 Application settings for database maintenance (*continued*)

Item	Description
Reports	Indicates the number of days to keep report data in the database before it is deleted. Report data includes property information about report jobs that were generated. The report itself is not deleted.
Audit logs	Indicates the number of days to keep audit log data in the database before it is deleted. The audit log includes information about operations that are performed in Backup Exec. See “About audit logs” on page 205.
Perform database consistency check	Checks the logical and physical consistency of the data in the database. The option is not checked by default. It is recommended that you run a consistency check periodically at a time when there is minimal activity from Backup Exec.
Save contents of database to the Backup Exec data directory	Places the data that is contained in the database into the Backup Exec data directory so that the database backup file (BEDB.bak) can be backed up. The dump file will be maintained in the data directory until the next database maintenance process is performed and then this file will be overwritten. Selecting this option enables you to recover the database in the event of failure.
Optimize database size	Reorganizes fragmented pages and decrease the size of the physical database to 10 percent above what is actually used.

Viewing the location of Backup Exec databases

On the media server’s advanced properties, you can view information about the location of the databases for Backup Exec, which include the Backup Exec database, the device and media database (ADAMM), and the catalog database.

During Backup Exec installation, if you chose the default option to create a local Backup Exec SQL Express instance on which to store the Backup Exec database, the databases are all located on the local media server. If you chose another instance on the network on which to store the Backup Exec database, then the databases are all located on the Microsoft SQL Server that contains that instance.

In a SAN SSO configuration, on the secondary servers, the database locations are the same as the database locations displayed on the primary server.

In a Central Admin Server Option configuration, if the device and media database location is on the central administration server, that information is displayed.

See [“How CASO works”](#) on page 1504.

Note: Advanced properties are displayed only for the media server that the Backup Exec Administration Console is connected to.

To view the location of Backup Exec databases

- 1 On the navigation bar, click **Devices**.
- 2 Select the media server from the tree view.
- 3 Under **General Tasks** in the task pane, select **Properties**.
- 4 On the **Advanced** tab, view the properties.

See [“Advanced properties for a media server”](#) on page 212.

Advanced properties for a media server

On the media server’s advanced properties, you can view information about the location of the databases for Backup Exec, which include the Backup Exec database, the device and media database (ADAMM), and the catalog database.

See [“Viewing the location of Backup Exec databases”](#) on page 211.

Table 3-13 Advanced properties for a media server

Item	Description
Server	Shows the name of the Microsoft SQL Server that contains the Backup Exec database.
Instance	Shows the name of the instance that the Backup Exec database is installed on.
Name	Shows the Microsoft SQL Server name of the Backup Exec database.
Path	Shows the path of the Backup Exec database.
Server	Shows the name of the Microsoft SQL Server that contains the Advanced Device and Media Management (ADAMM) database.
Instance	Shows the name of the instance that the Advanced Device and Media Management (ADAMM) database is installed on.

Table 3-13 Advanced properties for a media server (*continued*)

Item	Description
Name	Shows the Microsoft SQL Server name for the Advanced Device and Media Management (ADAMM) database.
Path	Shows the path of the Advanced Device and Media Management (ADAMM) database.
Server	Shows the name of the Microsoft SQL Server that contains the Backup Exec catalog database.
Instance	Shows the Database instance that contains the catalog database.
Name	Shows the Microsoft SQL Server name for the Backup Exec catalog database.
Path	Shows the path of the Backup Exec catalog database.

About the Symantec RSS Reader

The Symantec RSS Reader is on the **Home** view on the Backup Exec Administration Console. You can customize the Symantec RSS Reader and select the default Backup Exec feeds that display in the reader. You can also add additional Symantec RSS feeds to the reader. The Symantec RSS Reader refreshes the RSS feeds every 15 minutes when the item is open in the **Home** view.

See [“About the Home view”](#) on page 95.

See [“Viewing an article in the Symantec RSS Reader”](#) on page 213.

See [“Customizing the Symantec RSS Reader”](#) on page 214.

See [“Removing a default RSS feed from the Symantec RSS reader”](#) on page 215.

Viewing an article in the Symantec RSS Reader

The Symantec RSS Reader sorts articles by the date and the time. The reader displays the last entry of an article in the RSS feed, however you can also choose to view the full article.

See [“About the Symantec RSS Reader”](#) on page 213.

To view an article in the Symantec RSS Reader

- 1 On the navigation bar, click **Home**.
- 2 In the task pane, under **Help and Technical Support Items**, select the **Symantec RSS Reader** check box.

- 3 In the **Symantec RSS Reader**, click the arrow next to the RSS feed that contains the article.
- 4 Click the hyperlink for the article that you want to open.
The Symantec RSS Reader opens a new window that contains a portion of the article from the RSS feed.
- 5 Click **Go to full Article** to open Internet Explorer and view the entire contents of the article.

Customizing the Symantec RSS Reader

You can customize and add Symantec RSS feeds to the reader in addition to the default Backup Exec RSS feeds. You can only add Backup Exec and Symantec RSS feeds to the Symantec RSS Reader.

See [“About the Symantec RSS Reader”](#) on page 213.

To customize the Symantec RSS feeds to the Symantec RSS Reader

- 1 On the navigation bar, click **Home**.
- 2 In the task pane, under **Help and Technical Support Items**, select the **Symantec RSS Reader** check box.
- 3 In the **Symantec RSS Reader**, click the pencil icon to add an RSS feed.
- 4 Type the URL and the name of the RSS feed that you want to add. Or click **More** to view additional Symantec RSS feeds.

See [“Add Symantec RSS Feed options”](#) on page 214.

- 5 Click **OK**.

Add Symantec RSS Feed options

You can enter property information for an RSS feed that you want to add to the Symantec RSS Reader.

See [“Customizing the Symantec RSS Reader”](#) on page 214.

Table 3-14 Symantec RSS Feed options

Item	Description
URL	Indicates the location of the RSS feed that you want to add to the Symantec RSS Reader.
Name	Indicates the name of the RSS feed that you want to display in the Symantec RSS Reader.

Table 3-14 Symantec RSS Feed options (*continued*)

Item	Description
More	Shows a list of Symantec RSS feeds that you can add to the Symantec RSS Reader.

Removing a default RSS feed from the Symantec RSS reader

You can remove the RSS feeds that display in the Symantec RSS reader. If the RSS feed is not open in the reader, the RSS feed does not refresh.

See [“About the Symantec RSS Reader”](#) on page 213.

To remove an RSS feed from the Symantec RSS reader

- 1 On the navigation bar, click **Home**.
- 2 Do one of the following:

To remove a default Backup Exec RSS feed Select the check box of the Backup Exec RSS feed.

To remove an RSS feed that you added to the Symantec RSS reader Click the red X next to the name of the RSS feed.

Hiding columns

Backup Exec uses panes to present information to you in a structured and organized manner. Each pane contains several columns that specify the type of information that is displayed. You can remove the columns that may not interest you.

See [“Showing a hidden column”](#) on page 215.

To hide columns

- 1 Right-click any column title.
- 2 Click **Configure Columns**.
- 3 Click the column title you want to hide.
- 4 Click **Hide**.

Showing a hidden column

If you hide a column, you can show it again at any time.

See “[Hiding columns](#)” on page 215.

To show a hidden column

- 1 Right-click any column title.
- 2 Click **Configure Columns**.
- 3 Select a column that you want to show.
- 4 Click **Show**.

Rearranging columns

You can change the position of columns to suit your needs. You can also change the size of the columns to better match the size of the information in the columns.

To rearrange columns

- 1 Right-click any column title.
- 2 Click **Configure Columns**.
- 3 Select a column title, and then click **Move Up** or **Move down**.

Each click of the Move Up option moves the column name one column to the left in the pane, while each click of the Move Down option moves the selected column name one column to the right.

- 4 If you want to change the width of a column, do the following:
 - Select the column.
 - In the Width of selected column (in pixels) field, enter the column width.
- 5 Click **OK**.

Sorting column information

You can choose the order in which Backup Exec sorts the information in columns.

To sort column information

- 1 Right-click any column title.
- 2 Click **Multi-Column Sort**.
- 3 In the **Sort by** list, select the column titles on which you want to sort information.
- 4 Click **Ascending** to sort the information in ascending order or click **Descending** to sort the information in descending order.

- 5 To sort by additional columns, repeat step 3 and step 4 in the **Then by** lists.
- 6 Click **OK**.

Viewing properties

Properties provide detailed information, such as statistics, dates, and settings.

To view properties

- ◆ Do one of the following:
 - Right-click the item for which you want to view properties, and then click **Properties**.
 - Select the item for which you want to view properties, and then in the task pane under **General Tasks**, click **Properties**.

Managing media

This chapter includes the following topics:

- [About media in Backup Exec](#)
- [About media overwrite protection](#)
- [Selecting settings for media management](#)
- [Viewing audit log entries for media operations](#)
- [Configuring specific media operations to appear in the audit log](#)
- [Media labeling](#)
- [About WORM media](#)
- [Creating a new catalog](#)
- [Creating a restore job while reviewing media or devices](#)
- [Media locations and vaults](#)
- [About moving media to a vault or to the offline media location](#)
- [About removing damaged media](#)
- [General properties for media](#)
- [Statistics properties for media](#)
- [Media rotation strategies](#)

About media in Backup Exec

With Backup Exec's media management tools you can do the following:

- Protect data from being overwritten.
- Set up media rotation strategies.
- Track the location of media.
- Label media automatically.
- Read and track media labels with bar codes.
- Collect and report media statistics.

With Backup Exec, you are not required to select media for jobs; it is done for you by the Advanced Device and Media Management (ADAMM) component. Backup Exec tracks all media that is loaded into attached storage devices, media that is offline, and media that has been placed in media vaults.

The following table lists and describes the **Media** nodes in the **Media** view.

Table 4-1 A description of media nodes in the Media view

Media node	Media node description
All Media	Lists all media that has been introduced into Backup Exec. Any media that is available for overwriting in backup operations, such as scratch or recyclable media, displays in blue. See “General properties for media sets” on page 230.
Media Sets	Lists default system media sets and media sets that you create. A media set is a set of rules that manage media. These rules include the append and the overwrite protection periods, and vault rules, which allow you to set dates for when media should be moved to or returned from a media vault Media that are associated with a media set are allocated media. Allocated media have current append and overwrite protection periods. Media that are associated with a media set, but have expired overwrite protection periods are recyclable media. See “About creating media sets” on page 226.
Cleaning Media	Lists all cleaning media. See “Defining a cleaning slot ” on page 471.

Table 4-1 A description of media nodes in the Media view (*continued*)

Media node	Media node description
Imported Media	<p>Lists all media created by a product other than this installation of Backup Exec. By default, imported media have an overwrite protection period of Infinite, but can still be overwritten if the media overwrite protection level is set to Partial or None. You can overwrite imported media using several methods. Data can be restored from imported media until that media is overwritten.</p> <p>See “Selecting settings for media management” on page 237.</p>
Backup Exec and Windows NT Backup Media	<p>Lists all media from another installation of Backup Exec.</p> <p>See “Creating a new catalog” on page 248.</p>
Foreign Media	<p>Lists all media from a product other than Backup Exec.</p> <p>See “Creating a new catalog” on page 248.</p>
Retired Media	<p>Lists all media that you have taken out of service, usually because of an excessive number of errors. After a media has been associated with the retired media set, it is not selected for backup jobs by Backup Exec. It is still available for restore operations, if it has not been damaged. Retired Media protects media from being used (overwritten).</p> <p>You can delete media that is in Retired Media to remove it from Backup Exec. You may want to delete media, for example, when you have a lot of offsite media that you do not want to recycle or if you throw away the media. If you decide to use deleted media in Backup Exec, it is recognized as Imported Media and must be cataloged before you can restore from it.</p> <p>See “About removing damaged media” on page 259.</p>
Scratch Media	<p>Lists all media that can be overwritten. New, blank, and erased media are automatically associated with the Scratch Media set.</p> <p>See “About media overwrite protection” on page 223.</p>

Table 4-1 A description of media nodes in the Media view (*continued*)

Media node	Media node description
Keep Data for 4 Weeks	<p>Lists all media that you associate with this media set. As long as you use the backup job defaults that are set when you install Backup Exec, the media set Keep Data for 4 Weeks is the default media set for all backup jobs. This media set protects data from being overwritten for four weeks and allows the media to be appended to for six days.</p> <p>You can edit and rename Keep Data for 4 Weeks after installation. Therefore, it may not continue to appear in the Media view or in the backup job defaults as Keep Data for 4 Weeks.</p> <p>See “About the default media set” on page 226.</p> <p>See “About media overwrite protection” on page 223.</p>
Keep Data Infinitely - Do Not Allow Overwrite	<p>Lists all media that you associate with this media set.</p> <p>When you associate media with this media set, data is not overwritten unless you perform any of the following actions on the media:</p> <ul style="list-style-type: none"> ■ Erase. ■ Label. ■ Format. ■ Associate the media with the scratch media set. <p>You can append data to this media for an infinite period (until the media is full).</p> <p>You can edit and rename Keep Data Infinitely - Do Not Allow Overwrite after installation. Therefore, it may not continue to appear in the Media view or in the backup job defaults as Keep Data Infinitely - Do Not Allow Overwrite.</p> <p>See “About the default media set” on page 226.</p> <p>See “About media overwrite protection” on page 223.</p>
Media Location	<p>Lists the location of media when it is online, offline, or in a user-defined media vault.</p> <p>See “Media locations and vaults” on page 250.</p>

See [“Creating media sets by using the Media Set Wizard”](#) on page 227.

See [“About creating media sets”](#) on page 226.

See [“Associating media with a media set”](#) on page 229.

About media overwrite protection

Each media is associated with a media set, which is a set of rules that manage media.

These rules include the following:

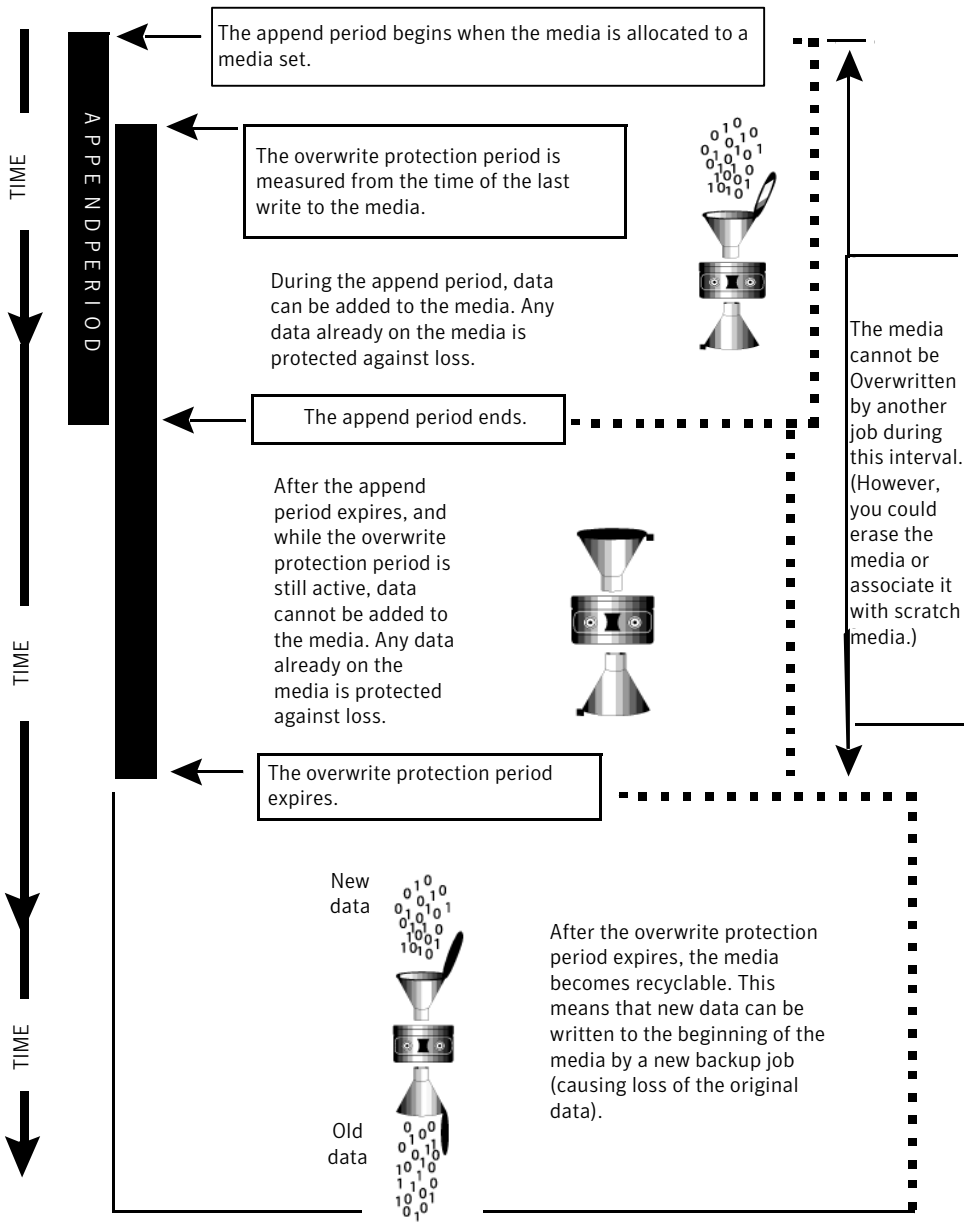
Table 4-2 Rules specified in the media set

Rule	Description
Append period	The amount of time that data can be appended (added) to media. It is measured from the time the media was first allocated. It can be specified in hours, days, weeks, or years.
Overwrite protection period	<p>The amount of time that media is protected from being overwritten. It is measured from the time of the last write to the media, that is, at the end of the last append or overwrite job. It can be specified in hours, days, weeks, or years. When the overwrite protection period is over, the media becomes recyclable and can be overwritten.</p> <p>The overwrite protection period begins when the backup job is completed. If there is an append period, the overwrite protection period begins again each time an append job completes. Because the overwrite protection period does not begin until the job completes, the amount of time that the job takes to complete affects the amount of time until the media can be overwritten. You may shorten the overwrite protection period to take into account the amount of time a job may run.</p> <p>For example, setting the overwrite protection period for seven days and the append period for four days ensures that data will not be overwritten for at least seven days, and that data can be appended to the media for the next four days. The last data appended to this media is retained for at least seven days.</p>

Your media rotation strategy must balance between your need to save useful data as long as possible, and the fact that media are not in infinite supply. The compromise between the longevity of stored backup data and the cost of more media is controlled in Backup Exec by the rules specified in the media set, which allows Backup Exec to identify which media can be written to and which media is overwrite-protected.

The following graphic shows the relationship between the append period and overwrite protection periods.

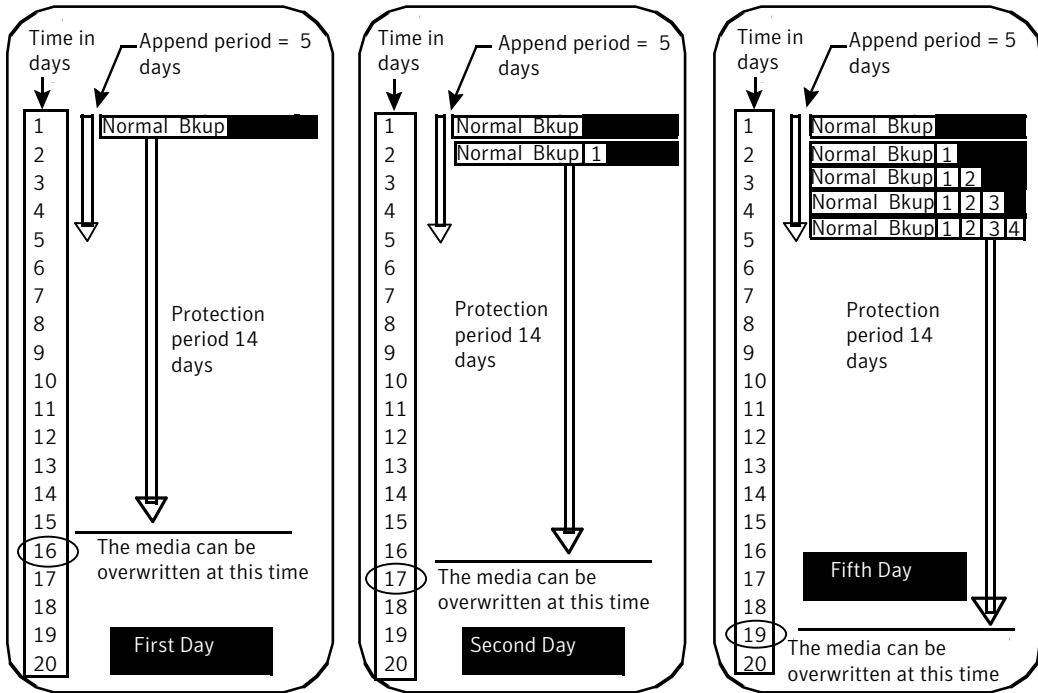
Figure 4-1 Append periods and overwrite protection periods



The append and overwrite protection periods that you specify apply to all the data on the media.

Each time data is written to a media, the time remaining in the overwrite protection period is reset and the countdown restarted.

Figure 4-2 How overwrite protection periods are reset



Because the overwrite protection period does not begin until the job completes, the amount of time that the job takes to complete affects the amount of time until the media can be overwritten.

For example, suppose that you create a media set named Weekly with an overwrite protection period of seven days, and an append period of 0 days, and you schedule a full backup job to run each Friday at 20:00. When it is time for the full backup to run at 20:00 the following Friday, the job cannot run because the first backup job that ran the previous Friday did not complete until 21:10 p.m. The overwrite protection period for the Weekly media set still has 70 minutes remaining.

Typically, to prevent this situation, you would shorten the overwrite protection period to account for the amount of time a job may run. For this example, the scheduled job recurring at 20:00 can run if the overwrite protection period is set to 6 days instead of 7 days.

About the default media set

When you install Backup Exec, a default media set named **Keep Data for 4 Weeks** is created. All of the media that you use for the backup jobs that you create are automatically associated with the media set **Keep Data for 4 Weeks**. The backup data is protected from overwrite for four weeks. The media can be appended to for six days.

You can change the media set by doing one of the following:

- Create new media sets that have the append and overwrite protection periods set to the time intervals that accommodate your data retention strategy. Then, specify the media set that is most appropriate when you create a backup job. For example, you can create a media set that keeps data for 60 days, and a media set that keeps data for 90 days.
- Select the media set **Keep Data Infinitely - Do Not Allow Overwrite**. The risk that is associated with the media set **Keep Data Infinitely - Do Not Allow Overwrite** is that you can use all of your scratch media. You must continually add new media to Backup Exec.

Note: Symantec recommends that if you need to keep data longer than four weeks, create a duplicate backup job. A duplicate backup job copies the backup data from the original storage device to tape, which you can then send for long-term or off-site storage.

See [“About duplicating backed up data”](#) on page 369.

See [“About creating media sets”](#) on page 226.

See [“Associating media with a media set”](#) on page 229.

About creating media sets

A media set consists of rules that specify append periods, overwrite protection periods, and vaulting periods.

When you create a new media set, you specify an append period and an overwrite protection period for the set. When media is associated with the media set, the append and overwrite protection periods apply to that media.

You can also specify vault rules, which allow you to set dates for when media should be reported as ready to be moved to or returned from a media vault. The **Vault Wizard** logically moves the media, and then exports the media, but you must physically move the media to and from the vault. You can use the **Vault Wizard** to print or view reports that contain details on which media are ready to be moved, and to update the media location.

If your environment includes remote sites, you should create separate media sets for each remote site, so that if vault rules are enabled, the reports contain details on which media are ready to be moved for just that site.

Note: Even if the overwrite protection period is current, media can still be overwritten if the overwrite protection level is set to **None**.

See [“Creating media sets”](#) on page 227.

See [“Media overwrite protection levels”](#) on page 232.

See [“About media overwrite protection”](#) on page 223.

See [“Deleting a media set”](#) on page 228.

See [“Renaming a media set”](#) on page 228.

See [“Editing general properties for media sets”](#) on page 230.

See [“Configuring vault rules for media sets”](#) on page 252.

See [“Using the Vault Wizard to move media”](#) on page 256.

Creating media sets

Create a media set to set up the rules that specify append periods, overwrite protection periods, and vaulting periods for media.

To create a media set

- 1 On the navigation bar, click **Media**.
- 2 Under **Media Set Tasks** in the task pane, click **New Media Set**.
- 3 On the **General** tab, select the appropriate options, and then click **OK**.
See [“General properties for media sets”](#) on page 230.
- 4 On the **Vault Rules** tab, select the appropriate options, and then click **OK**.
See [“Properties for vault rules for media sets”](#) on page 252.

Creating media sets by using the Media Set Wizard

The **Media Set Wizard** guides you through the process of creating a new media set, and may be helpful if you are new to Backup Exec or are unfamiliar with the concepts and terminology of media sets. Before you run this wizard, you should understand media overwrite protection and append periods.

See [“About media overwrite protection”](#) on page 223.

If you do not need a wizard to set up a new media set, you can create it manually.

See [“About creating media sets”](#) on page 226.

If you set the media overwrite protection level to Full, you must have blank media ready and online before the first backup job runs. If you plan to use imported media, you must inventory it first.

To create media sets by using the Media Set Wizard

- 1 On the **Tools** menu, point to **Wizards**.
- 2 Click **Media Set Wizard**, and then follow the instructions.

Deleting a media set

Use **Delete** to remove a media set from the **Media Sets** category. If you delete a media set to which scheduled jobs are targeted, you are prompted to retarget the jobs to another media set.

You cannot delete a media set that has associated media. You must associate the media with another media set first.

Caution: Make sure that the media set to which you associate the media has the appropriate overwrite protection and append periods.

To delete a media set

- 1 On the navigation bar, click **Media**.
- 2 In the **Media** selection pane, under **Media Sets**, select the media set that you want to delete.
- 3 Under **General Tasks** on the task pane, click **Delete**.
- 4 When prompted to delete the media set, click **OK**.
- 5 If there are scheduled jobs allocated to the deleted media set, you are prompted to redirect the jobs to another media set.

See [“Retarget Job options ”](#) on page 520.

Renaming a media set

When you rename a media set, any jobs that use that media set will display the new media set name.

To rename a media set

- 1 On the navigation bar, click **Media**.
- 2 In the **Media** selection pane, under **Media Sets**, select the media set you want to rename.

- 3 Under **General Tasks** on the task pane, click **Rename**.
- 4 In the **Name** field, type the new name you want to assign to this media set, and then click **OK**.

Associating media with a media set

When you associate media with a media set, the media uses the append and overwrite protection period properties of that media set.

Note: Associating scratch or imported media with a media set is not recommended. Backup Exec automatically associates scratch or imported media with a media set when a backup job requires it.

To associate media with a media set

- 1 On the navigation bar, click **Media**.
- 2 Expand **All Media** to display a list of media.
- 3 Select the media that you want to associate with a media set.
- 4 Do one of the following:
 - Drag the media to the media set.
 - Under **Media Tasks** on the task pane, click **Associate with media set**, select a media set to associate the media with, and then click **Yes** or **Yes to all**.

Associate Media with Media Set options

The overwrite protection period, the append period, and any configured vault rules for a media may change when you associate it with a different media set.

See “[Associating media with a media set](#)” on page 229.

Table 4-3 Options to associate media with a media set

Item	Description
Name	Displays the name of the media set that the selected media is currently associated with.
Description	Displays a description of the media.
Associate with	Displays the name of the media set with which you want to associate the selected media.

Editing general properties for media sets

On general media set properties, you can change the following:

- Name of a media set
- Overwrite protection and append periods for a media set
- Media vault and the vaulting periods associated with a media set.

To edit the general properties for media sets

- 1 On the navigation bar, click **Media**.
- 2 In the **Media** selection pane, under **Media Sets**, select a media set.
- 3 Under **General Tasks** on the task pane, click **Properties**.
- 4 To change the media set name or overwrite or append periods, click the **General** tab.
See [“General properties for media sets”](#) on page 230.
- 5 Select the appropriate options, and then click **OK**.

General properties for media sets

General properties for media sets provide information about the retention period and the append period for the media.

See [“Creating media sets”](#) on page 227.

See [“Viewing properties”](#) on page 217.

Table 4-4 General properties for media sets

Item	Description
Name	Displays the name of the media set.
Creation date	Displays the date and time when the media set was created. Backup Exec sets the date and time automatically. You cannot change them.

Table 4-4 General properties for media sets (*continued*)

Item	Description
Overwrite protection period	<p>Displays the length of time in hours, days, weeks, or years to retain the data on the media before the media can be overwritten.</p> <p>Note: Regardless of the overwrite protection period that is set, media can be overwritten if it is erased, formatted, labeled, associated with Scratch Media, or if the Media Overwrite Protection Level is set to None</p> <p>Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.</p> <p>The default is Infinite - Don't Allow Overwrite, which protects the media from being overwritten for 1,000 years, unless the media is erased, formatted, labeled, moved to Scratch Media, or if the media overwrite protection level is set to None.</p> <p>See “About media overwrite protection” on page 223.</p>
Append period	<p>Displays the length of time in hours, days, or weeks, that data can be added to media. Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.</p> <p>The append period starts when the first backup job is written to this media.</p> <p>The default is Infinite - Allow Append, which allows data to be appended until the media capacity is reached.</p>

- See [“Creating media sets”](#) on page 227.
- See [“About creating media sets”](#) on page 226.
- See [“Deleting a media set”](#) on page 228.
- See [“Renaming a media set”](#) on page 228.
- See [“Editing general properties for media sets”](#) on page 230.
- See [“Configuring vault rules for media sets”](#) on page 252.
- See [“Using the Vault Wizard to move media”](#) on page 256.

Media overwrite protection levels

The media overwrite protection level is a global setting that supersedes the media set's overwrite protection period. Although the terms are similar, the media overwrite protection level and the media overwrite protection period are different. The media overwrite protection period is a time interval that changes from one media set to another. The media overwrite protection level specifies whether to overwrite scratch, imported, or allocated media, regardless of the media's overwrite protection period.

Use the media overwrite protection level to specify the type of media, such as scratch or imported media, that you want to be available for overwrite backup jobs.

The options for the media overwrite protection level are as follows:

- **Full** - Overwrites scratch media, which are media that contain data you are willing to discard, and recyclable media, which are media that are associated with media sets and have expired overwrite protection periods.
- **Partial** - Overwrites imported media, which are media that was created by another installation of Backup Exec or some other backup product, and overwrite scratch media.
- **None** - Overwrites all media, including those that have current overwrite protection periods (allocated media).

Caution: The None option is not recommended. It does not protect data from being overwritten.

See [“Selecting settings for media management”](#) on page 237.

See [“Media locations and vaults”](#) on page 250.

See [“About media overwrite protection”](#) on page 223.

About overwriting allocated or imported media

Backup Exec protects allocated and imported media from being overwritten when full or partial overwrite protection is used. However, if necessary, you can allow allocated and imported media to be overwritten by Backup Exec before the data overwrite protection period expires, and without setting the media overwrite protection level to None.

The following methods are available:

- Move the media to **Scratch Media**. The media is overwritten when it is selected for an overwrite job.

- Erase the media. Erased media is automatically recognized as scratch media and will be overwritten immediately.
- Label the media. The Label Media operation immediately writes a new media label on the media, which destroys any data contained on the media.
- Format the media. Formatting destroys any data contained on the media.
- Change the overwrite protection period for the media set so that it is expired.

See [“About media in Backup Exec”](#) on page 219.

See [“About deleting media”](#) on page 260.

See [“Editing general properties for media sets”](#) on page 230.

See [“How Backup Exec searches for overwritable media”](#) on page 233.

How Backup Exec searches for overwritable media

Media overwrite options set the order in which Backup Exec searches for overwritable media. When Backup Exec searches for overwritable media for a backup job, it searches for either scratch media or media that has an expired overwrite protection period.

You are prompted to select one of the following types of media that you want Backup Exec to use first:

- Overwrite scratch media before overwriting recyclable media contained in the targeted media set.
If you choose to overwrite scratch media before recyclable media, more media may be required for the same number of jobs, but the recyclable media may be preserved longer for possible recovery.
- Overwrite recyclable media contained in the targeted media set before overwriting scratch media.
If you choose to overwrite recyclable media before scratch media, you will re-use the same media more frequently than if you choose to overwrite scratch media before recyclable media.

In a device pool, Backup Exec selects the oldest recyclable media in all of the devices in the device pool to use first.

In a robotic library, Backup Exec selects the oldest recyclable media in the library to use first. If the robotic library is partitioned, Backup Exec searches for the oldest recyclable media in the targeted partition only.

Caution: It is recommended that you physically write-protect media containing critical data by using the write-protect tab on the media cartridge to protect against unintentional move or erase operations, or expired overwrite protection periods.

The following table describes the order in which Backup Exec searches for media to use for an overwrite job, depending on the combination of the overwrite protection level and the media overwrite option you select.

Table 4-5 How Backup Exec searches for overwritable media

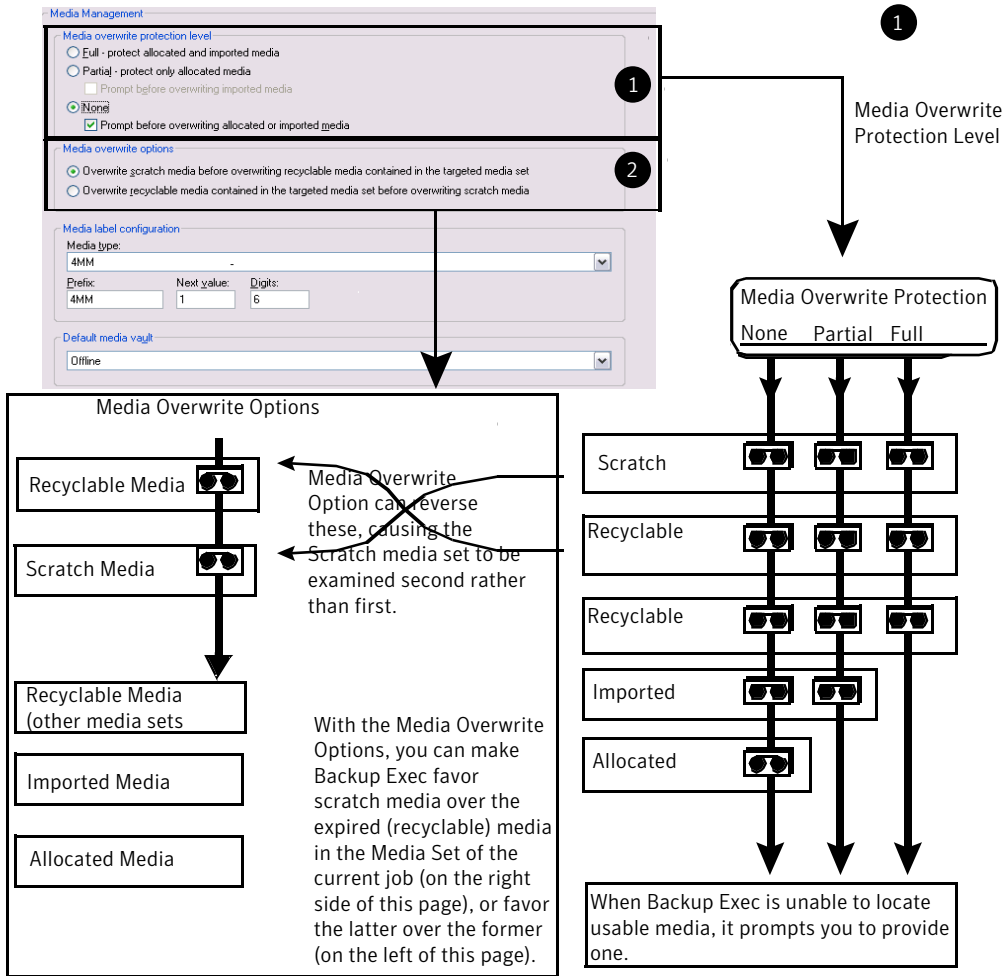
Overwrite protection level and overwrite option:	Media is overwritten in this order:
Full + Overwrite scratch media first Note: This combination provides the most protection against overwriting media.	<ul style="list-style-type: none"> ■ Scratch media ■ Recyclable media in the targeted media set ■ Recyclable media in any media set
Full + Overwrite recyclable media first	<ul style="list-style-type: none"> ■ Recyclable media in the targeted media set ■ Scratch media ■ Recyclable media in any media set
Partial + Overwrite scratch media first	<ul style="list-style-type: none"> ■ Scratch media ■ Recyclable media in the targeted media set ■ Recyclable media in any media set ■ Imported media
Partial + Overwrite recyclable media first	<ul style="list-style-type: none"> ■ Recyclable media in the targeted media set ■ Scratch media ■ Recyclable media in any media set ■ Imported media
None - No overwrite protection + overwrite scratch media first Warning: This options is not recommended because it does not protect data from being overwritten.	<ul style="list-style-type: none"> ■ Scratch media ■ Recyclable media in the targeted media set ■ Recyclable media in any media set ■ Imported media ■ Allocated media in any media set

Table 4-5 How Backup Exec searches for overwritable media (*continued*)

Overwrite protection level and overwrite option:	Media is overwritten in this order:
<p>None - No overwrite protection + overwrite recyclable media first</p> <p>Warning: This options is not recommended because it does not protect data from being overwritten.</p>	<ul style="list-style-type: none"> ■ Recyclable media in the targeted media set ■ Scratch media ■ Recyclable media in any media set ■ Imported media ■ Allocated media in any media set

In addition to setting overwrite protection levels, you must set overwrite options, which set the order in which Backup Exec searches for overwritable media.

Figure 4-3 Media overwrite protection



The most obvious candidates for backup jobs requiring overwritable media are scratch media and recyclable media (media with expired overwrite protection periods). These are the first types of media for which Backup Exec searches when a backup requires media to overwrite. The search pattern is different according to whether you have chosen Full, Partial, or None. The media indicate that a type of media set is examined for availability.

See [“Selecting settings for media management”](#) on page 237.

See [“Media locations and vaults”](#) on page 250.

Selecting settings for media management

Use this procedure to select settings for the media overwrite protection level, the media overwrite options, and media labeling.

To select settings for media management

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Settings**, click **Media Management**.
 See “[Settings for media management](#)” on page 237.
- 3 Select the appropriate options, and then click **OK**.

Settings for media management

Settings for media management include the media overwrite protection level, the media overwrite options, and media labeling.

See “[Selecting settings for media management](#)” on page 237.

Table 4-6 Settings for media management

Item	Description
Full - protect allocated and imported media	<p>Prevents media in media sets and imported media from being overwritten.</p> <p>See “Media overwrite protection levels” on page 232.</p> <p>This is the safest option to choose because the media being protected cannot be overwritten until the following occurs:</p> <ul style="list-style-type: none"> ■ The overwrite protection period for the media expires. ■ You move media that belongs to an active media set to Scratch Media. ■ You erase, format, or label the media. ■ You move media from Imported Media to Scratch Media.
Partial - protect only allocated media	<p>Allows imported and scratch media to be overwritten. Media in a media set that has an overwrite protection period that has not expired (allocated media), cannot be overwritten.</p> <p>This option is recommended if you have media from an earlier version of Backup Exec or another product (imported media) that you want to reuse.</p>

Table 4-6 Settings for media management (*continued*)

Item	Description
<p>Prompt before overwriting imported media</p>	<p>Prompts you before Backup Exec overwrites imported media when Partial has been selected.</p> <p>The job will not run until you respond to this prompt.</p>
<p>None</p> <p>This option is not recommended because it does not protect data from being overwritten.</p>	<p>Disables the media overwrite protection feature. With this option, you are responsible for making sure that the media in your storage devices are not accidentally overwritten.</p> <p>For example, when an overwrite job is submitted to a device, and the media overwrite protection level is set to None, the media in that device is overwritten.</p>
<p>Prompt before overwriting allocated or imported media</p>	<p>Prompts you before Backup Exec overwrites allocated or imported media. If you selected None (no overwrite protection), it is highly recommended that you select this option to be prompted before overwriting allocated or imported media.</p> <p>The job will not run until you respond to this prompt.</p>
<p>Overwrite scratch media before overwriting recyclable media contained in the targeted media set</p>	<p>Lets Backup Exec overwrite scratch media first when an overwrite job occurs.</p> <p>See “How Backup Exec searches for overwritable media” on page 233.</p> <p>If no scratch media are found in any of the storage devices, Backup Exec overwrites recyclable media in the targeted media set.</p> <p>If no recyclable media are found in the targeted media set, Backup Exec searches for recyclable media in any media set.</p> <p>If no recyclable media are found, Backup Exec automatically searches for other media to overwrite. The media that is overwritten depends on the level of overwrite protection that you set (Full, Partial, or None). If you select this option, more media may be required for the same number of jobs than if you choose to overwrite recyclable media first.</p> <p>Because this option affects the order in which Backup Exec overwrites media, choosing to overwrite scratch media first may allow the recyclable media to be preserved longer for possible recovery.</p>

Table 4-6 Settings for media management (*continued*)

Item	Description
Overwrite recyclable media contained in the targeted media set before overwriting scratch media	<p>Lets Backup Exec overwrite recyclable media in the targeted media set first when an overwrite job occurs.</p> <p>If no recyclable media are found in any of the storage devices, Backup Exec overwrites scratch media.</p> <p>If no recyclable or scratch media are found, Backup Exec searches for media to overwrite. The media that is overwritten depends on the level of overwrite protection that you set (Full, Partial, or None).</p> <p>See “How Backup Exec searches for overwritable media” on page 233.</p> <p>If you choose to overwrite recyclable media in the targeted media set first, you will re-use the same media more frequently than if you choose to overwrite scratch media first.</p>
Media type	<p>Displays the types of media for which you can create default labels.</p> <p>See “Imported media labeling” on page 244.</p> <p>For example, if you select 4mm, then all 4mm-type media that are entered for the first time into this installation of Backup Exec are assigned a label according to what you specify in the following fields.</p>
Prefix	<p>Displays the current default prefix for the selected cartridge type. To specify a new prefix on the label, type from one to eight alphanumeric characters.</p>

Table 4-6 Settings for media management (*continued*)

Item	Description
Next value	<p>Displays the next number that will be included in the label of the next media that matches the selected cartridge type when that media is entered for the first time into this installation of Backup Exec. This number is incremented by one each time a media that matches the selected cartridge type is entered into this installation of Backup Exec.</p> <p>For example, if Cartridge type is set to 4mm, and Next value is set to 1, the first time a 4mm media is entered into this installation of Backup Exec, its label will include the number 1. The label on the next 4mm media entered will include the number 2.</p> <p>To enter a new value, type from one to eight numeric characters. This number must not exceed the number specified in the Digits field.</p>
Digits	<p>Displays the length of Next value, including placeholder zeroes. This field defines the minimum size of the numeric portion of the label.</p> <p>For example, if Next value is set to 1, and Digits is set to 6, then the Next value for the media label is 000001, 000002, 000003, and so on.</p> <p>If the Next value exceeds the entry in the Digits field, the extra digit is added. Using the previous example, if label numbering continued until 999,999, the next label would be 1,000,000 even though the value specified in Digits is 6.</p> <p>Rolling over the label numbering to 1,000,000 and 1,000,001 rather than 000,000 and 000,001 prevents the duplication of labels.</p> <p>The number entered in the Digits field must be in the range of three to eight.</p>
Default media vault	<p>Displays the default vault that you want media moved to when a job is run to move media to a vault or to export media. The default media vault that you select here is displayed on the job properties dialog box, in Options.</p> <p>See “Scheduling a job to move media” on page 255.</p> <p>See “Exporting expired media from a robotic library” on page 492.</p>

See “[How Backup Exec searches for overwritable media](#)” on page 233.

Viewing audit log entries for media operations

The audit log provides information about media operations, such as when media are overwritten or appended to. This information can help you find all of the media that is required for a restore job.

The following options for media operations are enabled by default in the audit log:

- Delete media
- Delete media set
- Erase media (long)
- Erase media (quick)
- Format media
- Format media (WORM)
- Label media
- Move media
- Overwrite media

To view the audit log entries for media operations

- 1 On the **Tools** menu, click **Audit Log**.
- 2 In the **Select category to view** field, click **Devices and Media**.
- 3 View the entries in the **Audit Log** window.

Configuring specific media operations to appear in the audit log

You can enable some or all media operations to appear in the audit log.

To configure specific media operations to appear in the audit log

- 1 On the **Tools** menu, click **Audit Log**.
- 2 On the **Audit Logs** dialog box, click **Configure Logging**.
- 3 Expand the **Devices and Media** category.

- 4 Select the operations that you want to log, or clear the check box of any item or operation that you do not want to log.
- 5 Click **OK**.

Media labeling

Media used in Backup Exec is identified by its media label. When new, blank, or unlabeled media is used during a backup operation, Backup Exec automatically labels the media. This label consists of a prefix that identifies the cartridge type, and an incrementing number. For example, if the media is a 4mm tape, then the prefix would be 4M, followed by 000001. The next media label generated for an unlabeled 4mm tape would be 4M000002, and so on.

You can allow the media label to be assigned automatically by Backup Exec, or you can specify a label prefix and number to be assigned for a type of media. For example, you can specify that all 4mm media that are entered for the first time into this installation of Backup Exec are labeled with a prefix of ACCT, and with numbering starting from 1000. You can specify another media type to be labeled with a prefix of FIN, and with numbering starting at 10,000. Customizing labels in this manner can help you recognize and organize media.

Another type of media label used by Backup Exec is the media ID, which is a unique label assigned by Backup Exec to individual media used in Backup Exec. The media ID is used internally by Backup Exec in order to keep statistics on each media. Because the media label or bar code label for media can be changed, Backup Exec must use the media ID, which cannot be changed or erased, to preserve continuity in record keeping for each individual media. The media ID has no effect on the media label, or on your ability to rename, label, or erase media.

At times, you may need to use the media ID to distinguish media that have duplicate media labels. Duplicate labels can be automatically generated in instances when Backup Exec is reinstalled or media from another Backup Exec installation is used. Use the media ID to distinguish between duplicate labels. You can view the media ID in a media's property page.

Write the media label on an external label fixed to the outside of the physical media. Whenever you change the media label, you should also change the external label to match.

The following methods are available in Backup Exec to change a media label:

- **Label Media operation.** Writes a new media label on the media. This write operation destroys any data on the media. This option is available on the **Devices** view.

- **Rename operation.** Changes the name of the media in the display, but does not write the new label to the media until an overwrite operation occurs. The data on the media is viable until the media is overwritten.
- **Edit the label in the media's property page.** Editing the label changes the name of the media in the display, but does not write the new label to the media until an overwrite operation occurs. The data on the media is viable until the media is overwritten.

See [“Media locations and vaults”](#) on page 250.

See [“Labeling media”](#) on page 486.

See [“Bar code labeling”](#) on page 244.

See [“Renaming a media label”](#) on page 243.

See [“General properties for media sets”](#) on page 230.

Renaming a media label

Use **Rename** to assign a new label to media. The new label is not actually written to the media until an overwrite operation occurs. All data on the media is preserved until the next overwrite job. However, the new media label is stored in the database and is displayed for that media. To write a new media label to the media immediately, use **Label Media** on the device's property page. The media's contents will be erased.

If you rename a media, and then use it in another installation of Backup Exec, that media is treated as imported media, and the media's original media label is displayed; the renamed label is not transferred to other installations of Backup Exec.

To rename media

- 1 Do one of the following:
 - If the media is in a device, then from the navigation bar, click **Devices**, and then click that device to display the media.
 - If you don't know where the media is, from the navigation bar, click **Media**, and then click **All Media** to display all media.
- 2 Select the media you want to rename.
- 3 Under **General Tasks** in the task pane, click **Rename**.
- 4 In **Name**, type a new media label, and then click **OK**.
- 5 Write this media label on an external label fixed to the outside of the physical media.

Imported media labeling

Backup Exec does not automatically relabel imported media. The imported media's existing label is read and displayed in the Media view, in one of the Imported Media sets. Additionally, the imported media's original media label is displayed under the heading **Media Description** in the Results pane of the **Media** view. You can edit the media description in the media's property page to make it a more descriptive label.

If the media overwrite protection level is set to Partial or None, the imported media may be selected for a job and be overwritten. The imported media is automatically labeled when it is overwritten during a job.

If you want to label a specific imported media while maintaining full media overwrite protection for other imported media, erase the specific media and then label it.

See [“General properties for media sets”](#) on page 230.

Bar code labeling

If there is a bar code label on the physical cartridge, and the cartridge is in a robotic library that has a bar code reader, the bar code label automatically becomes the media label.

You can change the media label in Backup Exec, but as long as the media has a bar code label that can be read, the bar code label takes precedence over the media label. To use the media label you entered using Backup Exec, you must remove the physical bar code label from the media cartridge, or use the media in a device without a bar code reader.

For example, Robotic Library 1 has bar code support. During a backup operation, Backup Exec requests a new or overwritable media for the operation. A new media with the bar code label 'ABCD' is inserted in the robotic library magazine and the bar code reader scans the bar code ID on the media label. Backup Exec selects this media for the operation and detects that a bar code label has been assigned to the media. Backup Exec automatically uses the bar code label and continues the operation.

When you change magazines or insert new media in a magazine, use the Scan option to quickly update slot information.

See [“Media labeling”](#) on page 242.

See [“Bar code rules in mixed media libraries”](#) on page 245.

Bar code rules in mixed media libraries

If you have bar code support for a robotic library that uses different types of drives, you can create a bar code rule so that Backup Exec can identify which media type to use in a drive. When Backup Exec reads the bar code rule, it locates the type of media that corresponds to the prefix or suffix and then mounts the media into a drive that accepts that type of media.

Bar code rules for robotic libraries do not go into effect until you enable them.

See [“Enabling bar code rules for robotic libraries”](#) on page 470.

See [“Creating bar code rules in mixed media libraries”](#) on page 245.

Creating bar code rules in mixed media libraries

If a robotic library has bar code support, then you can create bar code rules. Bar code rules specify the type of media that Backup Exec should use in a robotic library drive.

To create bar code rules in a mixed media library

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Settings**, click **Bar Code Rules**.
- 3 Click **New**.
- 4 Select the appropriate options.
See [“Add Bar Code Rule options”](#) on page 247.
- 5 Click **OK** to save the bar code rule for the media.
- 6 Verify that bar code rules are enabled for the robotic library. The bar code rules do not go into effect until you enable them for the robotic library.

See [“Enabling bar code rules for robotic libraries”](#) on page 470.

Editing a bar code rule

You can change the settings of a bar code rule.

To edit a bar code rule

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Settings**, click **Bar Code Rules**.
- 3 Click **Edit**, and then change the options as needed.

See [“Add Bar Code Rule options”](#) on page 247.

- 4 Click **OK** to save the changes, and then click **OK** to exit.

Deleting a bar code rule

You can delete a bar code rule.

To delete a bar code rule

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** Pane, under **Settings**, click **Bar Code Rules**.
- 3 Select a bar code rule to delete, and then click **Delete**.
- 4 Click **Yes** to verify that you want to delete the rule, and then click **OK**.

Application settings for bar code rules

Default bar code rules appear in the list. You can add, edit, or delete bar code rules.

Table 4-7 Application settings for bar code rules

Item	Description
Vendor	Displays the name of this library's manufacturer.
Prefix	Displays a prefix that is placed before the bar code. Only media with bar codes that have this prefix are used in the specified drive.
Suffix	Displays a suffix that is placed after the bar code. Only media with bar codes that have this suffix are used in the specified drive.
Media type	Displays the media type.
New	Lets you add a new bar code rule. See “Creating bar code rules in mixed media libraries” on page 245.
Edit	Lets you edit a bar code rule. See “Editing a bar code rule” on page 245.
Delete	Lets you delete a bar code rule. See “Deleting a bar code rule” on page 246.

See [“Bar code rules in mixed media libraries”](#) on page 245.

Add Bar Code Rule options

Bar code rules specify the type of media that Backup Exec should use in a robotic library drive.

See [“Creating bar code rules in mixed media libraries”](#) on page 245.

Table 4-8 Add Bar Code Rule options

Item	Description
Select a media type	Displays the types of media that you can select to include in the bar code rule.
Vendor	<p>Displays the name of this library’s manufacturer. You can find the name of the library manufacturer on the library’s property page. This field is not case-sensitive.</p> <p>By typing a vendor name here, you restrict the bar code rule to that vendor’s libraries. If you are creating a general bar code rule that applies to libraries from different vendors, leave this field blank.</p>
Bar code prefix	Displays a code that is placed before the bar code that represents a media type. The code can be up to 16 characters, and any combination of letters and numbers. This field is not case-sensitive.
Bar code suffix	Displays a code that is placed after the bar code that represents a media type. The code can be up to 16 characters, and any combination of letters and numbers. This field is not case-sensitive.

See [“Bar code rules in mixed media libraries”](#) on page 245.

About WORM media

Write once, read many (WORM) data storage is used for archiving data that requires a long retention period. Data can be written to WORM media one time only. After that, the media can be appended to, but it cannot be overwritten, erased, or reformatted.

When WORM media is used in a media set, the overwrite protection period is not applied to it, but the append period is applied.

New WORM media is WORM media that has not been written to. When new WORM media is introduced into Backup Exec, it is placed in the Scratch Media set.

After the WORM media has been written to one time, you cannot move it to the Scratch media set. You can move WORM media to the Retired Media set to delete it from Backup Exec, but you cannot erase it or reformat it.

You can use WORM media for ad hoc backup jobs and for backup jobs created from policies. When you select the option Use Write once, read many (WORM) media, Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent.

See [“Device and media options for backup jobs and templates”](#) on page 339.

Creating a new catalog

Catalog media to do the following:

- Log the contents of a media created by a product other than this installation of Backup Exec (imported media)
- Create a new catalog on the local hard drive if the catalog for the media no longer exists.

Before you can restore or verify data, the media must have a catalog. If the media is being used by this computer for the first time, you may need to inventory the media first.

See [“About inventoring media”](#) on page 448.

To create a new catalog

- 1 On the navigation bar, click **Media** or **Devices**.
- 2 Select the media set or location that is associated with the media to be cataloged, or select the device containing the media to be cataloged, and then select the media that you want to catalog.
- 3 On the task pane, under **Media Tasks**, click **Catalog media**.
- 4 To specify a device, or a password for a media that is password-protected, on the **Properties** pane, under **Target**, click **Device**, and then select the appropriate options.

See [“Device options for catalog jobs”](#) on page 249.

- 5 To specify a name for the job, or to specify a job priority, on the **Properties** pane, under **Settings**, click **General**, and then select the appropriate options.

See [“General options for utility jobs”](#) on page 482.

- 6 If you want a person or group to be notified when the job completes, on the **Properties** pane, under **Settings**, click **Notification**, and then select the options you want.

See [“Sending a notification when a job completes”](#) on page 682.

- 7 Do one of the following:

To run the job now Click **Run Now**.

To set the scheduling options you want to use On the **Properties** pane, under **Frequency**, click **Schedule**.
 See [“Scheduling jobs”](#) on page 356.

You can monitor or cancel the catalog job on the **Job Monitor**.

- 8 If the job requires that media be inserted into a robotic library, you are prompted to create an **Import Library** job.

See [“Importing media to a robotic library”](#) on page 489.

- 9 When you create the **Import Library** job, under **Import Job Properties**, click **Options**, and then select **Auto-inventory after import is completed**.

Device options for catalog jobs

Device options include the device on which to run a catalog job, and a password if necessary.

See [“Creating a new catalog”](#) on page 248.

Table 4-9 Device options for catalog jobs

Item	Description
Device	Displays the device on which this job will run.
Password	Displays the password to use if this media is password-protected and is being cataloged by this system for the first time.
Confirm password	Confirms that the same password is retyped.

Creating a restore job while reviewing media or devices

You can create a restore job while reviewing media or devices in Backup Exec. You must catalog the media before you can select the files that you want to restore.

The catalog for media that was backed up at other Backup Exec installations does not exist on the media server. You must catalog media that was backed up at other Backup Exec installations on the local media server before you can view data in the **Restore Job Properties** dialog box.

Creating a restore job while reviewing media or devices

- 1 On the navigation bar, click **Media** or **Devices**.
- 2 Double-click the media set or location that is associated with the data to be restored.
- 3 Select the media that you want to restore.
- 4 On the task pane, under **Media Tasks**, click **Restore data**.
- 5 Select restore job properties.
See [“Restoring data by setting job properties”](#) on page 605.
- 6 If the job requires that media be inserted into a robotic library, you are prompted to create an **Import Library** job.
See [“Importing media to a robotic library”](#) on page 489.
- 7 When you create the **Import Library** job, under **Import Job Properties**, click **Options**, and then select **Auto-inventory after import is completed**.

Media locations and vaults

Media in Backup Exec can be located in any of the following:

- **Online media.** This location lists media that reside in a storage device, robotic library slot, or backup-to-disk folder. Online media is defined by Backup Exec, so you cannot delete or rename it, and you cannot add or move media to it.

Note: If you move media from the online media location, its overwrite protection period and append period remain in effect.

- **Offline media.** This location displays all media that are onsite but are not in devices or slots, or media vaults. Media are automatically moved to this location if you use Backup Exec to remove media from a device or slot. You can add

media to the offline location from another media location. To move offline media back to online, run an inventory of the devices or slot, or catalog the media. You cannot delete or rename the offline location.

- **User-defined media vault.** A media vault is a logical representation of the actual physical location of media. You can create media vaults to keep track of where media is physically stored, such as a special media room, a scratch bin, or an offsite location. For example, you could create a media vault where media to be sent offsite are moved. Then, print the **Media Vault Contents** report, which lists the media contained in that vault, to accompany the physical media to their offsite storage. You can also create vault rules to help you track when media should be moved to or returned from a vault.

See [“Creating media sets”](#) on page 227.

See [“Finding media in a location or vault”](#) on page 254.

See [“Configuring vault rules for media sets”](#) on page 252.

See [“Using the Vault Wizard to move media”](#) on page 256.

See [“Renaming a media vault”](#) on page 253.

See [“About moving media to a vault or to the offline media location”](#) on page 254.

See [“Deleting a media vault”](#) on page 253.

See [“Media Vault Contents Report”](#) on page 751.

Creating media vaults

Create a media vault so that you can track media that are stored in specific sites. The new vault is displayed under **Media Location** in the **Media** view.

To create media vaults

- 1 On the navigation bar, click **Media**.
- 2 Under **Media Location Tasks** on the task pane, click **New media vault**.
- 3 Type the name and a description of the new vault.
See [“Media Vault Properties”](#) on page 251.
- 4 Click **OK**.

Media Vault Properties

Properties for media vaults include the name and a description of the media vault.

See [“Creating media vaults”](#) on page 251.

Table 4-10 Properties for media vaults

Item	Description
Name	Displays the name of the media vault.
Description	Displays a description of the media vault.

Configuring vault rules for media sets

On the vault rule properties for media sets, you can add or change the following:

- The media vault to which you want to send media that are associated with this media set.
- The amount of time to wait between when the media is allocated and when it is sent to the vault.
- The amount of time to wait between returning the media from the vault and when it was last written to.

Backup Exec does not update the vault automatically. You must use the **Update vault using wizard** task to print or view reports that contain details on which media are ready to be moved to and from the vault, and to update the media location. You can also schedule a job called **Move Media to a Vault** to export media from a device, and update the media location for any media that are moved to a vault.

To configure vault rules for media sets

- 1 On the navigation bar, click **Media**.
- 2 In the **Media** selection pane, under **Media Sets**, select a media set.
- 3 Under **General Tasks** on the task pane, click **Properties**.
- 4 Click the **Vault Rules** tab.
See [“Properties for vault rules for media sets”](#) on page 252.
- 5 Select the appropriate options, and then click **OK**.

Properties for vault rules for media sets

Properties for vault rules provide information on the dates on which media is moved to or returned from a media vault.

See [“Configuring vault rules for media sets”](#) on page 252.

See [“Creating media sets”](#) on page 227.

Table 4-11 Properties for vault rules for media sets

Item	Description
Select the media vault to use with this media set:	<p>Displays the media vault that stores the media that is associated with this media set.</p> <p>Before the media location can be updated, even if the move and return dates are overdue, you must run the task Update vault using wizard.</p> <p>See “Using the Vault Wizard to move media” on page 256.</p> <p>This wizard can print reports that detail which media are ready to move to and return from the vault, and can update the location of the media if you choose to move them. However, you must physically collect the media, and move the media to and from the vault.</p>
Move media to this vault x after it is allocated (first written to or overwritten)	<p>Displays the time period after which this media will be reported as ready to be moved to this vault.</p>
Return media from this vault x after it is last written to	<p>Displays the time period after which this media will be reported as ready to be returned from this vault.</p>

See [“About creating media sets”](#) on page 226.

Deleting a media vault

You can delete an empty media vault. If there is any media in the vault, you must move it before you can delete the vault. You cannot delete the online or offline locations.

To delete a media vault

- 1 On the navigation bar, click **Media**.
- 2 Select the media vault that you want to delete.
- 3 Under **General Tasks** in the task pane, click **Delete**.
- 4 When prompted if you are sure that you want to delete the media vault, click **OK**.

Renaming a media vault

You can rename a media vault. You cannot rename the **Backup Exec Media Location** defaults **Online Media** and **Offline Media**.

To rename a media vault

- 1 On the navigation bar, click **Media**.
- 2 On the **Media** selections pane, click the media vault that you want to rename.
- 3 Under **General Tasks** in the task pane, click **Rename**.
- 4 Type the new name, and then click **OK**.

Finding media in a location or vault

You can find where media is located by searching for the name on the media label.

To find media in a location or vault

- 1 On the navigation bar, click **Media**.
- 2 On the **Media** selection pane, click **Media Location**.
- 3 Under **Media Location Tasks** on the task pane, click **Find media**.
- 4 Type the name from the media label of the media you want to find, and then click **OK**.

About moving media to a vault or to the offline media location

Several operations are available for you to logically move media to a vault or to the offline media location. Some move operations also prompt you to export the media as part of the operation. While these operations logically move the media, you must physically move the media to an actual location that is represented by the vault name.

Move media to a vault or to the offline media location using any of the following methods:

- Scan the bar code label or type the media label to logically move media to a vault or to the offline media location.
See [“Scanning bar code labels to move media”](#) on page 255.
- Schedule a job to logically move media to a vault after it is successfully exported from a device.
See [“Scheduling a job to move media”](#) on page 255.
- In a policy, create a job template to export media.
See [“Adding an export media template to a policy”](#) on page 537.
- Run the **Vault Wizard** to export media from a device and to logically move media to a vault.

See [“Using the Vault Wizard to move media”](#) on page 256.

- Drag and drop media to a vault or to the offline media location, and then export the media from the device.
See [“Drag and drop methods to move media”](#) on page 258.
- Use the **Move to vault** option to logically move media to a vault or to the offline media location.
See [“Using the Move to vault task to move media”](#) on page 257.

Scanning bar code labels to move media

If you have a bar code scanner, this is an efficient method for moving media to a vault or to the offline media location. You can also type a media label into the dialog box.

To scan bar code labels to move media

- 1 On the navigation bar, click **Media**.
- 2 In the **Media** selections pane, select the media location or vault to which you want to move media.
- 3 In the task pane, under **Media Location Tasks** in the task pane, click **Add media to selected vault**.
See [“Move Media to Vault options”](#) on page 258.
- 4 Enter each label on a separate line.
- 5 Click **OK**.
- 6 Repeat steps 4 and 5 for all the media you want to add.
- 7 Click **OK**.

Scheduling a job to move media

If you set up vault rules for a media set, you can schedule a job called **Move Media to Vault**. This job exports the media from the device, and then logically moves the media to the specified vault.

You can specify a vault for a media set in the media set vault rules, or you can specify a default vault.

There must be an existing user-defined media vault under the **Media Location** node in the **Media** view.

See [“Creating media vaults”](#) on page 251.

To schedule a job to move media

- 1 On the navigation bar, click **Job Setup**.
- 2 On the task pane, under **Backup Strategy Tasks**, click **New job to move media to a vault**.
- 3 To specify a name for the job, or to specify a job priority, on the **Properties** pane, under **Settings**, click **General**, and then select the appropriate options.
See [“General options for utility jobs”](#) on page 482.
- 4 To move the media to a media vault after a successful export, on the **Properties** pane, under **Settings**, click **Options**, and select a media vault.
- 5 If you want a person or group to be notified when the job completes, on the **Properties** pane, under **Settings**, click **Notification**, and then select the options you want.
See [“Assigning recipients to alert categories for notification”](#) on page 680.
- 6 If you want to run the job now, click **Run Now**. Otherwise, on the **Properties** pane, under **Frequency**, click **Schedule** to set the scheduling options you want to use.
See [“Scheduling jobs”](#) on page 356.

Using the Vault Wizard to move media

Use the **Vault Wizard** to print or view reports that contain details on which media are ready to be moved to and from a vault, and to update vaults.

After a vault is updated, if Backup Exec detects a robotic library, you are prompted to export media. If you choose to export the media, an **Export Media** job runs.

See [“Exporting expired media from a robotic library”](#) on page 492.

Note: If your environment includes remote sites, you should create separate media sets for each remote site, so that if vault rules are enabled, the reports contain details on which media are ready to be moved for just that site.

To use the Vault Wizard to move media to a vault

- 1 On the navigation bar, click **Media**.
- 2 Select the vault that you want to update.

- 3 Under **Media Location Tasks** on the task pane, click **Update vault using wizard**.
 The **Vault Wizard** appears.
- 4 Follow the instructions on the wizard.

Using the Move to vault task to move media

You can select media in the **Media** view, and then choose a vault or the offline media location to move the media to. The media location is updated in the Backup Exec database, but the media is not ejected or exported.

To use the Move to vault task to move media

- 1 On the navigation bar, click **Media**.
- 2 Expand **All Media** to display a list of media, and then select the media that you want to move to a media vault.
- 3 Under **Media Tasks** on the task pane, click **Move to vault**.
 See [“Move Media options”](#) on page 257.
- 4 Select a media vault to which you want to move this media, and then click **Yes** or **Yes to All**.

Move Media options

You can move media to a vault or to the offline media location.

See [“Using the Move to vault task to move media”](#) on page 257.

Table 4-12 Move Media options

Item	Description
Name	Displays the media label that you selected to move.
Description	Displays the media description, if there is one.
Move to	Displays a media vault or a media location to which the media is moved.
Yes	Moves a single media to the new location and updates the Backup Exec database.
No	Unselects the media and does not move it.

Table 4-12 Move Media options (*continued*)

Item	Description
Yes to All	Moves all media that was selected to the new location and updates the Backup Exec database.

Move Media to Vault options

You can use a bar code scanner to enter the media labels of the media that you want to move to this vault. You can also type the media label for any media that you want to move to this vault.

See [“About moving media to a vault or to the offline media location”](#) on page 254.

Drag and drop methods to move media

To move media to a vault or to the offline media location, drag and drop it from one location to another location. The drag and drop method also prompts you to export the media from a device.

The following table lists the drag-and-drop rules for media:

Note: If you move media from an online location, its overwrite protection period and append period remain in effect.

Table 4-13 Media Drag-and-Drop Rules

From/To	Online Location	Offline Location	Media Vaults	Media Pools/Sets	All Media
Online location	No	Yes, with a warning that the media is not physically moved from the online location.	Yes, with a warning that the media is not physically moved from the online location.	No	No
Offline location	No	No	Yes	No	No
User-defined vault	No	Yes	Yes	No	No
Media Pools/Sets	No	Yes, with a warning	Yes, with a warning	Yes	No

Table 4-13 Media Drag-and-Drop Rules (*continued*)

From/To	Online Location	Offline Location	Media Vaults	Media Pools/Sets	All Media
All Media	No	Yes, with a warning that the media is not physically moved from the online location.	Yes, with a warning that the media is not physically moved from the online location.	Yes	No

Using drag and drop methods to move media

After you drag and drop media to a new location, an **Export Media** job runs. After the job runs, you are prompted to remove the media from the device.

See [“Drag and drop methods to move media”](#) on page 258.

To use the drag and drop method to move media

- 1 On the navigation bar, click **Media**.
- 2 On the **Media** selections pane, click the vault containing the media.
- 3 Select the media from the Results pane, and drag it to the new location.
- 4 When you are prompted to export the media, click **Yes** or **Yes to All**.

About removing damaged media

Media that meets or exceeds the discard thresholds determined by the media manufacturer should be associated with the **Retired Media** media set. Based on a measurement of soft errors generated by the storage device firmware, media that exceeds acceptable levels of these errors are reported to Backup Exec as potential candidates to be discarded.

To decide which media to retire, run a **Media Sets** report to see the total number of errors for media, or view the properties for a specific media.

Associate any media with an unacceptable level of errors to **Retired Media** so that you are protected against using defective media before critical backup operations begin. After you mark media as retired, it will not be used by Backup Exec for future backup jobs. The media is still available to be restored from if it is not damaged.

See [“About deleting media”](#) on page 260.

See [“Statistics properties for media”](#) on page 263.

See [“Failed Backup Jobs Report”](#) on page 744.

About deleting media

When you delete media from Backup Exec, all records of the media are removed from the Backup Exec Database. These records include catalog information, media statistics, and other information that is associated with the media. You can only delete media when it belongs to the **Retired Media** set.

You may want to delete media when the following occurs:

- You have a lot of offsite media that you do not want to recycle.
- You throw away damaged or old media.

Media can only be deleted from Backup Exec after it has been associated with the **Retired Media** set.

When deleted media is reused in Backup Exec, it is recognized as imported media. Before you can restore from the media, you must catalog it.

Note: Deleting media from Backup Exec is not the same operation as erasing media.

See [“Deleting media”](#) on page 260.

See [“Statistics properties for media”](#) on page 263.

See [“Failed Backup Jobs Report”](#) on page 744.

Deleting media

You can delete media from the Backup Exec Database.

To delete media

- 1 On the navigation bar, click **Media**.
- 2 Associate the media that you want to delete with the **Retired Media** set by doing one of the following:
 - Drag the media to the **Retired Media** icon.
 - Under **Media Tasks** on the task pane, click **Associate with media set**, select the **Retired Media** set to associate the media with, and then click **Yes** or **Yes to all**.
- 3 Double-click the **Retired Media** icon, and then select the media that you want to delete.

- 4 Under **General Tasks** in the task pane, click **Delete**.
 If **Delete** is unavailable, the media is not associated with the **Retired Media** set. You must associate the media with **Retired Media** before **Delete** is available.
- 5 Click **Yes** or **Yes to All** to delete the media that are displayed.

General properties for media

General properties for media provides information about the media.

See [“Viewing properties”](#) on page 217.

Table 4-14 General properties for media

Item	Description
Media label	<p>Displays the media label that was assigned automatically by Backup Exec, or that was assigned or changed by the administrator, or that was a pre-assigned bar code label.</p> <p>You can edit the media label, which is limited to 32 characters. Editing the label changes the name of the media in the display, but does not write the new label to the media until an overwrite operation occurs. When you edit a media label, try to make it a concise identifier that will remain constant even when the media is reused. You should write this media label on a label fixed to the outside of the physical media.</p> <p>Duplicate labels can be automatically generated. For example, reinstalling Backup Exec or bringing media from another Backup Exec installation could cause duplication in labels. Duplicate labels are allowed, but not recommended.</p> <p>If a bar code is available, and a bar code-equipped device is used, then the media label automatically defaults to that bar code.</p>
Description	<p>Displays the original media label if the media is imported media.</p> <p>You can edit the media description, which is limited to 128 characters, to make it a more descriptive label.</p>

Table 4-14 General properties for media (*continued*)

Item	Description
Media type	Displays the media type and subtype (if available). Click the button next to the field to change the media type or subtype.
Export pending	Displays Yes when a job runs that has an associated Export Media template to export this media. See “About export media templates” on page 536.
Media set	Displays the name of the media set this media belongs to.
Media location	Displays the name of the device or vault where this media is located.
Creation date	Displays the date and time when the media was first entered into Backup Exec.
Allocated date	Displays the date and time when the media was added to a media set as a result of an overwrite operation.
Modified date	Displays the date and time when data was last written to the media.
Overwrite protection until	Displays the date and time after which the media can be overwritten.
Appendable until	Displays the date and time after which the media can no longer be appended to.
Supports HW encryption	Displays Yes if this media supports hardware encryption. See “About hardware encryption” on page 416.

See [“Media labeling”](#) on page 242.

See [“Creating a test run job”](#) on page 386.

See [“Properties for robotic library slots”](#) on page 472.

See [“Statistics properties for media”](#) on page 263.

Statistics properties for media

You can view statistics about a media.

See “[Viewing properties](#)” on page 217.

Table 4-15 Statistics properties for media

Item	Description
Hours in use	Displays the total number of hours that this media has been in use.
Used capacity	Displays the amount of raw capacity on the media that has been used. Used capacity is calculated by subtracting available capacity from total capacity . Used capacity may or may not equal bytes written .
Available capacity	Displays the amount of expected raw capacity on the media that remains unused. Some tape devices support the ability to read the amount of remaining capacity of the media that is currently loaded in the device. If a tape device supports reading of the remaining capacity amount, then available capacity is derived from the remaining capacity amount. Otherwise, available capacity is calculated by subtracting bytes written from total capacity . Because free space is reported in terms of unused raw capacity, review bytes written and compression ratio to better estimate if there is enough free space to accommodate a specific job.
Total capacity	Displays the amount of expected total raw capacity of the media. Some tape devices support the ability to read the amount of total capacity of the media that is currently loaded in the device. If a tape device supports reading of the total capacity amount, then total capacity is derived from the total capacity amount. Otherwise, total capacity is estimated based on past usage of the media.
Compression ratio	Displays the ratio of bytes written to used capacity . Compression ratio will show the overall effect that data compression and media flaws are having on the amount of data that is being stored on the media.

Table 4-15 Statistics properties for media (*continued*)

Item	Description
Bytes written	Displays the amount of data that has been written into blocks on the media. Bytes written may differ from used capacity due to the effects of data compression and media flaws. Data compression will tend to increase bytes written when compared to used capacity . Media flaws will decrease bytes written when compared to used capacity .
Bytes read	Displays the number of bytes that have been read from this media.
Mounts	Displays the number of times this media has been mounted.
Seeks	Displays the total number of seek operations that have been performed on this media. Seek operations run to locate a specific piece of information on the media.
Seek errors	Displays the number of errors encountered while trying to locate data.
Soft write errors	Displays the number of recoverable write errors encountered. If you receive soft errors, it may indicate the beginning of a problem. If you receive excessive errors for your environment, check the media for damage.
Hard write errors	Displays the number of unrecoverable write errors encountered. If you receive hard errors, check the media for damage.
Soft read errors	Displays the number of recoverable read errors encountered. If you receive soft errors, it may indicate the beginning of a problem. If you receive excessive errors for your environment, check the media for damage.
Hard read errors	Displays the number of unrecoverable read errors encountered. If you receive hard errors, check the media for damage.

See [“Editing general properties for media sets”](#) on page 230.

See [“About creating media sets”](#) on page 226.

Media rotation strategies

There are many media rotation strategies you can use to protect your data. They differ mostly by the number of media required and by the amount of time the media is kept before it is rotated back into the schedule.

The most commonly used media rotation strategies include the following:

- Son, which uses the same media each day to run a full backup.
 See “[Son media rotation strategy](#)” on page 265.
- Father/Son, which uses multiple media, includes a combination of weekly full and daily differential or incremental backups for a two-week schedule, and provides backups for offsite storage.
 See “[Father/son media rotation strategy](#)” on page 266.
- Grandfather, which uses multiple media, includes a combination of weekly and monthly full and daily differential or incremental backups, and provides backups for offsite storage.
 See “[Grandfather media rotation strategy](#)” on page 267.

Son media rotation strategy


The Son media rotation strategy requires the following:

Table 4-16 Son media rotation strategy

Item	Description
Number of media required	1 (minimum)
Overwrite protection period	Last backup

The Son strategy involves performing a full backup every day.

Figure 4-4 Son Backup Strategy



Although the Son strategy is simple to administer, backing up with a single media is not an effective method of backup. This is because magnetic media eventually wears out after many uses and the data you can restore only spans back to your last backup.

Father/son media rotation strategy

The Father/son media rotation strategy requires the following:

Table 4-17 Father/son media rotation strategy


Item	Description
Number of media required	6 (minimum)
Overwrite protection period	Two weeks

The Father/Son media rotation strategy uses a combination of full and differential or incremental backups for a two-week schedule.

In the Father/Son scenario, four media are used Monday through Thursday for differential or incremental backups. The other two media containing full backups are rotated out and stored offsite every Friday.

The Father/Son strategy is easy to administer and allows you to keep data longer than the Son strategy, but it is not suitable for the stringent data protection needs of most network environments.

Figure 4-5 Father/Son Backup Strategy



When this backup strategy is first implemented, you must start with a full backup.

Grandfather media rotation strategy

The Grandfather media rotation strategy requires the following:

Table 4-18 Grandfather media rotation strategy


Item	Description
Number of media required	19 (minimum)
Overwrite protection period	One year

The Grandfather method is one of the most common media rotation strategies. It is simple to administer and comprehensive enough to allow easy location of files when they need to be restored.

In the Grandfather scenario, four media are used Monday through Thursday for incremental or differential backups; another three media are used every Friday for full backups.

The remaining 12 media are used for monthly full backups and are kept offsite.

Figure 4-6 Grandfather Backup Strategy



The Grandfather strategy is recommended because it offers a good media number to storage life ratio (19 media/1 year). It is also easy to modify if you want to incorporate more media. For example, you could perform a full backup on the last Saturday of the month to keep permanently.

Preparing for backup

This chapter includes the following topics:

- [How to prepare for backup](#)
- [About backup strategies](#)
- [How to choose a backup strategy](#)
- [About selecting data to back up](#)
- [About using fully qualified computer names in backup selections](#)
- [About the Computer Name node in the backup selections list](#)
- [About the Favorite Resources node in the backup selections list](#)
- [About the Domains node in the backup selections list](#)
- [Adding an Active Directory domain to the Active Directory Domains node](#)
- [Deleting an Active Directory domain from the Active Directory Domains node](#)
- [About the User-defined Selections node in the backup selections list](#)
- [Adding a user-defined selection to the User-defined Selections node](#)
- [Deleting a user-defined selection from the User-defined Selections node](#)
- [About managing Microsoft Virtual Hard Disk \(VHD\) files in Backup Exec](#)
- [How to back up user-defined Microsoft Windows Distributed File System data](#)
- [About selection lists](#)
- [About resource discovery](#)
- [About the Backup Exec Shadow Copy Components file system](#)

- [How to restore individual items by using Granular Recovery Technology](#)

How to prepare for backup

Before you back up your data, you should become familiar with how to do the following:

- Use backup strategies
See [“About backup strategies”](#) on page 270.
- Select data to back up
See [“About selecting data to back up ”](#) on page 280.
- Create user-defined selection lists
See [“About selection lists”](#) on page 295.
- Protect specific types of data, such as Windows Server systems and utility partitions
See [“About selecting data to back up ”](#) on page 280.

About backup strategies

A backup strategy is the collection of procedures you implement for backing up your network, including what methods of backups are performed, when backups are performed, and how media is rotated back into use for your regular backups. A good backup strategy allows minimal time to get a system up and running in the event of a disaster.

Backup Exec offers flexible solutions for protecting the data on your network. Use the media rotation feature and let Backup Exec do all the administrative work for you, or design and manage your own media rotation strategy, the procedures for reusing media, that meets your exact specifications.

You can create backup job templates that define your backup strategy, and then reuse the templates to implement your strategy for all resources being protected by your media server.

See [“Media rotation strategies”](#) on page 265.

See [“How to choose a backup strategy”](#) on page 270.

How to choose a backup strategy

In order to develop a secure and effective plan for managing your data, you should consider the following:

- The importance of the data you are backing up.
- How often your system needs to be backed up.
- How much storage media you will use.
- When you will use certain storage media.
- How you will keep track of your backup information.

See [“About backup strategies”](#) on page 270.

How to determine your backup schedule

While there is no requirement on how often to back up your data, there is one consideration that can help you decide: The cost of re-creating data that was added or modified since the last backup.

Calculate the manpower, lost time and/or sales, and other costs that would be incurred if the file server or workstation crashed right before the next backup was scheduled to take place (always assume the worst scenario). If the cost is excessive, the strategy should be adjusted accordingly.

For example, the cost to re-create an extensive database system that is continually updated by several database operators would be quite substantial. On the other hand, the cost to re-create the data for a user creating one or two inter-office memos would be considerably less. In this scenario, the network administrator would probably opt to back up the database several times daily, and set up daily jobs for the user’s workstation.

In an ideal environment, one full backup should be performed on workstations every day and servers should be fully backed up more often. Important data files and directories that constantly change may need to be backed up several times a day. Because of time and media constraints, this is not feasible for many environments, so a schedule including incremental or differential backups must be implemented. For safety reasons, a full backup should always be performed before adding new applications or changing the server configuration.

How to determine the amount of data to back up

The amount of data to be backed up is a key determinant of the media rotation strategy you choose. If you are backing up large amounts of data that needs to be retained on media for long periods of time, you will need to select a strategy that is suitable for these requirements.

How to determine a schedule for data storage

The amount of time the data needs to be stored is directly related to the media rotation scheme you use. For example, if you use one media and back up every day, your backups will never be more than a day old.

Since storage media is relatively inexpensive when compared to the value of your data, it is a good idea to periodically back up your system on media not used in the media rotation scheme and store it permanently. Some administrators may choose to do this every week, while others may choose to store only one permanent backup per month.

The threat of viruses is an issue also. Some viruses take effect immediately, while others may take days or weeks to cause noticeable damage.

You should have at least the following backups available to restore at any time:

- 3 daily backups (for example, Monday, Tuesday, Wednesday).
- A one-week-old full backup.
- A one-month-old full backup.

Having these backups available should allow you to restore your system prior to when it became infected.

How to determine which devices to back up

Since Backup Exec can back up servers, workstations, and agents, you should consider which resources you want to protect. You will need to coordinate times that are suitable to back up different resources. For example, you may want to back up file servers during the evening and back up workstations at lunchtime.

How to determine the number of resources to back up in a job

When you are setting up jobs for the network, you must decide if you want to create one job that includes many resources or a job for each resource.

Here are some of the advantages and disadvantages of each method.

Table 5-1 Advantages and disadvantages of including more than one system in a job

Method	Advantages	Disadvantages
One job per device	<ul style="list-style-type: none"> ■ If a job fails, you know immediately which resource was not backed up. ■ If a resource is turned off or moved, backups of other resources are not affected. ■ When resources are added to the network you can simply set up new jobs for each resource. 	<ul style="list-style-type: none"> ■ You have more jobs to keep track of (for example, reviewing job histories, and so forth).
Multiple resources per job	<ul style="list-style-type: none"> ■ There are fewer jobs to keep track of and create. ■ You know the order in which the data is backed up. ■ You could make it an overwrite job and thus be able to use the same name for the media and the job. 	<ul style="list-style-type: none"> ■ If any of the resources in the job are not available during the backup, the job results in an abnormal completion status.

About the archive bit and backup methods

Whenever a file is created or changed, the operating system activates the Archive Bit or modified bit. Unless you select to use backup methods that depend on a date and time stamp, Backup Exec uses the archive bit to determine whether a file has been backed up, which is an important element of your backup strategy.

Selecting the following backup methods can affect the archive bit:

- Full - Back up files - Using archive bit (reset archive bit)
- Differential - Back up changed files since last full - Using archive bit (does not reset archive bit)
- Incremental - Back up changed files since last full or incremental - Using archive bit (reset archive bit)

Whenever a file has been backed up using either the Full - Back up files - Using archive bit (reset archive bit) or Incremental - Changed Files - Reset Archive Bit backup method, Backup Exec turns the archive bit off, indicating to the system that the file has been backed up. If the file is changed again prior to the next full or incremental backup, the bit is turned on again, and Backup Exec will back up the file in the next full or incremental backup. Backups using the Differential - Changed Files backup method include only files that were created or modified

since the last full backup. When this type of differential backup is performed, the archive bit is left intact.

Consider the following backup strategy scenario:

Fred wants to implement a backup strategy for the office fileserver. Fred knows that all backup strategies begin with a full backup (backup of an entire device using the full backup method), so he creates a Selection List for his server and submits the job to run at the end of the day on Friday.

Since most files on the server, such as operating system files and application files, seldom change, Fred decides that he can save time and media by incorporating incremental or differential backups in his media rotation scheme. Fred opts to use incremental backups, so he schedules the script to run at the end of the day, Monday through Thursday, with the incremental backup method.

Here's what happens: Fred's Friday tape contains all of the data on the fileserver and Backup Exec changes all of the files' statuses to backed up. At the end of the day on Monday, the incremental job runs and only the files that were created or changed (had the archive bit re-set by the operating system) are backed up. When the incremental job completes, Backup Exec will turn the archive bit off, showing that the files have been backed up. On Tuesday through Thursday, the same events happen.

If Fred's fileserver crashed on Thursday morning, after he got it running, he would restore each backup in the order in which it was created (for example, Friday, Monday, Tuesday, and so forth).

If Fred had decided to perform differential backups on Monday through Thursday, he would have only needed Friday's and Wednesday's tapes: Friday's tape because it included all of the data, and Wednesday's tape because it included every file that had been created or changed since Friday's backup.

About backup methods

Before you can develop your media rotation strategy, you will need to decide whether you want to perform only full backups or use a strategy that includes Full backups and one of the modified backup methods (differential, incremental or working set backups).

Note: You need to perform a full backup of your server to establish a baseline for disaster recovery.

The backup methods used by Backup Exec are as follows:

- Full

See [“About the full backup method”](#) on page 275.

- Differential

See [“About the differential backup method”](#) on page 276.

- Incremental

See [“About the incremental backup method”](#) on page 276.

- Working Set

See [“About the working set backup method”](#) on page 277.

There are advantages and disadvantages to each backup method.

See [“About backup method advantages and disadvantages”](#) on page 277.

About the full backup method

Full backups include all of the data that was selected for backup. Backup Exec detects the device as having been backed up. You can use either archive bit or modified time to determine if a file has been backed up.

Note: You will need to perform a full backup of your server to establish a baseline for disaster recovery.

Full backups also include Copy backups, which include all selected data and do not affect any media rotation scheme because the archive bit is not reset. Copy backups are useful when you need to:

- Back up data for a special purpose (for example, to send to another site).
- Back up specific data.
- Perform an additional backup to take off-site.
- Back up data that belongs to a media rotation job without affecting the rotation cycle.

Another Full backup option is the **Back up and delete the files** option. This option deletes the selected files and folders from the volume after a successful copy backup. This backup option moves data from disk to storage media to free valuable disk space and to reduce clutter on your server volume. You should not use this option as part of a regular backup schedule.

Note: For data to be deleted, rights to delete a file must be granted; otherwise data is backed up but not deleted. Backup Exec does not delete data from remote computers on which remote agents are installed when you select the **Back up and delete the files** option.

See [“About backup methods”](#) on page 274.

About the differential backup method

Differential backups include all files that have changed since the last full or incremental backup. The difference between differential and incremental backups is that incremental backups include only the files that have changed since the last full or incremental backup.

Backup Exec provides two differential backup methods, one that uses archive bit and one that uses modified time to determine if the file was backed up. If you select to use the Full - Back up files - Using archive bit (reset archive bit) and want to run differential backups, you must use the Differential - Back up changed files since last full - Using archive bit (does not reset archive bit) method. If you select to use the Full - Back up files - Using modified time, you must use the Differential - Back up changed files since last full - Using modified time method.

Note: If you use modified time to determine if files have been backed up, the full and differential backups must use the same backup selection list.

In most schemes, differential backups are recommended over incremental backups. Differential backups allow much easier restoration of an entire device than incremental backups since only two backups are required. Fewer required media also decreases the risk of not being able to restore important data because of media errors.

You should not mix differential and incremental backups together.

See [“About backup methods”](#) on page 274.

About the incremental backup method

Incremental backups include only the files that have changed since the last full or incremental backup.

Backup Exec provides two incremental backup methods, one that uses archive bit and one that uses modified time to determine if the file was backed up. If you select to use the Full - Back up files - Using archive bit (reset archive bit) and want to run incremental backups, you must use the Incremental - Back up changed files since last full or incremental- Using archive bit (reset archive bit) method. If you select to use the Full - Back up files - Using modified time, you must use the Incremental - Back up changed files since last full or incremental - Using modified time method. The advantages and disadvantages described in this section pertain to either type of differential backup.

Note: If you use modified time to determine if files have been backed up, the full and incremental backups must use the same backup selection list.

See [“About backup methods”](#) on page 274.

About the working set backup method

The working set backup method includes two options; Changed today and Last accessed in x days. The Changed today option was called Daily backup in previous versions of Backup Exec. You can perform backups using the Changed today method in addition to the media rotation scheme selected. The Changed today method backs up all files with today’s date (created or changed today). The Changed today method does not affect the files’ backup status because the archive bit is not reset.

If you select the Last accessed in x days method, you can then indicate in the Files accessed in x days field that you want to include data that has been accessed in a specific number of days.

This option is similar to a differential backup, in which files that have been created or changed since the last full backup are included; however, the difference lies in that you can also specify to include all files accessed within the last x number of days. This option can speed the recovery of a crashed server because you only need to restore the working set backup to get up and running again, and then restore the latest full backup at a later time (if necessary).

To effectively include the data needed to make your system operational after restoring a working set backup, specifying at least 30 days in the Files accessed in x days field is recommended. With a full/working set backup scheme, the non-full backups will require more media space than full/differential or full/incremental schemes. However, in environments where active data is migrated frequently between machines, or when restore times are especially critical, working set backups can make up the cost of extra media in time savings for restoring data.

See [“About backup methods”](#) on page 274.

About backup method advantages and disadvantages

There are advantages and disadvantages to each backup method.

See [“About backup methods”](#) on page 274.

Table 5-2 Backup Method Advantages and Disadvantages

Method	Advantages	Disadvantages
Full	<ul style="list-style-type: none"> ■ Files are easy to find - Since full backups include all data contained on a device, you don't have to search through several media to find a file that you need to restore. ■ There is always a current backup of your entire system on one media or media set - If you should need to restore your entire system, all of the most current information is located on the last full backup. 	<ul style="list-style-type: none"> ■ Redundant backups - since most of the files on your file server rarely change, each full backup following the first is merely a copy of what has already been backed up. This requires more media. ■ Full backups take longer to perform - Full backups can be time consuming, especially when you have other devices on the network that need to be backed up (e.g., agent workstations, remote servers).
Differential	<ul style="list-style-type: none"> ■ Files are easy to find - Restoring a system backed up with a differential strategy requires a maximum of two backups - the latest full backup and the latest differential backup. This is less time consuming than backup strategies that require the latest full backup and all incremental backups created since the full backup. ■ Less time required for backup and restore - Differential backups take less time to restore than full backups. Faster recovery is possible in disaster situations because you only need the latest full and differential backup media to fully restore a device. 	<ul style="list-style-type: none"> ■ Redundant backups - All of the files created or modified since the last incremental backup are included; thus creating redundant backups.

Table 5-2 Backup Method Advantages and Disadvantages (*continued*)

Method	Advantages	Disadvantages
Incremental	<ul style="list-style-type: none"> ■ Better use of media - Only files that have changed since the last backup are included, so there is much less data storage space required. ■ Less time required for backup - Incremental backups take much less time than full and differential backups to complete. 	<ul style="list-style-type: none"> ■ Backups are spread across multiple media - Since multiple media is required in a disaster situation, this can cause recovery of a device to take longer. In addition, the media must be restored in the correct order to effectively bring the system up to date.
Working Set	<ul style="list-style-type: none"> ■ Restoring a system backed up with a working set strategy requires only the media containing the latest working set backup media and the media containing the most recent full backup. ■ You can perform a working set backup, restore the data to a new system, and be up and running faster than if you had to restore a full backup followed by all of the incremental or differential backups. ■ Working set backups take less time to run than full backups. 	<ul style="list-style-type: none"> ■ The Last accessed in (x) days method is available only on platforms that support the last accessed date (Windows, NetWare, and UNIX). Working set backups will work as differential backups when selected for other platforms.

About modified time and backup methods

If you select to use Full - Back Up Files - Using modified time, Differential - Using modified time, or Incremental - Using modified time, Backup Exec uses a file's modified time rather than the archive bit to determine if it needs to be backed up.

When Backup Exec runs a full or incremental backup, the time the backup launches is recorded in the Backup Exec database. The next time an incremental or differential backup launches, Backup Exec compares the file system time to the backup time recorded in the Backup Exec database. If the file system time is later than the database time, the file is backed up.

Note: A file's last modified date and timestamp does not change when the file is copied or moved. If the file's modified time is older than the previous backup's modified time, that file is not backed up. To ensure that the files are protected, run a full backup after you copy or you move files. If you have the Advanced Disk-based Option, you can run synthetic backups to ensure that any copied or moved files are protected.

When an incremental backup is run, a new time is recorded in the Backup Exec database. The database time is not updated with differential backups.

Using the modified time allows Backup Exec to run differential backups on file systems, such as UNIX, which do not have an archive bit.

If you want Backup Exec to use modified time to determine if a file has been backed up, the full and incremental (or full and differential) backups must use the same backup selection list.

See [“About selection lists”](#) on page 295.

The Full - Using modified time backup method adds the time of the backup to the Backup Exec database only if the full backup job completes successfully. If the full backup job does not complete successfully, any subsequent backup jobs that use the differential or incremental modified time backup methods back up all of the data instead of just the data that changed since the last full backup.

About using the Windows NTFS Change Journal to determine changed files

For Windows 2000 or later systems, you can enhance incremental and differential backup performance by selecting to have Backup Exec use the information recorded in the NTFS Change Journal. NTFS logs all file system changes in the Change Journal. If you select to use the Change Journal and select Differential - Using modified time or Incremental - Using modified time as the backup method, Backup Exec will scan the journal to get a list of changed files rather than scan all files, reducing the amount of time required to perform the incremental or differential backup.

About selecting data to back up

When you are setting up a backup job, select the data you want to back up. Make your selections from the backup selections pane on the **Backup Job Properties** dialog box.

You can find a list of icons that appear in the backup selections pane at the following URL:

<http://entsupport.symantec.com/umi/V-269-12>

There are several ways you can select data to back up. You can select an entire drive, a folder, files, System State, network share, Backup Exec Agent volume, or user-defined selection. You also can use the Advanced File Selection feature to include or exclude specific files or specific types of files. Or you can set up a selection list that you can reuse for several backups.

Note: If the account to which you are logged on does not have sufficient rights, you are required to supply another logon account that can be used to view files for backup.

To expand the view for an item, click the plus sign (+) next to it or double-click the item name. To collapse the view, click the minus sign (-) next to an item or double-click the item name.

To view the contents of an item, double-click the item's icon. The item's contents appear in the right frame of the backup selections view. For all items (except System State), you can traverse file levels from either side of the window by clicking folders and subfolders as they appear.

When you are browsing remote selections, Backup Exec requires a valid logon account to expand the resources and devices. If the default logon account does not enable access to a remote selection, Backup Exec prompts you to select another existing logon account or create a new logon account that can access the selection.

To select data, select the check box next to the drive or directory you want to back up.

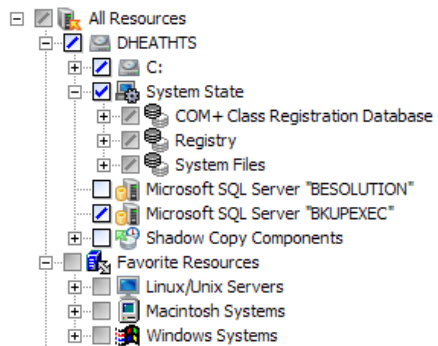
Figure 5-1 Data selection

A slash in a shaded check box means some items below the check box are selected, but the item itself cannot be selected.

A slash in a check box means some items below the directory or drive level are selected.

A check mark in a check box means all items at or below the directory or drive level are selected.

A clear check box means the item can be selected.



About using fully qualified computer names in backup selections

You can enter fully qualified computer names in Backup Exec anywhere that you can enter a computer name. In addition, Backup Exec can show fully qualified computer names where computer names are listed.

For fully qualified computer names, the following rules apply:

- The maximum number of characters for each label (the text between the dots) is 63
- The maximum total number of characters in the fully qualified name is 254, including the dots, but excluding the \\
- The name cannot include the following characters: * | < > ?

To find a fully qualified computer name, from the Control Panel, select System > Computer Name. The fully qualified name appears in the Full computer name field.

Symantec does not recommend using both fully qualified computer names and non-qualified computer names in selection lists. Symantec recommends using fully qualified computer names.

For example, if you have a computer named Test_Computer, you can have two selections for it. One selection is called Test_Computer. The fully qualified selection is called Test_Computer.domain.company.com. In this case, Backup Exec treats each selection as a separate computer, even though both selections are for the same computer. For backup jobs that use the short computer name, the catalog contains the short computer name. For backup jobs that use the fully qualified name, the catalog contains the fully qualified name.

About the Computer Name node in the backup selections list

The first node under **All Resources** shows the name of the computer on which Backup Exec is installed.

The **Computer Name** node includes the following sub-nodes:

Table 5-3 Sub-nodes of the Computer Name node

Sub-node name	Description
Local Drives	Includes hard drives as well as CD-ROM drives that physically reside on the media server.
Shadow Copy Components (Windows Server 2003/2008)	Uses Microsoft's Volume Shadow Copy Service to protect critical operating system and application service data, and third-party application and user data on Windows Server 2003/2008 resources.
System State	<p>Lists a collection of system-specific data that is backed up whenever the computer name node is selected. Symantec recommends that you back up System State. However, you can clear the check box next to System State if you do not want to back it up with the resources on the server. You cannot select or expand the System State resources individually. They are backed up only as a collection, never individually.</p> <p>You can only perform a full backup on System State. However, if you select other items at the same time, you can perform other backup methods on those items. You can back up System State remotely on other computers if Backup Exec Remote Agent is installed on the remote computer. For more information about System State, refer to your Microsoft Windows documentation.</p>
Active Directory Application Mode	Appears only when Active Directory Application Mode (ADAM) resources are available for backup. Even though ADAM is a Shadow Copy component, the ADAM node is not selected automatically when you select the Shadow Copy Components node. You must select the ADAM node if you want to back up the ADAM resources.
Backup Exec database	Includes job, schedule, job history, notification, alerts, device, media, and catalog indexes data for Backup Exec.
Utility Partition	Includes the utility partitions that are installed on the system and available for backup. Individual utility partition objects are named Utility Partition on Disk <i>disk_number</i> (for example, Utility Partition on Disk 0), and cannot be expanded. Backing up utility partitions is recommended when a full system backup is done, such as for disaster recovery preparation. Utility partitions can be backed up individually. If there are not utility partitions on the system, this resource is not available. Administrative rights are required to browse and back up utility partitions.

Table 5-3 Sub-nodes of the Computer Name node (*continued*)

Sub-node name	Description
EFI System Partition	Appears if an Extensible Firmware Interface (EFI) system partition is on the computer. In most cases, each computer will have only one EFI system partition. However, if more than one exists on a computer, Backup Exec displays only the active partition.

See [“How the Active Directory Recovery Agent works”](#) on page 882.

About the Favorite Resources node in the backup selections list

The **Favorite Resources** node lists the remote computers that are set up to publish information to the media server. Several nodes may appear under the **Favorite Resources** node, depending on the type of remote computers that you use.

The nodes that may appear include the following:

- **Linux/Unix Servers**

See [“About publishing Linux, UNIX, and Macintosh computers to media servers”](#) on page 1882.

- **Macintosh Servers**

See [“About publishing Linux, UNIX, and Macintosh computers to media servers”](#) on page 1882.

- **NetWare Agents**

See [“About publishing NetWare servers to the NetWare agents list”](#) on page 1933.

- **Windows Systems**

See [“About publishing the Remote Agent for Windows Systems to media servers”](#) on page 1951.

From the **Windows Systems** node, you can add or delete remote Windows computers. When you add or delete a remote computer, it may take a few minutes for the computer name to appear or to be removed from the **Windows Systems** node.

Backup Exec automatically deletes a remote computer from the **Windows Systems** node in the following situations:

- The remote computer becomes disconnected from the network.
- The Remote Agent is uninstalled from the remote computer.

If the media server receives published information from the remote computer again, Backup Exec adds the computer name to the **Windows Systems** node again.

If the media server does not receive published information within a 24-hour period, a user attention icon appears next to the remote computer's name. This icon is displayed for 13 days.

Some of the reasons why the media server may not receive published information include the following:

- The publishing option is disabled on the Remote Agent.
- A media server is removed from the list to publish to.

Adding a Windows system to the Favorite Resources node in the backup selections list

To add a Windows system to the **Favorite Resources** node, you must know the name or IP address of the remote computer.

Note: It may be several minutes before the computer name appears under the node.

To add a Windows system to the Favorite Resources node in the backup selections list

- 1 On the navigation bar, click **Job Setup**.
- 2 Do one of the following:

<p>To work with a backup job that is associated with a policy</p>	<p>In the Backup Selection Lists pane, click the backup job with which you want to work.</p>
<p>To work with a backup job that is not associated with a policy</p>	<p>In the Jobs pane, click the backup job with which you want to work.</p>
- 3 In the Task pane, under **General Tasks**, click **Properties**.
- 4 On the backup selection tree, expand the **Favorite Resources** node.
- 5 Right-click **Windows Systems**.
- 6 Click **Add Windows System**.
- 7 In the **System Name** field, type the name of the Windows computer that you want to add.

8 Do one of the following:

To install the Remote Agent on a Windows computer and then add it to the Favorite Resources node	Select Install the Remote Agent, and then add to Favorite Resources .
--	--

If you select this option, the installation wizard appears when you complete the dialog box.

See [“About installing the Remote Agent for Windows Systems”](#) on page 139.

To add a Windows computer that already has the Remote Agent installed on it	Select Add a system that already has the Remote Agent installed .
---	--

9 In the **Logon Account** field, select the logon account you use for the remote computer. Click **New** to add a new logon account instead.

10 Click **OK**.

Deleting a Windows system from the Favorite Resources node in the backup selections list

You can delete a Windows system from the **Favorite Resources** node at any time.

Note: It may be several minutes before the computer name is removed from the node.

To delete a Windows system from the Favorite Resources node in the backup selections list

1 On the navigation bar, click **Job Setup**.

2 Do one of the following:

To work with a backup job that is associated with a policy

In the **Backup Selection Lists** pane, click the backup job with which you want to work.

To work with a backup job that is not associated with a policy

In the **Jobs** pane, click the backup job with which you want to work.

3 In the Task pane, under **General Tasks**, click **Properties**.

4 On the backup selection tree, expand the **Favorite Resources** node, and then expand the **Windows Systems** node.

- 5 Right-click the Windows system that you want to delete.
- 6 Click **Delete from Favorite Resources**.

Add Windows System options

You can add a Windows system to the **Favorite Resources** in the backup selections list.

See [“Adding a Windows system to the Favorite Resources node in the backup selections list”](#) on page 285.

Table 5-4 Add Windows System options

Item	Description
System Name	Specifies the name or IP address of the remote computer that you want to add.
Install the Remote Agent, and then add to Favorite Resources	Installs the Remote Agent on the remote computer, and then adds the remote computer to the Favorite Resources .
Add a system that already has the Remote Agent installed	Adds the remote computer to the Favorite Resources node
Logon Account	Specifies the logon account to use for the remote computer. This option is not available if you selected Install the Remote Agent, and then add to Favorite Resources .
New	Allows you to select a new logon account to use for the remote computer. This option is not available if you selected Install the Remote Agent, and then add to Favorite Resources .

About the Domains node in the backup selections list

The Domains resource includes Active Directory Domains and the Microsoft Windows Network. The Microsoft Windows Network node enables you to browse to resources in a Microsoft Windows network.

The Active Directory Domains node enables you to browse Active Directory domains. Backup Exec automatically discovers the Active Directory domain to which the media server belongs and displays it in the backup selection list. Also, you can manually add Active Directory domains to the backup selections list.

You cannot select the Active Directory Domains node itself for backup. You must expand the node to browse Active Directory domains.

When you add an Active Directory domain, you must use a fully qualified domain name. An example of a fully qualified domain name is domain.companyname.com.

For fully qualified domain names, the following rules apply:

- The maximum number of characters for each label (the text between the dots) is 63
- The maximum total number of characters in the fully qualified domain name is 254, including the dots, but excluding the \\
- The name cannot include the following characters: * | < > ?

Adding an Active Directory domain to the Active Directory Domains node

You must know the fully qualified domain name of the Active Directory domain.

To add an Active Directory domain

- 1 On the navigation bar, click **Job Setup**.
- 2 Do one of the following:

To work with a backup job that is associated with a policy

In the **Backup Selection Lists** pane, click the backup job with which you want to work.

To work with a backup job that is not associated with a policy

In the **Jobs** pane, click the backup job with which you want to work.

- 3 In the Task pane, under **General Tasks**, click **Properties**.
- 4 On the backup selections tree, expand the **Domains** node.
- 5 Right-click **Active Directory Domains**.
- 6 Click **Manage Active Directory Domains**.
- 7 In the **Name** box, type the fully qualified domain name.
- 8 Click **Add**.
- 9 Click **Close**.

Deleting an Active Directory domain from the Active Directory Domains node

You can delete an Active Directory domain from the Active Directory Domains node if you no longer need it.

To delete an Active Directory domain

- 1 On the navigation bar, click **Job Setup**.
- 2 Do one of the following:

To work with a backup job that is associated with a policy

In the **Backup Selection Lists** pane, click the backup job with which you want to work.

To work with a backup job that is not associated with a policy

In the **Jobs** pane, click the backup job with which you want to work.

- 3 In the Task pane, under **General Tasks**, click **Properties**.
- 4 On the backup selections tree, expand the **Domains** node.
- 5 Right-click **Active Directory Domains**.
- 6 Click **Manage Active Directory Domains**.
- 7 Select the domain you want to delete in the **Domains** list.
- 8 Click **Delete**.
- 9 Click **Close**.

Manage Active Directory Domains options

You can add or delete Active Directory domains from **Active Directory Domains** in the backup selections.

See [“Adding an Active Directory domain to the Active Directory Domains node”](#) on page 288.

See [“Deleting an Active Directory domain from the Active Directory Domains node”](#) on page 289.

Table 5-5 Manage Active Directory Domains options

Item	Description
Name	Specifies the name of the Active Directory domain you want to add or delete from Active Directory Domains in the backup selections.
Domains	Lists the domains that currently reside in Active Directory Domains in the backup selections.
Add	Adds the new domain you specified to Active Directory Domains in the backup selections.
Delete	Deletes the domain you selected in the Domains list from Active Directory Domains in the backup selections.

About the User-defined Selections node in the backup selections list

You can create shortcuts to shares and save them as user-defined selections. Use this feature to quickly access shares that have a very long path or that are unavailable when you set up a backup job. A share may be unavailable because the network resources that are used to locate the computer are offline, even though the computer may still be running and available. In some cases this happens because the computer is on the Internet and accessible from within the company's private network, but cannot be located by using just its name or normal browsing methods.

See [“Adding a user-defined selection to the User-defined Selections node”](#) on page 290.

See [“Deleting a user-defined selection from the User-defined Selections node”](#) on page 292.

Adding a user-defined selection to the User-defined Selections node

You can set up direct access to a share by entering its Universal Naming Convention (UNC) path name or computer name, or a fully qualified computer

name. The selections that you specify can be selected for backup operations from the **User-defined Selections** node.

Table 5-6 User-defined selection formats

Format type	Example
UNC name	\\mycomputer\shared\temp
Fully qualified computer name	\\mycomputer.domain.companyname.com\temp

See [“About using fully qualified computer names in backup selections”](#) on page 282.

To add a user-defined selection to the User-defined Selections node

- 1** On the navigation bar, click **Job Setup**.
- 2** Do one of the following:

<p>To work with a backup job that is associated with a policy</p>	<p>In the Backup Selection Lists pane, click the backup job with which you want to work.</p>
<p>To work with a backup job that is not associated with a policy</p>	<p>In the Jobs pane, click the backup job with which you want to work.</p>
- 3** In the Task pane, under **General Tasks**, click **Properties**.
- 4** On the backup selections tree, right-click **User-defined Selections**, and then click **Manage User-defined Selections**.
- 5** In the **Name** box, type the server name and volume name, the computer name, or fully qualified computer name.

You can provide TCP/IP addresses for user-defined selections, but Symantec does not recommended it. Backup Exec does not support user-defined selections for IP addresses in a Dynamic Host Configuration Protocol (DHCP) environment.
- 6** Click **Add**.
- 7** When you are finished adding selections, click **Close**.

Deleting a user-defined selection from the User-defined Selections node

You can delete a user-defined selection from the **User-defined Selections** node if you no longer need it.

See [“About the User-defined Selections node in the backup selections list”](#) on page 290.

To delete a user-defined selection from the User-defined Selections node

- 1 On the navigation bar, click **Job Setup**.
- 2 Do one of the following:

To work with a backup job that is associated with a policy	In the Backup Selection Lists pane, click the backup job with which you want to work.
To work with a backup job that is not associated with a policy	In the Jobs pane, click the backup job with which you want to work.
- 3 In the Task pane, under **General Tasks**, click **Properties**.
- 4 On the backup selections tree, right-click **User-defined Selections**, and then click **Manage User-defined Selections**.
- 5 Select the user-defined selection you want to delete in the **Selections defined** list.
- 6 Click **Delete**.
- 7 Click **Close**.

User-defined Selections options

You can add or delete user-defined selections from **User-defined Selections** in the backup selections.

See [“Adding a user-defined selection to the User-defined Selections node”](#) on page 290.

See [“Deleting a user-defined selection from the User-defined Selections node”](#) on page 292.

Table 5-7 User-defined Selections options

Item	Description
Name	Specifies the name of the user-defined selection you want to add or delete from User-defined Selections in the backup selections.
Selections defined	Lists the user-defined selections that are currently defined under User-defined Selections in the backup selections.
Add	Adds the new user-defined selection you specified to User-defined Selections in the backup selections.
Delete	Deletes the domain you selected in the Selections defined list from User-defined Selections in the backup selections.

About managing Microsoft Virtual Hard Disk (VHD) files in Backup Exec

Microsoft Windows 2008 R2 gives users the ability to create native Virtual Hard Disk (VHD) files. VHD files are virtual hard disks contained in a single file. For more information about VHD files, see your Microsoft Windows documentation.

Backup Exec gives you the ability to back up and restore native VHD files. If a native VHD file is not mounted, you can back up the volume on which it resides normally.

If a native VHD file is mounted to a drive letter or to an empty folder path, the file is skipped during backup jobs. You cannot include a mounted VHD as part of a selection list. To back up the data in a mounted VHD file, select its mount point in the backup selections.

See [“Creating a backup job by setting job properties”](#) on page 332.

You can restore native VHD files as part of any normal restore job. You can also redirect a restore job to a native VHD if you use Microsoft Windows 2008 R2. When you redirect a restore job to a native VHD, Backup Exec creates a VHD file that expands dynamically as you save data to it. The file expands until it reaches 2040 GB, which is the maximum size for a native VHD file. You can create one VHD file with data from all redirected backup sets or you can create a VHD file for each backup set.

See [“About redirecting restore jobs to native Microsoft Virtual Hard Disk \(VHD\) files”](#) on page 636.

How to back up user-defined Microsoft Windows Distributed File System data

The Microsoft Distributed File System (DFS) feature consists of DFS Namespaces and DFS Replication technologies. To back up user-defined DFS configuration settings and file system data, Symantec recommends specific backup selections.

Backup Exec supports the following:

- DFS for Windows Server 2003 and earlier
- DFS Namespace for Windows Server 2003 R2 and later
- DFS Replication for Windows Server 2003 R2 and later
- File Replication Service (FRS) for Windows Server 2003 and earlier

The following backup selections are recommended for DFS:

Table 5-8 Recommended backup selections for DFS

DFS item to back up	Recommended backup selections
Stand-alone DFS or DFS Namespaces configuration settings	The following selections should be backed up: <ul style="list-style-type: none">■ The System State registry of the server that hosts the DFS root■ The System State registry of all remote servers that host target shares
Domain-based DFS or DFS Namespaces configuration settings	The following selections should be backed up: <ul style="list-style-type: none">■ The System State registry of the target server■ The Active Directory of the domain controller that hosts the DFS root■ The System State registry of all the remote servers that host target shares <p>Note: You cannot restore domain DFS or DFSN configuration settings from Active Directory backups for which the Granular Recovery Technology option was enabled.</p>

Table 5-8 Recommended backup selections for DFS (*continued*)

DFS item to back up	Recommended backup selections
DFS Namespaces shared data, if Microsoft replication technologies are not used	The system volume of the server that hosts the shared folders or targets
FRS configuration settings for Windows Server 2003 and earlier	The System State registry and the Active Directory of the domain controller that hosts the replicated data Note: You cannot restore FRS configuration settings from Active Directory backups for which the Granular Recovery Technology option was enabled.
FRS data for Windows Server 2003 and earlier	The system volume on any server that hosts the replicated data

About selection lists

Selection lists provide a quick and easy way of selecting files that you back up often. After you choose devices, directories and files, you can save the selections as a selection list that you can use in regularly scheduled operations or once-only operations. Selection lists, which define what is to be backed up, are also automatically created when you create a backup. You can combine a selection list with a policy and quickly create a backup job.

Backup Exec detects and notifies you about items in a selection list that are no longer on the resource. Notification occurs as a selection list is loaded for local selections, and as any remote server is expanded in the tree.

You can also choose to notify recipients when a job completes that contains a particular selection list. This feature allows you to notify users who may be interested that a particular set of selections was backed up. The completion status of the job is included in the notification.

You can view the job history of the jobs that are associated with a selection list.

See [“Viewing the history for backup selection lists”](#) on page 314.

See [“Creating selection lists”](#) on page 296.

See [“Merging selection lists”](#) on page 300.

See [“Replacing selection lists”](#) on page 301.

See [“Copying selection lists”](#) on page 302.

See [“Deleting selection lists”](#) on page 303.

See [“Editing selection lists”](#) on page 304.

See [“Creating separate selection lists for each computer or resource”](#) on page 309.

Creating selection lists

A backup selection list is a list of the resources that you want to back up. After you create a selection list you can use it with any backup job or policy.

Depending on how you set the Backup Exec default options for selection lists, Backup Exec will do one of the following when you create a new selection list:

- Create a separate selection list for each computer you choose.
- Create a separate selection list for each resource you choose.
- Create one selection list, regardless of the number of computers or resources you choose.

See [“About selection lists”](#) on page 295.

To create a selection list

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Selection Lists Tasks**, select **New backup selection list**.
- 3 Select the resources that you want to back up from the backup selections pane.
- 4 Select the appropriate options.
See [“New Backup Selection List options”](#) on page 297.
- 5 (Optional) To change the order in which the resources in the selection list are backed up, in the **Properties** pane, under **Source**, click **Resource Order**.
See [“Resource Order Backup options”](#) on page 338.
- 6 (Optional) To change or test a logon account for the resources, in the **Properties** pane, under **Source**, click **Resource Credentials**.
See [“Resource Credentials options”](#) on page 337.
- 7 (Optional) To set the priority for processing the jobs associated with the selection list, or to set a time range when the resources in the list will be available for backup, in the **Properties** pane, under **Source**, click **Priority and Availability**.

See [“Priority and Availability backup options”](#) on page 308.

- 8 (Optional) To notify users that a job containing this selection list was completed, in the **Properties** pane, under **Source**, click **Selection List Notification**.
 See “[Notification options for jobs](#)” on page 682.
- 9 (Optional) To select a preferred server or servers for the resources in the selection list, in the **Properties** pane, under **Source**, click **Preferred Servers**.
 See “[Preferred Servers backup options](#)” on page 437.
- 10 If you are creating a selection list for a CASO environment, do the following in the order listed:
 - In the **Properties** pane, under **Destination**, click **Device**.
 - Check **Restrict backup of the selection list to devices on the following media server or media servers in a pool**.
 - Select the media server from the drop-down list.
- 11 Click **OK**.

New Backup Selection List options

A backup selection list is a list of the resources that you want to back up.

See “[Creating selection lists](#)” on page 296.

The **New Backup Selection List** dialog box contains the following options:

Table 5-9 **New Backup Selection List options**

Item	Description
Selection list name	Designates the name of this selection list.
Load selections from existing list	Loads an existing selection list or merges multiple selection lists. See “ Merging selection lists ” on page 300.
Selection list description	Describes this selection list.
Include/Exclude	Lets you use the Advanced File Selection for selecting files for backing up. See “ Backup Include/Exclude Selections options ” on page 298.
Include subdirectories	Selects the contents of all the subfolders when a directory is selected.

Table 5-9 New Backup Selection List options *(continued)*

Item	Description
Show file details	Displays any details about the files available for selecting.
View by Resource	Lets you view resources in a tree view.
View Selection Details	Lets you view selections as a list of files and directories.

Backup Include/Exclude Selections options

Advanced file selection allows you to quickly select or de-select files for backup operations by specifying file attributes.

The **Include/Exclude Selections** dialog box contains the following fields:

Table 5-10 Include/Exclude Selections options

Option name	Description
General	Lets you include or exclude any type of resource other than NDMP.
NDMP	Lets you include or exclude NDMP resources.
Resources	Allows you to include or exclude files from a backup of a different drive than the one you selected previously on the Backup Job Properties dialog box.

Table 5-10 Include/Exclude Selections options (*continued*)

Option name	Description
Path	<p>Specifies the name of the folder and/or subfolder that contains any specific file you want to include or exclude.</p> <p>You can use wildcard characters. use a question mark (?) to represent any single character. Use two asterisks (**) to represent any number of characters.</p> <p>For example, on your C: drive you have a My Documents folder that contains a subfolder called Work Files. There are three Work Files subfolders called 1999, 2000, and 2001. Each one of those subfolders has a subfolder called Personnel.</p> <p>If you type the path as \My Documents**\Personnel, the backup will include or exclude the following:</p> <ul style="list-style-type: none"> ■ C:\My Documents\Work Files\2001\Personnel ■ C:\My Documents\Work Files\2000\Personnel ■ C:\My Documents\Work Files\1999\Personnel <p>In addition, every subfolder below the ** wildcard is included or excluded. However, the only files from the subfolders that are included or excluded are those that match the file name you type in the File field. So in the example above, every subfolder of C:\My Documents is included in or excluded from the backup, and only the files that match the file name you type in the File field are included or excluded.</p> <p>Note: The Remote Agent for NetWare Systems does not support wildcard characters.</p> <p>After you type the path, type the file name in the File field.</p>
File	<p>Specifies the file you want to include in or exclude from the backup.</p> <p>You can use wildcard characters. Use a question mark (?) to represent any single character. Use two asterisks (**) to represent any number of characters.</p> <p>For example, to include all files with the .exe extension, type *.exe.</p> <p>After you type the file name, indicate whether you want to include or exclude it.</p>
Include	<p>Specifies that the files that you selected should be included in the job. This is the default option.</p>
Exclude	<p>Specifies that the files that you selected should be excluded from the job.</p>

Table 5-10 Include/Exclude Selections options (*continued*)

Option name	Description
Include subdirectories	Includes the contents of all the subfolders when a directory is selected.
Only modified files	Includes or excludes modified files in the path you specify.
Only read-only files	Includes or excludes files that cannot be modified.
Files dated	Includes or excludes the files created or modified during a specific time period. Then select the beginning and ending dates. Note: The Remote Agent for NetWare Systems does not support this option.
Files not accessed in x days	Includes or excludes files that have not been accessed in a specified number of days. This is useful when you need to migrate older files from your system.

Merging selection lists

You can create a new selection list by merging two or more existing lists with new selections.

To merge and replace selection lists

- 1 On the navigation bar, click **Job Setup**.
- 2 Under **Selection Lists Tasks** in the task pane, select **New backup selection list**.
- 3 On the **New Backup Selection List** dialog box, select resources to include in the selection list.
See “[New Backup Selection List options](#)” on page 297.
- 4 Click **Load selections from existing list**.
- 5 Select the selection lists that you want to merge with the previously selected backup selections.
See “[Load Selections from Existing List options](#)” on page 301.
- 6 Click **Merge**.
- 7 Complete the other options on the **New Backup Selection List** dialog box.
See “[Creating selection lists](#)” on page 296.
- 8 Click **OK**.

Replacing selection lists

You can replace selections in the selection tree with other selection lists.

To replace selection lists

- 1 On the navigation bar, click **Job Setup**.
- 2 Under **Selection Lists Tasks** in the task pane, select **New backup selection list**.
- 3 On the **New Backup Selection List** dialog box, select resources to include in the selection list.
 See “[New Backup Selection List options](#)” on page 297.
- 4 Click **Load selections from existing list**.
- 5 Select the selection lists that you want to replace the previously selected backup selections.
 See “[Load Selections from Existing List options](#)” on page 301.
- 6 Click **Replace**.
- 7 Complete the other options on the **New Backup Selection List** dialog box.
 See “[Creating selection lists](#)” on page 296.
- 8 Click **OK**.

Load Selections from Existing List options

You can merge selection lists to create a new selection list. You can also replace the selections with an existing selection list.

Table 5-11 Load Selections from Existing List options

Item	Description
Name	Displays the names of existing selection lists.
Description	Displays the descriptions of existing selection lists.
Properties	Lets you view the properties of the selected selection list.
Replace	Replaces the items in the selection tree with the selection list you selected in the Name column.

Table 5-11 Load Selections from Existing List options (*continued*)

Item	Description
Merge	Merges the items in the selection tree with the selection list you selected in the Name column.

Copying selection lists

You can copy a selection list to reuse it on a new media server. You can also copy a selection list to the same media server and then edit its settings to create an entirely new selection list.

See [“Editing selection lists”](#) on page 304.

To copy a selection list

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Backup Selection Lists** pane, right-click the selection list you want to copy, and then click **Copy**.
- 3 Complete the appropriate options.
See [“Copy Selection List options”](#) on page 302.
- 4 Click **OK**.

Copy Selection List options

You can copy a selection list to one or more media servers.

See [“Copying selection lists”](#) on page 302.

Table 5-12 Copy Selection List options

Item	Description
Copy to this media server	Copies the selection list to the media server on which the selection list currently resides.
Copy to other media servers	Copies the selection list to other media servers.
Name	Indicates the name of the destination media servers to which you can copy the selection list.
Logon Account	Indicates the logon account for each destination media server.

Table 5-12 Copy Selection List options (continued)

Item	Description
Add	Lets you add a new media server to the list of destinations.
Edit	Lets you edit information about the selected media server.
Remove	Removes the selected media server from the list of destinations.
Import List	Imports a list of media servers.
Overwrite selection lists with identical names that already exist on the destination media server	Lets you overwrite selection lists on the destination media server if they have the same name as the selection list you select to copy.

Holding jobs that back up a selection list

You can place all jobs that back up a selection list on hold to prevent the jobs from running. The jobs do not run until you change the job's hold status.

To hold jobs that back up a selection list

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Backup Selection Lists** pane, right-click the selection list whose jobs you want to place on hold, and then click **Hold Jobs**. You can select multiple selection lists by selecting a selection list, and then pressing the <Ctrl> or <Shift> keys while you click other selection lists.

The jobs that back up this selection list are placed on hold.

- 3 To remove the hold and run the jobs according to the schedule, click **Remove Hold**.

Deleting selection lists

You can delete a selection list. However, if a selection list is associated with a policy, you must remove the selection list's association from the policy before you can delete the selection list.

See [“About selection lists”](#) on page 295.

You cannot delete the Excludes selection list.

See [“To edit the Excludes selection list”](#) on page 305.

To delete selection lists

- 1 On the **Edit** menu, click **Manage Selection Lists**.
- 2 Click the selection list that you want to delete.
- 3 Click **Delete**.
- 4 Click **Yes** to delete the selection list or click **No** to cancel the delete operation.
 If the selection list is being used by a job, you will not be able to delete it.
- 5 Click **Close**.

Manage Selection Lists options

You can delete or edit existing selection lists.

See [“Deleting selection lists”](#) on page 303.

See [“Editing selection lists”](#) on page 304.

Table 5-13 **Manage Selection Lists** options

Item	Description
Name	Displays the name of the selection list.
Type	Displays the type of selection list.
Edit	Lets you edit the selected selection list.
Delete	Deletes the selected selection list.

Editing selection lists

Editing a selection list affects all jobs that use the selection list. However, if you edit a selection list that is being used by an active job, the changes do not affect that job. If you want to only edit selections for a specific job, edit the job rather than the selection list.

If a resource on your selection list no longer exists and you want to delete it, you must use the **View Selection Details** tab.

To edit a selection list

- 1 On the **Edit** menu, click **Manage Selection Lists**.
- 2 Select the selection list that you want to edit.
- 3 Click **Edit**.

- 4 Edit the selection list properties.
 See [“New Backup Selection List options”](#) on page 297.
- 5 Click **OK**.

Editing the Excludes selection list

You can change the Excludes selection list at any time.

See [“How to include or exclude files for backup”](#) on page 355.

See [“About selection lists”](#) on page 295.

To edit the Excludes selection list

- 1 On the **Edit** menu, click **Manage Selection Lists**.
- 2 On the **Manage Selection Lists** dialog box, select **Excludes**.
- 3 Click **Edit**.
- 4 Do one of the following:

<p>If the Excludes list was edited previously</p>	<ul style="list-style-type: none"> ■ Select the selection rule that you want to edit. ■ Click Edit.
<p>If this is the first time the Excludes list is being edited</p>	<p>Click Insert.</p>
- 5 Edit the selection list properties.
 See [“New Backup Selection List options”](#) on page 297.
- 6 If you want to delete one of the selection rules:
 - Select the selection rule that you want to delete.
 - Click **Delete**.
- 7 Click **OK**.
- 8 Click **Close**.

Excludes Properties options

You can exclude or include new files or folders.

See [“Editing the Excludes selection list”](#) on page 305.

Table 5-14 Excludes Properties options

Item	Description
Selection list name	Displays the selection list name. If you edit the Excludes properties it should say "Excludes."
Selection list description	Describes the Excludes selection list. You can enter a description here to help remember the contents of your Excludes selection list.
View Selection Details	Displays the details about the Excludes selection list.
Edit	Lets you edit the Excludes selection list to add or remove files and folders. You can also edit the selection criteria.
Insert	Lets you create criteria and settings for the Excludes selection list.
Delete	Lets you delete criteria and settings from the Excludes selection list.

About priority and availability windows for selection lists

When you create a backup selection list, you can specify the priority for processing the jobs associated with the selection list. In addition, you can set a time range when the resources in the list will be available for backup. The time range is called the availability window. You can set a default availability window for selection lists. When you create a new selection list, the default availability window displays, but you must select the **Limit availability to this daily time window** option in order for the selection list to use the default window.

See [“Setting priority and availability windows for selection lists”](#) on page 307.

You can set one availability window per selection list, and the window is the same for each day of the week. If you merge two or more selection lists or replace a selection list, Backup Exec uses the availability window of the original list.

If you schedule a job to run outside of the availability window, the job does not run and Backup Exec displays an Invalid Schedule status for the job on the **Job Monitor**. For example, you set the availability window to allow resources to be available for backup between the hours of 11:00 p.m. and 6:00 a.m. If you schedule a backup job to run at 7:00 a.m., the job will not run because the resources are not

available at that time. When scheduling a job, be sure that the schedule overlaps the availability window for the resources.

Setting default priority and availability windows for all selection lists

You can set a default availability window for selection lists. When you create a new selection list, the default availability window displays, but you must select the **Limit availability to this daily time window** option in order for the selection list to use the default window.

See [“About priority and availability windows for selection lists”](#) on page 306.

You can also specify priority or set an availability window for specific selection lists.

See [“Setting priority and availability windows for selection lists”](#) on page 307.

To set default priority and availability windows for all selection lists

- 1 On the **Tools** menu, select **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Priority and Availability**.
- 3 Complete the appropriate options.

See [“Priority and Availability backup options”](#) on page 308.

Setting priority and availability windows for selection lists

When you create a backup selection list, you can specify the priority for processing the jobs associated with the selection list.

See [“About priority and availability windows for selection lists”](#) on page 306.

You can also set a default availability window for selection lists. When you create a new selection list, the default availability window displays, but you must select the **Limit availability to this daily time window** option in order for the selection list to use the default window.

See [“Setting default priority and availability windows for all selection lists”](#) on page 307.

To set priority and availability windows for selection lists

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Selection Lists Tasks**, select **New backup selection list**.
- 3 Select the data you want to back up.

- 4 In the **Properties** pane, under **Source**, click **Priority and Availability**.
- 5 Complete the appropriate options.
See “[Priority and Availability backup options](#)” on page 308.

Priority and Availability backup options

When you create a backup selection list, you can specify the priority for processing the jobs that are associated with the selection list. In addition, you can set a time range when the resources in the list will be available for backup.

See “[Setting priority and availability windows for selection lists](#)” on page 307.

Table 5-15 Priority and Availability backup options

Item	Description
Job priority	Displays the priority of the access to the devices for this job. See “ About job priority ” on page 195.
Limit availability of this selection list for backup to the following daily time window	Enables the availability window, which specifies when the selection list will be available for backup each day. If you do not select this option, the resources in the selection list are always available for backup. Backup Exec considers both the resource's availability window and the job's time window when it runs a job. If you schedule a job to run outside of the availability window, it does not run. Backup Exec displays an Invalid Schedule status for the job on the Job Monitor . When you schedule a job, be sure that the job's time window is within the availability window for the resources. See “ Schedule options ” on page 356.
First available date	Designates the first date when the selection list should be available to be backed up. The list will be available every day from this date onward.
Begin time	Designates the earliest time when this selection list will be available for backup.
End time	Designates the latest time when this selection list will be available for backup.

Table 5-15 Priority and Availability backup options (*continued*)

Item	Description
Enable automatic cancellation for this selection list	Cancels the job that are associated with this selection list if the job does not complete within the selected number of hours or minutes. Backup Exec starts timing the length of time the job takes to run when the job is queued, not when the job begins.
Cancel backup job if not completed within x	Designates the number of hours or minutes you want to allow jobs to complete before they are automatically canceled. The default amount of time is three hours.

Creating separate selection lists for each computer or resource

Backup Exec includes default settings that enable separate selection lists to be created for each resource or computer you select when you create a new backup selection list outside of a backup job. This feature does not apply when you create a selection list while creating a backup job.

If you set up Backup Exec to create a separate selection list for each resource or computer, the selection list name will contain either the default name or a user-defined name followed by the name of the computer or resource that you selected for backup.

To create separate selection lists for each computer or resource

- 1 On the **Tools** menu, select **Options**.
- 2 In the **Properties** pane, under **Settings**, select **Selection List**.
- 3 Select the appropriate option:
See “[Application settings for a selection list](#)” on page 309.
- 4 Click **OK**.

Application settings for a selection list

The application settings for a selection list let you create separate selection lists for each resource when you create a backup selection list outside of a job.

See “[Creating separate selection lists for each computer or resource](#)” on page 309.

Table 5-16 Application settings for a selection list

Item	Description
Separate backup selection list for each computer	Creates a different backup selection list for each computer that you select when you create a selection list outside of a backup job.
Separate backup selection list for each resource	Creates a different backup selection list for each resource that you select when you create a selection list outside of a backup job.
Single backup selection list for all selections	Creates one selection list, regardless of the number of resources or the number of computers that are selected for backup. This option is the default option.

Creating a custom filter for backup selection lists

You can filter backup selection lists based on the following criteria:

- Selection list name
- Selection list description
- Computers that are backed up by the selection list
- Resource type
- Policy
- Selection lists that are not protected

To create a custom filter for Backup Selection Lists

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Custom Filter Tasks**, click **Manage custom filters**.
- 3 Click **New**, and then select **Backup Selection Lists Custom Filter**.
- 4 Type a unique name and a description for this filter.

5 Complete the following options as necessary:

To filter by selection list name	In the Properties pane, under Criteria , select Selection List Name . See “New Backup Selection List Name Custom Filter options” on page 312.
To filter by selection list description	In the Properties pane, under Criteria , select Description . See “New Backup Selection List Custom Filter Description options” on page 312.
To filter by a server that is protected by Backup Exec	In the Properties pane, under Criteria , select Protected Server . See “New Backup Selection List Custom Filter Protected Server options” on page 312.
To filter by resource type	In the Properties pane, under Criteria , select Resource Type . See “New Backup Selection List Custom Filter Resource Type options” on page 312.
To filter by policy	In the Properties pane, under Criteria , select Policy . See “New Backup Selection List Custom Filter Policy options” on page 313.
To filter by selection lists that are not currently protected	In the Properties pane, under Criteria , select Not Protected . See “New Backup Selection List Custom Filter Not Protected options” on page 313.

6 Click **OK**.

New Backup Selection List Custom Filter options

You can create custom filters for backup selection lists.

See [“Creating a custom filter for backup selection lists”](#) on page 310.

Table 5-17 **New Backup Selection List Custom Filter options**

Item	Description
Name	Indicates the unique name of the custom filter.
Description	Indicates a description of the custom filter.

New Backup Selection List Name Custom Filter options

The **Selection List Name** field indicates the name of the selection list for which you want to create the custom filter.

See [“Creating a custom filter for backup selection lists”](#) on page 310.

New Backup Selection List Custom Filter Description options

The **Description** field indicates the selection list description for which you want to create the custom filter.

See [“Creating a custom filter for backup selection lists”](#) on page 310.

New Backup Selection List Custom Filter Protected Server options

The **Protected Server** field indicates the name of the protected server for which you want to create the custom filter.

See [“Creating a custom filter for backup selection lists”](#) on page 310.

New Backup Selection List Custom Filter Resource Type options

The **Resource Type** dialog lets you select which types of resources you want to include in the custom filter you create.

See [“Creating a custom filter for backup selection lists”](#) on page 310.

Table 5-18 **New Backup Selection List Custom Filter Resource Type options**

Item	Description
Enable this filter	Enables the resource type criteria for the custom filter you create.
Resource Type	Specifies the types of resources from which you can select. The resources you select are included in the custom filter.

Table 5-18 **New Backup Selection List Custom Filter Resource Type options**
(continued)

Item	Description
Check All	Selects all resources in the Resource Type field.
Uncheck All	Deselects all resources in the Resource Type field.

New Backup Selection List Custom Filter Policy options

The **Policy** dialog lets you select which types of resources you want to include in the custom filter you create.

See [“Creating a custom filter for backup selection lists”](#) on page 310.

Table 5-19 **New Backup Selection List Custom Filter Policy options**

Item	Description
Enable this filter	Enables the policy criteria for the custom filter you create.
Policy	Specifies the policies from which you can select. The policies you select are included in the custom filter.

New Backup Selection List Custom Filter Not Protected options

The **Filter for backup selection lists that are not protected** field lets you create a custom filter that displays backup selection lists that Backup Exec does not protect.

See [“Creating a custom filter for backup selection lists”](#) on page 310.

Filtering backup selection lists

Use filters to view backup selection lists that meet certain criteria.

To filter backup selection lists

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Backup Selection Lists** pane, in the **Filter** list, click the filter that you want to use.

Searching selection lists

Backup Exec includes a search feature for selection lists, which enables you to search for selection lists that back up a particular computer. This feature is helpful when you have a large number of selection lists.

When you complete the search, the **Search Selection Lists** dialog box expands to display the results. You can right-click a selection list to create a new job using policies, copy or delete the selection list, or view the selection list properties.

To search a selection list

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Selection List Tasks**, click **Search backup selection lists**.
- 3 Type the name of the computer for which you want to locate selection lists. You can type the complete name or part of the name.
- 4 Click **Find Now**.

Search Backup Selection Lists options

You can search for the selection lists that back up a particular computer. You can type the complete computer name or part of the name.

See [“Searching selection lists”](#) on page 314.

Viewing the history for backup selection lists

You can view the history of the jobs that use a specific backup selection list.

Backup Exec shows the following history information:

- Job name
- Device name
- Job type
- Job status
- Percent complete
- Start time
- End time
- Elapsed time
- Byte count
- Job rate

- Error code

To view the history for backup selection lists

- 1 On the navigation bar, click **Job Setup**.
- 2 Right-click the backup selection list for which you want to view history.
- 3 Click **View History**.
- 4 Click **OK**.

Viewing a summary for a selection list

You can view the following summary information for a selection list:

- Selections
- Resource order
- Credentials
- Priority and availability
- Notification

To view a summary for backup selection lists

- 1 On the navigation bar, click **Job Setup**.
- 2 Right-click the backup selection list for which you want to view a summary.
- 3 Click **View Summary**.
- 4 Click **OK**.

Selection list summary

You can view a summary of information about a particular selection list.

See “[Viewing a summary for a selection list](#)” on page 315.

You can view the following summary information for a selection list:

- Selections
- Resource order
- Credentials
- Priority and availability
- Notification

About resource discovery

Backup Exec's Resource Discovery feature allows detection of new backup resources within a Windows or Active Directory domain. Using this feature, you can create and schedule a job that searches for new server volumes or databases. You can specify which types of resources to include in the search, and can have Backup Exec send a notification when a new resource is discovered.

Using the discovered resources identified in the job log, you can then create a backup job to ensure that the new resource is protected.

When you set up a resource discovery job, Backup Exec lists all of the Windows domains it has discovered. If you have Active Directory domains, you must add them to the list manually. When the list of domains is in place, you select which domains you want to search for new resources.

The Remote Agent is required to discover resources on remote computers. However, installing a MAPI client on the media server enables Exchange resources to be discovered on remote resources on which the Remote Agent is not installed.

On Windows Server 2003/2008 resources, Backup Exec's Resource Discovery feature detects the Shadow Copy Components; it does not detect System State.

See ["Using resource discovery to search for new resources"](#) on page 316.

Using resource discovery to search for new resources

Backup Exec's Resource Discovery feature allows detection of new backup resources within a Windows or Active Directory domain. Using this feature, you can create and schedule a job that searches for new server volumes or databases. You can specify which types of resources to include in the search, and can have Backup Exec send a notification when a new resource is discovered.

See ["About resource discovery"](#) on page 316.

To use resource discovery to search for new resources

- 1 On the navigation bar, click **Job Setup**.
- 2 Under **Backup Strategy Tasks** in the task pane, select **New job to automatically discover resources**.
- 3 To add an Active Directory Domain to the list of domains to search for new resources, click **Add Active Directory Domain** and then complete the appropriate options.

See ["Add Active Directory Domains options"](#) on page 317.

- 4 Select the domain you want Backup Exec to search for new resources.

- 5 If you need to change the logon account for the domain, click **Change Logon Account** and enter or select the logon credentials to access this domain.
- 6 If you want to exclude computers from the search, in the **Properties** pane, under **Target**, click **Exclude** and then select the computers to exclude.
 See [“Exclude options for resource discovery jobs”](#) on page 318.
- 7 In the **Properties** pane, under **Settings**, click **General** and then complete the appropriate options.
 See [“General options for resource discovery jobs”](#) on page 318.
- 8 In the **Properties** pane, under **Settings**, click **Resources** and then complete the appropriate options.
 See [“Resources options for resource discovery jobs”](#) on page 319.
- 9 If you want Backup Exec to notify someone when this job completes, in the **Properties** pane, under **Settings**, click **Notification**.
 See [“Notification options for jobs”](#) on page 682.
- 10 If you want to run the job now, click **Run Now**. Otherwise, in the **Properties** pane, under **Frequency**, click **Schedule** to set the scheduling options you want to use
 See [“Schedule options”](#) on page 356.

Add Active Directory Domains options

You can add an Active Directory domain to the list of domains Backup Exec searches for resource discovery jobs.

See [“Using resource discovery to search for new resources”](#) on page 316.

Table 5-20 Add Active Directory Domains options

Item	Description
Name	Designates the fully qualified name for the Active Directory domain you want to add.
Domains	Displays the list of domains that Backup Exec uses to discover new resources.
Add	Adds the new Active Directory domain to the list of domains.
Delete	Deletes the selected Active Directory domain from the list of domains.

Domains options for resource discovery jobs

You can search the domains in your environment to discover new resources.

See [“Using resource discovery to search for new resources”](#) on page 316.

Table 5-21 Domains options for resource discovery jobs

Item	Description
Name	Displays the name of the domain.
Logon Account	Displays the logon account Backup Exec uses to access the domain.
Change Logon Account	Lets you change the logon account that Backup Exec uses to access the domain.
Add Active Directory Domain	Lets you add an Active Directory domain to the list of domains.

Exclude options for resource discovery jobs

You can search the domains in your environment to discover new resources. You may want to exclude certain servers or domains from the resource discovery job.

See [“Using resource discovery to search for new resources”](#) on page 316.

Table 5-22 Exclude options for resource discovery jobs

Item	Description
Domain	Displays the servers that are included in the resource discovery job.
Servers excluded	Displays the servers that are excluded from the resource discovery job.
Exclude	Lets you move a server to the list of servers that are excluded from the resource discovery job.
Include	Lets you move a server to the list of servers that are included in the resource discovery job.

General options for resource discovery jobs

You can search the domains in your environment to discover new resources.

See [“Using resource discovery to search for new resources”](#) on page 316.

Table 5-23 General options for resource discovery jobs

Item	Description
Job name	Displays the name for this job.
Job priority	Displays the priority of the access to the devices for this job. See “About job priority” on page 195.

Resources options for resource discovery jobs

You can search the domains in your environment to discover new resources.

See [“Using resource discovery to search for new resources”](#) on page 316.

Table 5-24 Resources options for resource discovery jobs

Item	Description
Network administrative shares	Searches for new administrative network shares or volumes.
Network user shares	Searches for new user-defined shares.
Microsoft SQL databases	Searches for new Microsoft SQL databases.
Microsoft Exchange servers	Searches for new Microsoft Exchange servers. Backup Exec searches for Information Store, Exchange Directory, or storage groups; it does not discover individual databases under storage groups.
Lotus Domino databases	Searches for new Lotus Domino databases.
System State and/or Shadow Copy Components	Searches for new System State resources or shadow copy components.
Oracle databases	Searches for new Oracle databases.
DB2 databases	Searches for new DB2 databases.
Send separate notification for each new resource found	Sends the separate notifications when each new resource is found.
Send one notification for all new resources found	Sends out a single notification for all new resources found.

Table 5-24 Resources options for resource discovery jobs *(continued)*

Item	Description
Include previously discovered resources when sending notification	Sends a notification that includes all resources previously found during resource discovery jobs.

About the Backup Exec Shadow Copy Components file system

The Backup Exec Shadow Copy Components file system uses Microsoft’s Volume Shadow Copy Service to protect critical operating system and application service data, and third-party application and user data on Windows Server 2003/2008 resources.

Volume Shadow Copy Service allows a computer to be backed up while applications and services are running by providing a copy of a volume when a backup is initiated. Applications do not need to be shut down to ensure a successful volume backup. Volume Shadow Copy Service enables third party vendors to create snapshot plug-ins, or Writers, for use with this shadow copy technology.

A Writer is specific code within an application that participates in the Volume Shadow Copy Service framework to provide point-in-time, recovery-consistent operating system and application data. Writers appear as Shadow Copy Components, which are listed as resources in backup and restore selections.

When expanded, the Backup Exec Shadow Copy Components file system includes the following types of Writers:

- Service State - Critical operating system and application service data, such as Event Logs, Windows Management Instrumentation (WMI), and others.
- User Data - Third party application and user data, and others.

Even though ADAM and System State are Shadow Copy component, the **ADAM** node and the **System State** node are not selected automatically when you select the **Shadow Copy Components** node. You must select those nodes if you want to back up those resources.

Only Writers that have been tested for use with Backup Exec are available for selection in the backup selection list. Other Writers may be displayed in the selection list, but cannot be selected for backup.

If you select a volume that contains Shadow Copy data for backup, Backup Exec determines which Shadow Copy files should not be included in a volume level backup. These files will be automatically excluded for backup by a feature called

Active File Exclusion. If this exclusion did not happen during a non-snapshot backup, these files would appear as in use - skipped. If this exclusion did not happen during a snapshot backup, the files would be backed up in a possible inconsistent state, which could create restore issues.

The Windows SharePoint Services feature pack utilizes a SQL (MSDE) instance called SHAREPOINT as a repository for shared information and collaboration data. On Windows Server 2003/2008, in the absence of a Symantec SQL Agent installation, the SQL SHAREPOINT instance can be protected by the Shadow Copy Components file system. If the SQL Agent is installed, then the SQL SHAREPOINT instance can be protected by the SQL Agent.

Note: If Windows SharePoint Services is installed using an instance name other than the default SHAREPOINT instance name, then it cannot be protected by the Shadow Copy Components file system. In that case, the Symantec SQL Agent must be used to protect the SQL SHAREPOINT instance.

Windows Small Business Server 2003 Standard and Premium contain a SQL (MSDE) instance called SBSMONITORING as a repository for server-related activity data. In the absence of a Symantec SQL Agent installation, the SQL SBSMONITORING instance can be protected by the Shadow Copy Components file system. If the SQL Agent is installed, then the SQL SBSMONITORING instance can be protected by the SQL Agent.

How to restore individual items by using Granular Recovery Technology

You can use Granular Recovery Technology (GRT) to restore certain individual items from backup sets. For example, you can use the Agent for Microsoft Exchange Server to restore an email message from a backup without having to restore the entire mailbox. Or, you can use the Agent for Microsoft SharePoint to restore a list without restoring the entire site.

To restore individual items, the Granular Recovery Technology feature must be enabled when you create a backup job.

GRT is enabled by default for backups for the following agents:

- Active Directory Recovery Agent
- Agent for Microsoft Exchange Server
- Agent for Microsoft Hyper-V
- Agent for Microsoft SharePoint

- Agent for VMware Virtual Infrastructure

You can restore either full backup sets or individual items from GRT-enabled backups.

By default, the **Agent for Microsoft Hyper-V** and the **Agent for VMware Virtual Infrastructure** use Granular Recovery Technology to protect files and folders at a granular level. You can also enable the granular recovery of Microsoft Exchange, SQL, and Active Directory application data that resides on virtual machines.

The following table lists the individual items you can restore for each agent.

Table 5-25 Individual items that can be recovered for each agent

Agent	Individual items
Active Directory Recovery Agent	<p>You can restore the following individual items:</p> <ul style="list-style-type: none"> ■ Active Directory objects and attributes ■ Active Directory Application Mode (ADAM) and Active Directory Lightweight Directory Services (AD LDS) objects and attributes
Agent for Microsoft Exchange Server	<p>You can restore the following individual items:</p> <ul style="list-style-type: none"> ■ Mailboxes ■ Mail messages and their attachments ■ Public folders
Agent for Microsoft Hyper-V	<p>You can restore drives, folders, and files from virtual machines that run a Windows operating system.</p> <p>You can also enable the granular recovery of Microsoft Exchange, SQL, and Active Directory application data that resides on virtual machines:</p> <p>See “How Backup Exec protects Microsoft Exchange, SQL, and Active Directory data on virtual machines” on page 1178.</p>

Table 5-25 Individual items that can be recovered for each agent (*continued*)

Agent	Individual items
Agent for Microsoft SharePoint	<p>You can restore the following individual items:</p> <ul style="list-style-type: none"> ■ Portal sites and their associated databases ■ Windows SharePoint Services sites and their associated databases ■ Document library stores (Web Storage System-based) ■ Individual documents that are contained in Document or Picture libraries (Web Storage System-based or Microsoft SQL Server-based) ■ Lists, sites, and sub-sites
Agent for VMware Virtual Infrastructure	<p>You can restore drives, folders, and files from virtual machines that run a Windows operating system.</p> <p>You can also enable the granular recovery of Microsoft Exchange, SQL, and Active Directory application data that resides on virtual machines:</p> <p>See “How Backup Exec protects Exchange, SQL, and Active Directory data on virtual machines” on page 1397.</p>

When you run a GRT-enabled backup job, Backup Exec creates media with an IMG prefix (for example, IMG00001). IMG media is a specific media type that Backup Exec creates only for GRT-enabled backup operations. When you run a GRT-enabled backup job, the IMG media stores the backup data.

Note: Backup-to-disk folders do not support encryption for GRT-enabled jobs.

You should consider which device you use for GRT-enabled backups before you begin. You should also consider any special requirements for the type of data you back up.

See [“Recommended devices for backups that use Granular Recovery Technology”](#) on page 324.

See [“About requirements for jobs that use Granular Recovery Technology”](#) on page 325.

See [“Recommendations for using backup-to-disk folders with backup jobs that use Granular Recovery Technology”](#) on page 511.

See [“How to reclaim disk space for backup jobs that use Granular Recovery Technology”](#) on page 513.

Recommended devices for backups that use Granular Recovery Technology

Symantec recommends that you select a backup-to-disk folder on a volume that does not have file size limitations as the destination for backups that are enabled for Granular Recovery Technology (GRT). An NTFS drive is an example of a volume without file size limitations. Some examples of volumes that have file size limitations include FAT and FAT32 volumes.

If you must use a backup-to-disk folder on a volume with file size limitations, Backup Exec requires a staging location. Backup Exec temporarily stores a small amount of metadata in the staging location during the backup job. It deletes the data from the staging location when the backup is finished. The staging location is not necessary, however, if you use a backup-to-disk folder on a volume without file size limitations as the destination.

The staging location's default path is C:\temp.

The volume that is used for a staging location for backup jobs should meet the following requirements:

- It is local to the media server
- It does not have any file size limitations

Additionally, Symantec recommends the following to avoid disk space problems:

- It should not be a system volume
- It should have at least 1 GB of available space

You can change the default staging location with the other default backup options.

See [“Setting default backup options”](#) on page 391.

Backup Exec also uses a staging location to restore GRT-enabled data from a tape or from a backup-to-disk folder on a volume with file size limitations. The staging location must be on a volume that does not have file size limitations and is local to the media server. The staging location is not necessary if you restore GRT-enabled data from a backup-to-disk folder on a volume without file size limitations, such as NTFS.

Backup Exec uses the staging area differently for the following types of restores:

Table 5-26 Staging processes

Location of data to be restored	Staging process
Tape	<p>Backup Exec copies the entire backup set or sets to the staging area. The staging area must have enough disk space for the entire backup set or sets from which you want to restore an individual item.</p> <p>Before you use a tape device for a GRT-enabled backup, ensure that sufficient disk space is available to perform a restore.</p> <p>Backup Exec deletes the data from the staging area when the restore job is complete.</p>
Backup-to-disk folder that is on a volume with file size limitations (such as FAT or FAT32)	<p>Backup Exec must copy a small amount of metadata that is associated with the backup set to the staging area to complete the restore.</p> <p>Backup Exec deletes the data from the staging area when the restore job is complete.</p>

The staging location's default path is C:\temp. You can change the default restore staging location with the other default restore options.

See [“Setting defaults for restore jobs”](#) on page 638.

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“About requirements for jobs that use Granular Recovery Technology”](#) on page 325.

About requirements for jobs that use Granular Recovery Technology

Keep in mind the following requirements when you use Granular Recovery Technology (GRT) with the agents listed:

Table 5-27 Granular Recovery Technology requirements

Agent	Restrictions
Active Directory Recovery Agent	You can run only full backups for GRT-enabled jobs.

Table 5-27 Granular Recovery Technology requirements (*continued*)

Agent	Restrictions
Agent for Microsoft Exchange Server	<p>Backup Exec must have access to a uniquely named mailbox within the Exchange organization for backup and restore of the Information Store.</p> <p>See “Requirements for accessing Exchange mailboxes” on page 1101.</p> <p>You cannot restore individual mailboxes and messages if both of the following conditions exist:</p> <ul style="list-style-type: none">■ The incremental or the differential backup method was used.■ The destination was a tape device. <p>If you create full, differential, or incremental backups, GRT-enabled jobs have the following restrictions:</p> <ul style="list-style-type: none">■ The full, differential, and incremental job templates must be part of a policy.■ The destination device must be a backup-to-disk folder.■ The backup sets from the full, differential, and incremental jobs must be on the same volume.

Table 5-27 Granular Recovery Technology requirements (*continued*)

Agent	Restrictions
Agent for Microsoft Exchange Server with CPS	<p>GRT-enabled jobs have the following restrictions:</p> <ul style="list-style-type: none"> ■ Backups must be sent to a backup-to-disk folder on a local NTFS drive. <p>Note: You should use the backup-to-disk folder exclusively for CPS Exchange jobs. Do not back up other resources to the backup-to-disk folder that is the destination for the GRT-enabled backup job.</p> <ul style="list-style-type: none"> ■ Backups must be sent to a specific backup-to-disk folder. You cannot select a device pool. ■ Backups cannot be sent to a backup-to-disk folder for which you selected the Allocate the maximum size for backup-to-disk files option.
Agent for Microsoft SharePoint	<p>GRT-enabled jobs have the following restrictions:</p> <ul style="list-style-type: none"> ■ You can run only full backups for GRT-enabled jobs. ■ You must have a current version of the Remote Agent for Windows Systems installed on the SharePoint server.

Table 5-27 Granular Recovery Technology requirements (*continued*)

Agent	Restrictions
Agent for Microsoft Hyper-V Agent for VMware Virtual Infrastructure	<p>GRT-enabled jobs have the following restrictions:</p> <ul style="list-style-type: none"> ■ You can run only full backups for GRT-enabled jobs. ■ You can recover only individual items to virtual machines that run a Windows operating system. <p>By default, the Agent for Microsoft Hyper-V and the Agent for VMware Virtual Infrastructure use Granular Recovery Technology to protect files and folders at a granular level. You can also enable the granular recovery of Microsoft Exchange, SQL, and Active Directory application data that resides on virtual machines.</p> <p>See “How Granular Recovery Technology works with the Agent for Microsoft Hyper-V” on page 1177.</p> <p>See “How Granular Recovery Technology works with the Agent for VMware” on page 1396.</p>

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“Recommended devices for backups that use Granular Recovery Technology”](#) on page 324.

See [“Recommendations for using backup-to-disk folders with backup jobs that use Granular Recovery Technology”](#) on page 511.

Backing up data

This chapter includes the following topics:

- [How to back up data](#)
- [Creating a backup job by using the Backup Wizard](#)
- [Creating a backup job by setting job properties](#)
- [About scheduling jobs](#)
- [About the full backup method for backing up and deleting files](#)
- [About duplicating backed up data](#)
- [Verifying a backup](#)
- [About test run jobs](#)

How to back up data

Backups are crucial for data protection, and Backup Exec offers you many choices for creating backup jobs to protect your data, including the following:

- **Using the Backup Wizard.** Use this wizard to submit a backup job if you are a new or inexperienced Backup Exec user. The wizard guides you through the process of creating a backup job using most of the default options. After you become more experienced with Backup Exec, you will probably create backups by configuring backup job properties.
- **Configuring backup job properties.** Experienced Backup Exec users can create customized backup jobs by selecting resources to protect and setting backup options. Using the backup job properties pages allows you to set some options, such as job priority and database options, that cannot be set per job using the Backup Wizard.

- Creating a selection list. Select the data you want to back up and save the selections as a selection list. You can then choose the selection list when creating a backup job. You can use selection lists for multiple jobs. You can also choose a selection list and combine it with a policy to create a job.

Backup Exec allows you to set default options for backup jobs, but also gives you the flexibility to override these options for specific jobs. You can direct all backup jobs to a specified network segment, isolating the backup data so that other connected networks are not affected when backup operations are performed, or you can specify a LAN for a single job.

Backup Exec also provides the option of setting up backup jobs that take place on a routine basis (scheduled jobs), or set up one-time backup jobs.

In addition to creating backup jobs to protect data, you can create the following:

- A test run of a scheduled backup job to determine whether or not it is likely to complete successfully.
- A job that duplicates backup sets either from previously backed up data or data scheduled to be backed up. If the backup sets are to be duplicated from a scheduled job, the duplicate backup data job runs automatically after the backup job completes.
- Verify jobs to test the integrity of the media.
- Backup jobs that use the **Back up and delete the files** method to free disk space on the server.
- Resource discovery jobs to find new resources that may need to be backed up on a regular basis.

Before you begin backing up data, you should develop a backup strategy that includes the method, frequency, and media rotation methods that are appropriate for your organization. You may have different strategies for different areas of the organization. You should also ensure that you have the proper user rights to run back up jobs.

See [“Required user rights for backup jobs”](#) on page 331.

You may want to configure device and media management before creating backup jobs. You can set up Backup Exec to use specific storage devices or logical groupings of devices, such as device pools.

Specifically, you might want to perform the following tasks to help you manage storage hardware and media most effectively:

- Set up drive pools for systems with more than one storage device.
- Create media sets.

Caution: To protect remote resources, you must install the Backup Exec Remote Agent for Windows Systems on the remote computer. The Remote Agent is a system service that runs on Windows servers and workstations and provides efficient backup processing by locally performing tasks that, in typical backup technologies, require extensive network interaction.

See [“Creating device pools”](#) on page 516.

See [“About creating media sets”](#) on page 226.

See [“Creating a backup job by using the Backup Wizard”](#) on page 331.

See [“About backup strategies”](#) on page 270.

Required user rights for backup jobs

To perform any backup operations, the following Windows user rights are required for the service account and any Backup Exec logon accounts:

- Act as part of the operating system
- Create a token object.
- Back up files and directories.
- Restore files and directories.
- Managed auditing and security log.
- Logon as a batch job (only for Windows Vista and later).

For more information about user rights in Windows operating systems, see your Microsoft documentation.

See [“About the Backup Exec service account”](#) on page 106.

See [“About configuring logon accounts”](#) on page 184.

Creating a backup job by using the Backup Wizard

If you are new to Backup Exec or are uncertain about how to set up a backup job, you can use the Backup Wizard.

If you have experience with Backup Exec, you can create a backup job by setting the properties you want.

See [“Creating a backup job by setting job properties”](#) on page 332.

To create a backup job by using the Backup Wizard

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job using Wizard**.
- 3 Follow the on-screen prompts.

Preventing the Backup Wizard from launching from the Backup button

By default, the Backup Wizard displays when you select Backup on the navigation bar. If you prefer to set up backup jobs manually, you can prevent the Backup Wizard from displaying.

To prevent the Backup Wizard from launching from the Backup button

- 1 On the navigation bar, click **Backup**.
- 2 Uncheck **Always launch the Backup Wizard from the Backup button**.
- 3 Click **Cancel**.

Configuring the Backup Wizard to launch from the Backup button

By default, the Backup Wizard displays when you select Backup on the navigation bar. If you disable the Backup Wizard, you can re-enable it at any time.

To configure the Backup Wizard to launch from the **Backup** button, select **Re-enable the Backup Wizard from the Backup Button** on the **Tools** menu.

Creating a backup job by setting job properties

If you have experience with Backup Exec, you can create a backup job by setting the properties you want.

If you are new to Backup Exec or are uncertain about how to set up a backup job, you can use the Backup Wizard.

See [“Creating a backup job by using the Backup Wizard”](#) on page 331.

To create a backup job by setting job properties

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 From the backup selections pane, select the data you want to back up.
See [“Selections options for backup jobs”](#) on page 336.
- 4 In the **Properties** pane, under **Destination**, click **Device and Media**.

- 5 Select the device and media information for this job.
See [“Device and media options for backup jobs and templates”](#) on page 339.

6 Complete the following options as necessary:

To determine the order in which resources are processed	In the Properties pane, under Source , click Resource Order . See “Resource Order Backup options” on page 338.
To set or test credentials for the resource that is being processed	In the Properties pane, under Source , click Resource Credentials . See “Resource Credentials options” on page 337.
To determine the job's priority and availability	In the Properties pane, under Source , click Priority and Availability . See “Priority and Availability backup options” on page 308.
To configure selection list notification	In the Properties pane, under Source , click Selection List Notification . See “Notification options for jobs” on page 682.
To select a preferred server	In the Properties pane, under Source , click Preferred Servers . See “Preferred Servers backup options” on page 437.
To configure general job settings	In the Properties pane, under Settings , click General . See “General options for backup jobs and templates” on page 342.
To configure advanced job settings	In the Properties pane, under Settings , click Advanced . See “Advanced options for backup jobs” on page 348.
To configure network and security options	In the Properties pane, under Settings , click Network and Security . See “Network and Security backup options” on page 407.

To create pre/post commands	<p>In the Properties pane, under Settings, click Pre/Post Commands.</p> <p>See “Pre/post commands for backup or restore jobs” on page 352.</p>
To configure backup settings for an agent	<p>In the Properties pane, under Settings, select the name of the agent.</p> <p>See “Advanced Open File options ” on page 949.</p> <p>See “Backup options for the Advanced Disk-based Backup Option” on page 926.</p> <p>See “SQL backup options” on page 1273.</p> <p>See “Microsoft Exchange backup options” on page 1134.</p> <p>See “Microsoft SharePoint backup options ” on page 1246.</p> <p>See “Active Directory Recovery Agent backup job options” on page 887.</p> <p>See “Lotus Domino backup job options” on page 1076.</p> <p>See “Oracle backup options ” on page 1339.</p> <p>See “DB2 backup options” on page 966.</p> <p>See “NetWare SMS backup options” on page 1939.</p> <p>See “Backup job options for Linux, UNIX, and Macintosh computers” on page 1918.</p> <p>See “NDMP backup options for NetApp/IBM/Fujitsu” on page 1852.</p> <p>See “Enterprise Vault backup options” on page 992.</p> <p>See “VMware backup options” on page 1391.</p> <p>See “Microsoft Hyper-V backup options” on page 1175.</p>
To configure backup settings for Archiving Option Components	<p>In the Properties pane, under Settings, click Archive.</p> <p>See “Backup job properties for archive jobs” on page 1482.</p>

To configure Backup Exec to notify someone when a backup job containing a specific selection list completes

In the **Properties** pane, under **Settings**, click **Notification**.
 See [“Notification options for jobs”](#) on page 682.

7 Do one of the following:

To run the backup job now

Click **Run Now**.

To schedule the backup job for later

In the **Properties** pane, under **Frequency**, click **Schedule**.
 See [“Schedule options”](#) on page 356.

Selections options for backup jobs

When the **Backup Job Properties** dialog box appears, **Selections** is chosen by default in the **Properties** pane. Through the **Selections** options, you choose the data you want to include in the backup job.

See [“Creating a backup job by setting job properties”](#) on page 332.

This dialog box includes the following options:

Table 6-1 Selections options for backup job

Item	Description
Selection list name	Designates the name of the selection list. If you create a job using an existing selection list, you can select the selection list you want to use. Otherwise, you can use the default Selection list name, which creates a new selection list using this name.
Load selections from existing list	Allows you to use a previously created selection list or merge existing selection lists. See “Load Selections from Existing List options” on page 301.
Selection list description	Describes the selection list.
Include/Exclude	Allows you to use the Advanced File Selection option for selecting files for backing up. See “Backup Include/Exclude Selections options” on page 298.

Table 6-1 Selections options for backup job (*continued*)

Item	Description
Include subdirectories	Selects the contents of all the subfolders when a directory is selected.
Show file details	Displays details about the files that you can select.
View by Resource	Allows you to view selections as a list of resources.
View Selection Details	Allows you to view selections as a list of files and directories.

See [“About selecting data to back up”](#) on page 280.

See [“About selection lists”](#) on page 295.

See [“Creating selection lists”](#) on page 296.

See [“Adding a user-defined selection to the User-defined Selections node”](#) on page 290.

Resource Credentials options

A logon account enables Backup Exec to access resources for backup or restore jobs. You can change or test logon accounts before you run a job.

See [“About configuring logon accounts”](#) on page 184.

This dialog box includes the following options:

Table 6-2 Resource Credentials options

Item	Description
Resource	Specifies the resource for the job.
Logon Account	Specifies the logon account Backup Exec uses for this backup or restore selection.
Test Results	Details the results of the credentials test.
Test All	Tests all listed resource credentials to verify that they can access the resource.
Test Selected	Tests only the selected resource credentials to verify that Backup Exec can access the resource or resources.
Cancel Test	Cancels the credentials test.

Table 6-2 Resource Credentials options (*continued*)

Item	Description
Change	<p>Lets you change the selected resource credentials.</p> <p>For remote selections, do not change the logon account information. They rely on the logon account used to connect to the server they reside on, and will ignore the additional logon account you specify. This applies to drives, Lotus, System State, and Exchange selections (except mailboxes, which can and do use logon accounts).</p>
Clear	Removes the selected resource credentials from the dialog box.

Resource Order Backup options

After you make selections for a backup job, you can set up Backup Exec to process those selections in a certain order.

Please note the following about the order in which selections can be backed up:

- You can order resources within a server, but you cannot alternate selections across servers. For example, you can select C: and D: from Server A followed by selections from Server B. However, you cannot order selections as C: from Server A and then C: from Server B and then D: from both servers.
- For any given server, system state must be ordered last.

Table 6-3 Resource Order Backup options

Item	Description
Make First	Designates the selected resource as the first resource Backup Exec should process during the backup job.
Move Up	Moves the selected resource up in the resource order, meaning that Backup Exec processes it sooner during the backup job.
Move Down	Moves the selected resource down in the resource order, meaning that Backup Exec processes it later during the backup job.

Table 6-3 Resource Order Backup options (*continued*)

Item	Description
Make Last	Designates the selected resource as the last resource Backup Exec should process during the backup job.

Enter Password options

You can change your password on this dialog box.

Table 6-4 Enter Password options

Item	Description
Password	Designates your new password.
Confirm	Confirms your new password.

Device and media options for backup jobs and templates

You select the storage device and media set on which the backup job will run.

See [“Creating a backup job by setting job properties”](#) on page 332.

This dialog box includes the following options:

Table 6-5 Device and Media options for backup jobs and templates

Item	Description
Device	<p>Designates a device pool, a robotic library drive, a stand-alone drive, a backup-to-disk folder, a removable backup-to-disk folder, or other type of supported storage device to which you want to send backup data.</p> <p>See “About tape drives and robotic libraries” on page 451.</p> <p>See “About backup-to-disk folders ” on page 496.</p> <p>See “About device pools” on page 515.</p> <p>See “About the All Virtual Disks device pool in the Storage Provisioning Option” on page 2029.</p> <p>See “About the Remote Media Agent for Linux Servers” on page 1966.</p>

Table 6-5 Device and Media options for backup jobs and templates (*continued*)

Item	Description
<p>Use client-side deduplication</p>	<p>Enables a remote computer that is configured as a Remote Agent for Deduplication to deduplicate data and then send only the unique data directly to the deduplication device that is selected in the Device field. This option is selected by default when you select an OpenStorage device or a deduplication storage folder in the Device field.</p> <p>This option is available only if the Deduplication Option is installed and a deduplication device is selected in the Device field.</p> <p>See “About client-side deduplication” on page 1590.</p>
<p>Use media server deduplication</p>	<p>Enables the media server to deduplicate data. Media server deduplication reduces the size of backups, which reduces storage requirements. If the media server does not support deduplication, the data is deduplicated on an intelligent disk device, such as Symantec PureDisk or a device from a third-party vendor.</p> <p>This option is available only if the Deduplication Option is installed and a deduplication device is selected in the Device field.</p> <p>See “About the Deduplication Option” on page 1568.</p>
<p>Restrict backup of the selection list to devices on the following media server or media servers in a pool</p>	<p>Specifies if you want a job to run on devices on a specific managed media server or on devices that are on a group of managed media servers.</p> <p>This check box displays only if you have the Central Admin Server Option installed. This is an additional filter that lets you control where certain jobs are delegated. For example, to always run backups of Exchange databases only on the devices that are attached to managed media servers in a pool named Exchange Backups, select this option, and then select the Exchange Backups media server pool.</p>

Table 6-5 Device and Media options for backup jobs and templates *(continued)*

Item	Description
Media set	<p>Specifies the media set for the backup. If you select Overwrite, the media in the drive is overwritten if the media is scratch, or if its overwrite protection period has expired. If allocated or imported media are in the drive, they may also be overwritten depending on the Media Overwrite Protection Level that is set.</p> <p>If you selected one of the append options, the backup will be added to an appendable media (if one exists).</p>
Overwrite media	<p>Places this backup on an overwritable media. Make sure that appropriate media is in the stand-alone drive or drive pool you select in the Device field in this dialog box.</p> <p>The media in the drive is overwritten if the media is scratch or recyclable (its overwrite protection period has expired). If allocated or imported media are in the drive, they may also be overwritten depending on the Media Overwrite Protection Level that is set.</p> <p>Depending on your configuration, overwritable media is selected from scratch media or recyclable media.</p> <p>See “How Backup Exec searches for overwritable media” on page 233.</p> <p>If the media in the drive is not overwritable, an alert appears requesting that you insert overwritable media.</p>
Append to media, overwrite if no appendable media is available	<p>Appends this backup to the media set listed in the Media Set field in this dialog box. The backup set is appended if an appendable media is available in the selected media set; if not, an overwritable media is used and added to the media set.</p> <p>If an append job fills a media, the job continues on another piece of overwritable media.</p> <p>If the media in the drive is not overwritable, an alert appears requesting that you insert overwritable media.</p>
Append to media, terminate job if no appendable media is available	<p>Appends this backup to the media set listed in the Media Set field in this dialog box. The backup set is appended if an appendable media is available in the selected media set; if not, the job is terminated.</p>
Eject media after job completes	<p>Ejects the media in the drive when the operation completes.</p>

Table 6-5 Device and Media options for backup jobs and templates (*continued*)

Item	Description
Retension media before backup	Runs the tape in the drive from beginning to end at a fast speed, which helps the tape wind evenly and run more smoothly past the tape drive heads. Retensioning is primarily for Mini Cartridge and quarter-inch cartridges and is not supported on most other types of tape drives.
Use Write once, read many (WORM) media	Specifies the use of WORM (write once, read many) media for this backup job. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent. See “About WORM media” on page 247.
Maximum number of devices to use for resources that support multiple data streams	Specifies the number of devices that this backup job can use. Only one device per stream can be used.
Minimum number of devices, terminate job if fewer devices are available	Specifies the minimum number of devices that can be used for this backup job. If the minimum number of devices is not available, the job ends without completing.

General options for backup jobs and templates

You can configure general options for backup jobs, including the name of the job and the backup method to be used.

See [“Creating a backup job by setting job properties”](#) on page 332.

This dialog box includes the following options:

Table 6-6 General settings options

Item	Description
Job name/Template name	Designates the name for this backup job or template. You can accept the default name that appears or enter a name. The name must be unique.
Backup set description	Describes the information in the backup set for future reference.

Table 6-6 General settings options (*continued*)

Item	Description
Backup method for files	

Table 6-6 General settings options (*continued*)

Item	Description
	<p>Designates one of the following backup methods:</p> <ul style="list-style-type: none"> <p>■ Full - Back up files</p> <ul style="list-style-type: none"> - Using archive bit (reset archive bit). Includes all of the files selected for backup and resets the archive bit to indicate that the files have been backed up. - Using modified time. Includes all of the files selected for backup and allows the use of incrementals and differentials using the modified date and time stamp. - Copy the files. Includes all selected data but does not reset the archive bit. It does not affect your backup strategy or media rotation scheme. - Back up and delete the files (delete selected files and folders after successful copy backup). Backs up the selected data, verifies the media, and then deletes the data from the volume. The logon account credentials that you use to run the job must have the rights to delete a file. To use the method to back up and delete the files on computers on which the Remote Agent for Linux or UNIX Servers or the Remote Agent for Macintosh Systems is installed, the Backup Exec logon account must have superuser privileges. Otherwise, the data is backed up, but is not deleted. The Backup Exec Archive Option offers more features for data archiving. See “About the Archiving Option” on page 1414. <p>■ Differential - Back up changed files since last full</p> <ul style="list-style-type: none"> - Using archive bit (does not reset archive bit). Includes all files that changed (based on the archive bit) since the last full backup. It does not affect your backup strategy or media rotation scheme because the archive bit is not reset. - Using modified time. Includes all files changed since the last full backup, using the files' last modified date and time stamp. Make sure that the same script or selection list is used for the differential backup that was used for the full backup. <p>Note: A file's last modified date and timestamp</p>

Table 6-6 General settings options (*continued*)

Item	Description
	<p>does not change when the file is copied or moved. If the file's modified time is older than the previous backup's modified time, that file is not backed up. To ensure that the files are protected, run a full backup after you copy or you move files. If you have the Advanced Disk-based Option, you can run synthetic backups to ensure that any copied or moved files are protected.</p> <ul style="list-style-type: none"> ■ Incremental - Back up changed files since last full or incremental <ul style="list-style-type: none"> - Using archive bit (reset archive bit). Includes only the files that have changed (based on the archive bit) since the last full or incremental backup and resets the archive bit to indicate that the files have been backed up. - Using modified time. Includes all files that have changed since the last full or incremental backup, using the files' last modified date and time stamp. Make sure that the same script or selection list is used for the incremental backup that was used for the full backup. <p>Note: A file's last modified date and timestamp does not change when the file is copied or moved. If the file's modified time is older than the previous backup's modified time, that file is not backed up. To ensure that the files are protected, run a full backup after you copy or you move files. If you have the Advanced Disk-based Option, you can run synthetic backups to ensure that any copied or moved files are protected.</p> ■ Working Set - Back up files <ul style="list-style-type: none"> - Changed today. Backs up all files that were created or modified today. - Last accessed in (x) days. If you select this backup method, you can then indicate in the Files accessed in x days field that you want to include data that has been accessed in a specific number of days. See "About backup methods" on page 274.

Table 6-6 General settings options (*continued*)

Item	Description
<p>Files accessed in x days</p>	<p>Specifies the number of days for which to include accessed files if you selected Last accessed in (x) days in the Backup method for files field.</p> <p>Symantec recommends that you specify at least 30 days in order to include the data needed to make your system operational if you have to restore a working set backup.</p>
<p>Use the Microsoft Change Journal if available</p>	<p>Uses the Microsoft Change Journal to determine which files have been modified since the last full backup. This option can only be used with NTFS volumes.</p> <p>This option is available when you select one of the following backup methods:</p> <ul style="list-style-type: none"> ■ Full - Back Up Files - Using modified time. This method is not available when performing offhost backup. ■ Differential - Back up changed files since last full - Using modified time ■ Incremental - Back up changed files since last full or incremental - Using modified time. <p>In addition, this option becomes available if you select the Collect additional information for synthetic backup and for true image restore check box.</p> <p>If you use the Change Journal with the option to collect additional information for synthetic backup and for true image restores, the archive bit is not reset, even if you selected a backup method that has "reset archive bit" in the name.</p> <p>If you are backing up volumes with junction points that were created by linkd.exe, you should not use the Microsoft Change Journal. Junction points are not followed properly in this situation.</p>
<p>Preserve tree on back up and delete</p>	<p>Retains the directory structure on the hard drive of the files that are backed up in a full backup job. This option is available only when you select the full backup method that backs up and deletes the files.</p>

Table 6-6 General settings options (*continued*)

Item	Description
Collect additional information for synthetic backup and for true image restore	<p>Displays only for templates. It is used with synthetic backup jobs and true image restore jobs. It specifies that Backup Exec collects the information required to detect files and directories that have been moved, renamed, or newly installed since the last backup, and then includes those files and directories in the backup jobs. If this option is not selected, Backup Exec skips these files and directories if their archive bits are unchanged. With this option selected, Backup Exec compares path names, file names, modified times, and other attributes with those from the previous full and incremental backups. If any of these attributes are new or changed, then the file or directory is backed up.</p> <p>Backups that have this option selected require more disk space, and take more time to run, than backups that do not.</p> <p>You must select this option for the baseline and incremental backup template in a synthetic backup policy.</p> <p>See “About the synthetic backup feature” on page 899.</p>
Verify after backup completes	<p>Performs a verify operation automatically to make sure the media can be read once the backup has been completed. Verifying all backups is recommended.</p>

Table 6-6 General settings options (*continued*)

Item	Description
Compression type	<p>Provides the following compression options:</p> <ul style="list-style-type: none"> ■ None. This option copies the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage media space. Hardware data compression should not be used in environments where devices that support hardware compression are used interchangeably with devices that do not have that functionality. In this situation, hardware compression is automatically disabled. You can manually re-enable hardware compression on the drives that support it, but this results in media inconsistency. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive. ■ Software. This option uses STAC software data compression, which compresses the data before it is sent to the storage device. ■ Hardware [if available, otherwise none]. This option uses hardware data compression (if the storage device supports it). If the drive does not feature data compression the data is backed up uncompressed. ■ Hardware [if available, otherwise software]. This option uses hardware data compression (if the storage device supports it). If the drive does not feature hardware data compression, STAC software compression is used.

See [“Creating a backup job by setting job properties”](#) on page 332.

Advanced options for backup jobs

You can customize your backup job with advanced options.

See [“Creating a backup job by setting job properties”](#) on page 332.

This dialog box includes the following options:

Table 6-7 Advanced options for backup job

Item	Description
Enable single instance backup for NTFS volumes	<p>Displays only if you use the Microsoft Windows Single Instance Store (SIS) feature. Single instance backup checks the NTFS volume for identical files. If Backup Exec finds multiple copies of a file, it only backs up one instance of that file, regardless of how many SIS links reference it.</p> <p>Single instance backup can considerably reduce the storage space that is required for your backups. Many applications automatically generate files that have identical content. The actual amount of space you save depends on the number of duplicate files on the volume.</p> <p>Warning: If the backup job does not run to completion, the file data may not be included in the backup set. Rerun the backup until it is successfully completed. If the incremental backup method was used, running the job again will not back up the same files. You must run a full or copy backup to ensure that all files are backed up completely. If the 'incremental - using modified time' backup method was used, running the same backup job to completion will back up the files correctly.</p>

Table 6-7 Advanced options for backup job (*continued*)

Item	Description
<p>Back up files and directories by following junction points</p>	<p>Backs up the information for the junction points and the files and directories to which they are linked. If this check box is not selected, then only the information for the junction points is backed up; the files and directories to which they are linked are not backed up.</p> <p>Backup Exec does not follow junction points automatically created by Microsoft Windows Vista/Server 2008 because it can cause the data to be backed up repeatedly.</p> <p>For more information, see the following Symantec Knowledge Base article: http://entsupport.symantec.com/umi/V-269-9</p> <p>Since Mounted Drives that do not have a drive letter assigned to them cannot be selected, the files and directories to which they are linked are backed up regardless of whether this option is selected.</p> <p>If this option is selected and the actual files and directories to which the junction points are linked are also included in the backup selections, then the files and directories are backed up twice; once during the full file and directory backup, and again via the junction point.</p> <p>Warning: If a junction point is linked to a location that encompasses it, then recursion (a situation where data is backed up repeatedly) will occur, resulting in an error and job failure. For example, if c:\junctionpoint is linked to c:\, recursion will occur when attempting to back up c:\junctionpoint, and the backup job will fail.</p>
<p>Back up files and directories by following symbolic links</p>	<p>Backs up the information for the symbolic links and the files and directories to which they are linked.</p> <p>If you do not select this option, only the information for the symbolic links is backed up. The files and directories to which they are linked are not backed up.</p> <p>If the symbolic link points to files and directories on a remote computer, the files and directories on the remote computer are not backed up.</p>

Table 6-7 Advanced options for backup job (*continued*)

Item	Description
Back up data in Remote Storage	<p>Backs up data that has been migrated from primary storage to secondary storage. The data is not recalled to its original location; it is backed up directly to the backup media.</p> <p>If this option is selected, you should not run a backup of your entire system because Backup Exec has to load the data that has been migrated to secondary storage and additional time is required for any set that includes migrated data.</p> <p>If this check box is cleared, only the placeholder that stores the location of the data on secondary storage will be backed up, not the data itself.</p> <p>This option should not be selected if the device used for secondary storage and backups contains only one drive because Remote Storage and Backup Exec will compete for use of the drive.</p>
Set Remote Agent priority	<p>Allows you to select the number of CPU cycles the media server will use to maintain optimal server performance while Remote Agent backups are running. The higher the priority, the more the protected server's CPU processing power is used during backup operations.</p> <p>Allocating fewer CPU cycles to a backup job may result in slower backup performance.</p> <p>This field contains the following options:</p> <ul style="list-style-type: none"> ■ Normal Priority. Select this option to allocate the default number of CPU cycles the protected server will use during a Remote Agent backup. ■ Below Normal Priority. Select this option to allocate fewer server CPU cycles to the backup job. ■ Lowest Priority. Select this option to allocate the fewest number of CPU cycles to the backup job.
Never	<p>Skips open files if they are encountered during the backup operation. A listing of skipped files appears in the job log for the backup.</p>

Table 6-7 Advanced options for backup job (*continued*)

Item	Description
If closed within x seconds	<p>Waits the specified time interval for files to close before skipping the open file and continuing the backup operation.</p> <p>If the file does not close during the specified interval, it is skipped. A listing of skipped files appears in the job log for the backup.</p> <p>If multiple files are open, Backup Exec waits the specified time interval for each file; depending on the number of open files, this may significantly increase the backup time.</p>
With a lock	<p>Attempts to open files that are in use. If Backup Exec is able to open a file, the file is locked while it is being backed up to prevent other processes from writing to it. Backing up open files is not as effective as closing applications and allowing the files to be backed up in a consistent state.</p>
Without a lock	<p>Attempts to open files that are in use. If Backup Exec is able to open the file, the file is NOT locked while it is being backed up. This allows other applications to write data to the file during the backup operation.</p> <p>Warning: This option allows files that contain inconsistent data and possibly corrupt data to be backed up.</p>

To back up the Removable Storage database in the \Ntmsdata subdirectory, the WMI repository in the wbem\Repository subdirectory, and the Terminal Services database in the default \LServer subdirectory, select the <Systemroot>\System32 directory. Files that you place in the Systemroot\System32\Ntmsdata subdirectory, the \wbem\Repository subdirectory, or the default \LServer subdirectory may not be backed up; only system files are included in the backup. It is recommended that you do not place user files in the Systemroot\System32 directory or subdirectories.

See [“Creating a backup job by setting job properties”](#) on page 332.

See [“Setting default options for the Advanced Open File Option”](#) on page 943.

Pre/post commands for backup or restore jobs

You can use pre/post commands to run commands before or after a job.

See [“About pre/post commands”](#) on page 399.

See [“Setting default pre/post commands”](#) on page 400.

This dialog box includes the following options:

Table 6-8 Pre- and post-command options

Item	Description
Pre-command	<p>Runs a command on the specified server before the backup or restore job is run. Use local paths, and make sure the paths exist on each server and are correct.</p> <p>Commands that require user interaction, such as prompts, are not supported.</p>
Post-command	<p>Runs a command on the specified server after the backup or restore job has run. Use local paths, and make sure the paths exist on each server and are correct.</p> <p>Commands that require user interaction, such as prompts, are not supported.</p>
Allow pre- and post-commands to be successful only if completed with a return code of zero	<p>Allows Backup Exec to check the return codes of the pre- and post-commands to determine if they completed successfully.</p> <p>An exit code of zero returned to the operating system by the pre- or post-command is interpreted by Backup Exec to mean that the command completed successfully. A non-zero exit code is interpreted by Backup Exec to mean the command ended with an error.</p> <p>After checking the return codes, Backup Exec continues processing according to selections you made for running the pre- and post-commands.</p> <p>If this option is not selected, the success of the pre- and post-commands is not determined based on the return code.</p>
Run job only if pre-command is successful	<p>Runs the backup or restore job only if the pre-command is successful. If the pre-command fails, the job does not run, and is marked as failed.</p> <p>If it is critical that the job does not run if the pre-command fails, then select Allow pre- and post-commands to be successful only if completed with a return code of zero. If a non-zero code is returned, it is interpreted by Backup Exec to mean that the pre-command did not run successfully. The job is not run and the job status is marked as Failed.</p>

Table 6-8 Pre- and post-command options (*continued*)

Item	Description
<p>Run post-command only if pre-command is successful</p>	<p>Runs the post-command only if the pre-command is successful.</p> <p>If it is critical that the post-command does not run if the pre-command fails, then select Allow pre- and post-commands to be successful only if completed with a return code of zero. If a non-zero code is returned for the pre-command, it is interpreted by Backup Exec to mean that the pre-command did not run successfully. The post-command does not run.</p> <p>If you also select Run job only if pre-command is successful, and both the pre-command and the job are successful, but the post-command returns a non-zero code, the job log reports both the job and the post-command as failed.</p>
<p>Run post-command even if job fails</p>	<p>Runs the post-command regardless of whether the job is successful or not.</p> <p>If you also select Allow pre- and post-commands to be successful only if completed with a return code of zero and the post-command returns a non-zero code, the job log reports the post-command as failed.</p>
<p>Run post-command after job verification completes</p>	<p>Runs the post-command after the verification completes if you selected the Verify after backup completes option on the General backup properties dialog box.</p>
<p>Cancel command if not completed within x minutes</p>	<p>Designates the number of minutes Backup Exec should wait before canceling a pre- or post-command that did not complete. The default time-out is 30 minutes.</p>
<p>On this media server</p>	<p>Runs the pre- and post-commands on this media server only.</p>
<p>On each server backed up</p>	<p>Runs the pre- and post-commands one time on each server backed up.</p> <p>The pre- and post-command selections apply to each server independently. If you select this option, the pre- and post-commands are run and completed for each server before processing begins on the next selected server.</p>

Job Summary properties

The **Job Summary** displays when you create a job. You should review the details to ensure that the job properties are accurate.

See [“Creating a backup job by setting job properties”](#) on page 332.

See [“Restoring data by setting job properties”](#) on page 605.

See [“About creating utility jobs to help manage devices and media”](#) on page 480.

Table 6-9 Job Summary properties

Item	Description
Do not display this summary again	Turns off the job summary so that it does not appear when you create jobs.
OK	Finalizes the creation of the job. The job runs as scheduled.
Cancel	Closes the Job Summary so that you can change the job settings before submitting it.
Print	Prints the job summary.

How to include or exclude files for backup

Advanced file selection allows you to quickly select or de-select files for backup operations by specifying file attributes.

See [“Creating a backup job by setting job properties”](#) on page 332.

See [“Backup Include/Exclude Selections options”](#) on page 298.

With this feature you can do the following:

- Include or exclude files by filename attributes. For example, you can select only files with .txt extensions, or exclude files with .exe extensions from a backup. If you exclude files by an attribute that does not exist, all files of that type are excluded. For example, excludes based on SQL database dates result in global SQL excludes since SQL databases do not have date attributes.
- Select only files that fall within a specified date range. For example, you can select files that were created or modified during the month of December.
- Specify the files that have not been accessed in a specified number of days. For example, you can select the files that have not been accessed in 30 days from your "My Documents" folder. Then, run a full backup job for which you select the method to back up and delete the files.

The Backup Exec Archive Option offers more features for data archiving.

See [“About the Archiving Option”](#) on page 1414.

About scheduling jobs

The schedule option enables you to configure the time and the frequency that you want to run jobs. You can configure a schedule for jobs such as backup, restore, inventory, and new catalog. During the job setup, you can choose to run jobs immediately, run once on a specific day and time, or run according to a schedule.

See [“Scheduling jobs”](#) on page 356.

When you create a backup selection list, you can set a time range when the resources in the list will be available for backup. The time range is called the availability window. If you schedule a job to run outside of the availability window, the job does not run and Backup Exec displays an Invalid Schedule status for the job on the Job Monitor. When scheduling a job, be sure that the schedule is within the availability window for the resources.

See [“Setting priority and availability windows for selection lists”](#) on page 307.

See [“Configuring default schedule options”](#) on page 366.

See [“Excluding dates from a schedule”](#) on page 366.

Scheduling jobs

The schedule option enables you to configure the time and the frequency that you want to run jobs. You can configure a schedule for jobs such as backup, restore, inventory, and new catalog. During the job setup, you can choose to run jobs immediately, run once on a specific day and time, or run according to a schedule.

See [“About scheduling jobs”](#) on page 356.

To schedule a job

- 1 Determine the type of job you want to schedule, and then on the navigation bar, click the appropriate button. For example, to schedule a backup job, click the arrow next to Backup.
- 2 In the **Properties** pane, under **Frequency**, click **Schedule**.
- 3 Select the appropriate options.
See [“Schedule options”](#) on page 356.
- 4 Click **Run Now**.

Schedule options

The following table lists the options that you can select for scheduling jobs.

See “[Scheduling jobs](#)” on page 356.

Table 6-10 Schedule options

Item	Description
Current date and time	Displays the current date and time that is set on this computer
Run now	Runs the job immediately.
Run on Date at Time	Lets you schedule the job to run one time on the selected date at the selected time.
Run according to schedule	Lets you configure a schedule for a recurring job.
Edit Schedule Details	Lets you select the run days for a recurring job.
Effective date	Displays the day that the schedule begins if Run according to schedule is selected.

Table 6-10 Schedule options (*continued*)

Item	Description
<p>Time Window</p>	<p>Displays the specified period of time during which a job can begin on any scheduled day if Run according to schedule is selected.</p> <p>When setting up the time during which a task runs, you can enter a time window that extends past midnight and into the next day. Bear in mind, however, that this may change the days on which the task runs. For example, if you schedule a task to run every Friday between 8:00 PM and 4:00 AM, it's possible for the task to run on Saturday morning sometime before or at 4:00 AM. If you don't want the task to run on Saturday, you must alter the time window, for example, by changing the ending value from 4:00 AM to 11:59:59 PM. Then, the task is confined to one day. When a time window crosses midnight, the start time is later in the day than the end time.</p> <p>Backup Exec considers both the job's time window and the resource's availability window when it runs a job. If you schedule a job to run outside of the availability window, it does not run. Backup Exec displays an Invalid Schedule status for the job on the Job Monitor. When you schedule a job, be sure that the job's time window is within the availability window for the resources.</p> <p>See "Priority and Availability backup options" on page 308.</p>
<p>Submit job on hold</p>	<p>Lets you submit the job with an on-hold status.</p> <p>You should select this option if you want to submit the job, but do not want the job to run until you change the job's hold status.</p>
<p>Delete the job if the job successfully completes</p>	<p>Deletes jobs that complete successfully, have been created to run once, either now or at a scheduled time, and have not been created from a template.</p>

Table 6-10 Schedule options (*continued*)

Item	Description
Delete the job after the job completes	Deletes any jobs that complete, even with errors, have been created to run once, and have not been created from a template. Jobs that are created to run once are deleted whether they run immediately or are scheduled for a later time.
Do not delete the job	Keeps the jobs that were created to run once, and were not created from a template, in the Job Setup view. This option is selected by default.

About the scheduling calendar

Both the **Calendar Schedule** tab and the **Exclude Dates** tab display a three-month calendar. The calendar provides a way to select days on which you want jobs to run and a way to view a summary of your schedule.

When you select a day to run a job, a green check mark displays on the calendar. In addition, when you are viewing a calendar for one type of schedule option, gray check marks indicate that other types of schedule options are set for those days.

See [“Scheduling a job to run on specific days”](#) on page 359.

See [“Scheduling a job to run on recurring week days”](#) on page 360.

See [“Scheduling a job to run on recurring days of the month”](#) on page 361.

See [“Scheduling a job to run on a day interval”](#) on page 362.

See [“Setting the effective date for a job schedule”](#) on page 363.

See [“Setting the time window for a scheduled job”](#) on page 364.

See [“Restarting a job during a time interval”](#) on page 365.

See [“Excluding dates from a schedule”](#) on page 366.

Scheduling a job to run on specific days

You can schedule a job to run on a single day or on multiple days.

To schedule a job to run on specific days

- 1 Determine the type of job that you want to schedule, and then on the navigation bar, click the appropriate button.
For example, to schedule a backup job, click the arrow next to **Backup**.
- 2 In the **Properties** pane, under **Frequency**, click **Schedule**.
- 3 Click **Run according to schedule**.
- 4 Do one of the following:

To select run days for a new job Continue to step 5.

To edit run days for an existing job Click **Edit Schedule Details**.

- 5 On the **Calendar Schedule** tab, under **Edit Calendar schedule by**, click **Specific Dates**.

- 6 Do one of the following:

- To select a single date
- Click **New**.
 - Enter the date.
 - Click **OK**.

To select multiple dates Click the dates on the calendar.

- 7 Click **OK**.

Scheduling a job to run on recurring week days

Use the recurring week days option to run jobs on the following types of schedules:

Table 6-11 Recurring schedule examples

Recurring job option	Example
The same day of the week, every week of the month	For example, every Wednesday.
Every day of the same week of every month	For example, every day during the second week of the month.
On selected days during selected weeks of the month	For example, the last Friday of every month.

To schedule a job to run on recurring week days

- 1 Determine the type of job that you want to schedule, and then on the navigation bar, click the appropriate button.

For example, to schedule a backup job, click the arrow next to **Backup**.

- 2 In the **Properties** pane, under **Frequency**, click **Schedule**.
- 3 Click **Run according to schedule**.

- 4 Do one of the following:

To select run days for a new job Continue to step 5.

To edit run days for an existing job Click **Edit Schedule Details**.

- 5 On the **Calendar Schedule** tab, under **Edit Calendar schedule by**, click **Recurring Week Days**.

- 6 Do one of the following:

To run a job on a single day Check the check box for that specific day.

To run a job on the same day every week Select the name of the day in the matrix. For example, to run a job every Monday, click **Mon**

To run a job every day of an entire week Select the row number for that week. For example, to select the first week of each month, click **1st**.

To run a job the last week of the month, regardless of the number of weeks in a month Click **Last**.

To run a job on all days of the month Click **Select All**.

To clear all existing selections Click **Deselect All**.

- 7 Click **OK**.

Scheduling a job to run on recurring days of the month

You can schedule jobs to run on specific days of the month, on the last day of the month, or on all days of the month.

To schedule a job to run on recurring days of the month

- 1 Determine the type of job that you want to schedule, and then on the navigation bar, click the appropriate button.
For example, to schedule a backup job, click the arrow next to **Backup**.
- 2 In the **Properties** pane, under **Frequency**, click **Schedule**.
- 3 Click **Run according to schedule**.
- 4 Do one of the following:

To select run days for a new job	Continue to step 5.
To edit run days for an existing job	Click Edit Schedule Details .
- 5 On the **Calendar Schedule** tab, under **Edit Calendar schedule by**, click **Recurring Days of the Month**.
- 6 Do one of the following:

To run jobs on specific days of the month	Click the button for each day.
To run jobs on the last day of the month, regardless of the actual date	Check Last Day .
To run a job on all days of the month	Click Select All .
To clear all existing selections	Click Deselect All .
- 7 Click **OK**.

Scheduling a job to run on a day interval

You can schedule a job to run every certain number of days calculated from a particular date. For example, you can set up a job to run every three days, starting on January 1, 2006. By default, the date from which the interval is calculated is the current date. However, you can set a date on which you want the schedule to go into effect. If the selection list that you are backing up has an availability window, Backup Exec uses the availability window instead of the date that you select here to calculate the starting date.

For example, you schedule a backup job to run every 7 days beginning on the 11th day of June. The associated selection list has an availability window that begins on the 12th day of June. The job is scheduled to run for the first time on June 12th. However, the calendar indicates that the starting date is June 11th.

Recurring tasks run during the specified time window.

To schedule a job to run on a day interval

- 1 Determine the type of job that you want to schedule, and then on the navigation bar, click the appropriate button.

For example, to schedule a backup job, click the arrow next to **Backup**.

- 2 In the **Properties** pane, under **Frequency**, click **Schedule**.

- 3 Click **Run according to schedule**.

- 4 Do one of the following:

To select run days for a new job Continue to step 5.

To edit run days for an existing job Click **Edit Schedule Details**.

- 5 On the **Calendar Schedule** tab, under **Edit Calendar schedule by**, click **Day Interval**

- 6 Check **Every**.

- 7 Enter the number of days on which you want the job to recur.

- 8 In the days calculated from box, select the date on which you want the schedule to go into effect.

The date you enter here does not override the effective date that you set up using the **Effective Date** option on the **Calendar Schedule** tab.

- 9 Click **OK**.

Setting the effective date for a job schedule

The effective date determines when your schedule goes into effect. A job cannot run prior to its effective date. By default, the effective date is the current date.

To set the effective date for a job schedule

- 1 Determine the type of job that you want to schedule, and then on the navigation bar, click the appropriate button.

For example, to schedule a backup job, click the arrow next to **Backup**.

- 2 In the **Properties** pane, under **Frequency**, click **Schedule**.

- 3 Click **Run according to schedule**.

- 4 Do one of the following:
 - To select run days for a new job Continue to step 5.
 - To edit run days for an existing job Click **Edit Schedule Details**.
- 5 On the **Calendar Schedule** tab, under **Edit Calendar schedule by**, click **Effective Date**.
- 6 Verify that the **Make the schedule go into effect on** check box is checked.
- 7 Select the date on which you want the schedule to go into effect.
- 8 Click **OK**.

About time windows

The time window is the period of time during which a job can begin. The time window does not extend beyond 23 hours, 59 minutes, and 59 seconds. For example, you cannot set a time window to start at 03:00 AM and end at 05:00 AM the next day.

The default time window is from 11:00 PM to 10:59:59 PM. If you use the default setting, a job that is scheduled to run on Monday can begin at or anytime after 11:00 PM on Monday. It cannot start after 10:59:59 PM on Tuesday night.

You can set a time window that extends past midnight and into the next day, which may change the day on which the job runs. For example, if you schedule a job to run every Friday between 10:00 PM and 04:00 AM, the job may run on Saturday before or at 04:00 AM. If you do not want the job to run on Saturday, you must alter the time window to start the job no later than 11:59:59.

See [“Setting the time window for a scheduled job”](#) on page 364.

Setting the time window for a scheduled job

You can set a time window to establish a period of time during which a job can begin.

See [“About time windows”](#) on page 364.

To set the time window for a scheduled job

- 1 Determine the type of job that you want to schedule, and then on the navigation bar, click the appropriate button.
 - For example, to schedule a backup job, click the arrow next to **Backup**.
- 2 In the **Properties** pane, under **Frequency**, click **Schedule**.

- 3 Click **Run according to schedule**.
- 4 Do one of the following:
 - To select run days for a new job Continue to step 5.
 - To edit run days for an existing job Click **Edit Schedule Details**.
- 5 On the **Calendar Schedule** tab, under **Edit Calendar schedule by**, click **Time Window**.
- 6 In the **Start no earlier than** box, select the time after which the job can start.
- 7 In the **and no later than** box, select the time by which the job must start.
- 8 Click **OK**.

Restarting a job during a time interval

You can set up a job to run multiple times on the scheduled run day during the specified time interval. You specify the interval at which the job repeats during the time window. For example, if there is a 12-hour time window for a job, you can set the job to run every two hours during that time window. The job runs at the interval you specify, relative to the start time of your time window. The interval must be greater than zero, and less than 23:59:59. In addition, the restart interval must be less than the amount of time set for the time window. For example, if you have a two-hour time window, you cannot specify a restart interval greater than 01:59:59.

To restart a job during a time interval

- 1 Determine the type of job that you want to schedule, and then on the navigation bar, click the appropriate button.
 - For example, to schedule a backup job, click the arrow next to **Backup**.
- 2 In the **Properties** pane, under **Frequency**, click **Schedule**.
- 3 Click **Run according to schedule**.
- 4 Do one of the following:
 - To select run days for a new job Continue to step 5.
 - To edit run days for an existing job Click **Edit Schedule Details**.
- 5 On the **Calendar Schedule** tab, under **Edit Calendar schedule by**, click **Restart Time Interval**.

- 6 Check **Restart task every**.
- 7 Select the time interval in hours, minutes, and seconds.
- 8 Click **OK**.

Excluding dates from a schedule

You can exclude specific dates, such as holidays, from a schedule. When you select a date, the symbol on that date on the calendar changes to a red circle with a line through it.

To exclude dates from a schedule

- 1 Determine the type of job that you want to schedule, and then on the navigation bar, click the appropriate button.
For example, to schedule a backup job, click the arrow next to **Backup**.
- 2 In the **Properties** pane, under **Frequency**, click **Schedule**.
- 3 Click **Run according to schedule**.
- 4 Do one of the following:

To select run days for a new job	Continue to step 5.
To edit run days for an existing job	Click Edit Schedule Details .
- 5 Click the **Exclude Dates** tab.
- 6 Do one of the following:

To exclude a single date	<ul style="list-style-type: none">■ Click New.■ Enter the date.■ Click OK.
To exclude multiple dates	Click the dates on the calendar.
To add an excluded date back into the schedule	<ul style="list-style-type: none">■ In the Exclude Dates box, click the date.■ Click Delete.
- 7 Click **OK**.

Configuring default schedule options

You can configure default scheduling parameters for all new jobs that you create. If you want to keep a static schedule for all new jobs that you run according to a

schedule, you can set a default schedule for all jobs, and then use the Run according to schedule option during job setup to make changes, if necessary.

To configure default schedule options:

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Schedule**.
- 3 Select the appropriate options.
 See “[Default schedule options](#)” on page 367.
- 4 Click **OK**.

Default schedule options

You can configure default scheduling parameters for all new jobs that you create. See “[Configuring default schedule options](#)” on page 366.

Table 6-12 Default schedule options

Item	Description
Edit Schedule Details	Lets you set or change the existing default schedule options.
Delete the job if the job successfully completes	Deletes jobs that complete successfully, have been created to run once, either now or at a scheduled time, and have not been created from a template.
Delete the job after the job completes	Deletes any jobs that complete, even with errors, have been created to run once, and have not been created from a template. Jobs that are created to run once are deleted whether they run immediately or are scheduled for a later time.
Do not delete the job	Preserves jobs that were created to run once, and were not created from a template, in the Job Setup view.

About the full backup method for backing up and deleting files

When you run a full backup, you can select the method to back up and delete the files. This backup method lets you free disk space on your server volume by moving

files and folders from the server to media. Backup Exec backs up the selected data as a copy backup, verifies the media, and then deletes the data from the volume.

The credentials in the Backup Exec logon account that you use to run the job must have the rights to delete a file. To use the method to back up and delete the files on computers on which the Remote Agent for Linux or UNIX Servers or the Remote Agent for Macintosh Systems is installed, the Backup Exec logon account must have superuser privileges. Otherwise, the data is backed up, but is not deleted.

Backup Exec performs a verify operation after the data is backed up. If the verify operation fails, the job stops and you are notified. If you get a verification failure, view the job log. Try to correct the problem, and then retry the job. After the data is backed up and verified, Backup Exec deletes the selected data. The job log contains a list of the data that is deleted.

You can enable the checkpoint restart option for a full backup job that uses the method to back up and delete the files. If a cluster failover occurs and the job is resumed, the files are not deleted from the source volume after the backup completes.

The Backup Exec Archive Option offers more features for data archiving.

See [“About the Archiving Option”](#) on page 1414.

See [“Backing up and deleting files”](#) on page 368.

Backing up and deleting files

When you run a full backup, you can select the method to back up and delete the files. Backup Exec backs up the selected data as a copy backup, verifies the media, and then deletes the data from the volume.

See [“About the full backup method for backing up and deleting files”](#) on page 367.

The Backup Exec Archive Option offers more features for data archiving.

See [“About the Archiving Option”](#) on page 1414.

To back up and delete files

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 Select the data that you want to back up and delete.
- 4 Click **General**.

- 5 In the **Backup method for files** field, select **Back up and delete the files (delete selected files and folders after successful copy backup)**.
- 6 Complete the backup job options.
See [“Creating a backup job by setting job properties”](#) on page 332.

About duplicating backed up data

You can create a job to duplicate backup data, selecting either to duplicate existing backup sets or to duplicate backup sets immediately following a scheduled job.

You can use a duplicate backup job to copy data directly from a virtual device to a physical device. Software encryption cannot be applied to a duplicate backup job when you copy data directly from a virtual device to a physical device. You must either disable DirectCopy or select not to encrypt the job.

See [“How to copy data directly from a virtual tape library to a physical tape device”](#) on page 381.

If you select to duplicate existing backup sets, the backup sets you select from catalogs are read from the source media and written to the selected destination, such as a drive, drive pool, or backup folder. You can schedule when this type of job runs.

If you duplicate Oracle or DB2 backup sets that were created with multiple data streams, note the following:

- Backup Exec converts the multiple data streams to a sequential data stream during the duplication job.
- A restore job from the duplicated copy may be slower than a restore job from the original media.

If you select to duplicate backup sets following a job, you select a scheduled backup job as the source. That backup job runs first, and then the backup sets it created are copied to the destination you selected for the duplicate job. To duplicate backup sets following a job, the backup job must be scheduled to run and must not be associated with any other duplicate jobs. You cannot schedule this job; instead, the duplicate job runs only after the related, or linked, backup job completes.

See [“Duplicating backed up data”](#) on page 369.

Duplicating backed up data

You can create a job to duplicate backup data, selecting either to duplicate existing backup sets or to duplicate backup sets immediately following a scheduled job.

See [“About duplicating backed up data”](#) on page 369.

To duplicate backup data

- 1 From the navigation bar, click **Job Setup**.
- 2 Under **Backup Tasks**, select **New job to duplicate backup sets**.
- 3 If you want to copy existing backup sets to another destination, do the following in the order listed:
 - Select **Duplicate existing backup sets**, and then click **OK**.
 - Select the backup sets you want to copy. For Oracle or DB2 jobs that were created with multiple data streams, under the instance name, select the date on which the backup set was created.
- 4 If you want to duplicate backup sets created when a scheduled backup job runs, do the following in the order listed:
 - Select **Duplicate backup sets following a job**, and then click **OK**.
 - Select the scheduled backup job to be used as the source.
- 5 In the **Properties** pane, under **Destination**, select **Device and Media**.
See [“Device and media options for duplicate backup jobs in a non-CASO environment”](#) on page 372.
See [“Device and media options for a duplicate backup job in a CASO environment”](#) on page 376.
- 6 In the **Properties** pane, under **Settings**, click **General**, and complete the appropriate options.
See [“General options for new duplicate backup set jobs”](#) on page 379.
- 7 In the **Properties** pane, under **Settings**, click **Advanced**, and complete the appropriate options.
See [“Advanced options for new duplicate backup set jobs”](#) on page 379.
- 8 To encrypt the duplicated data, do the following in the order listed:
 - In the **Properties** pane, under **Settings**, click **Network and Security**.
 - Select an encryption type from the list.
 - Select an encryption key from the list or click **Manage keys** to create a new key.
- 9 If you want Backup Exec to notify someone when the backup job completes, in the **Properties** pane, under **Settings**, click **Notification**.
See [“Notification options for jobs”](#) on page 682.

- 10 If you are duplicating data from a scheduled backup job, click **Run Now**.
 The duplicate job will launch immediately following the completion of the scheduled backup job.
- 11 If you are duplicating data from an existing backup set, either click **Run Now** or under **Frequency**, click **Schedule** to set the scheduling options you want to use.
 See [“Schedule options”](#) on page 356.

New job options to duplicate backup sets

When you create a job to duplicate backup data, you have two options. You can duplicate existing backup sets as a new job or you can duplicate an existing job's backup sets when the job is completed.

See [“Duplicating backed up data”](#) on page 369.

Table 6-13 New job options to duplicate backup sets

Item	Description
Duplicate existing backup sets	Creates a duplicate backup of existing backup sets. The backup sets you select in the catalogs are read from the source media and written to the selected destination.
Duplicate backup sets following a job	Creates a duplicate backup of a job's backup sets when the job is completed. The backup job you select runs first, then the backup sets it created are copied to the selected destination.

Selections options for new duplicate backup set jobs

You can create a job to duplicate existing backup sets.

See [“Duplicating backed up data”](#) on page 369.

Table 6-14 Selections options for new duplicate backup set jobs

Item	Description
Selection list	Designates the selection list you want to use for the duplicate backup set job.
Load selections from existing list	Lets you merge existing selection lists.

Table 6-14 Selections options for new duplicate backup set jobs *(continued)*

Item	Description
Search Catalogs	Enables you to find files or other items that you want to back up as part of the duplicate job.
Include/Exclude	Lets you include or exclude files based on file name attributes.
Include subdirectories	Selects the contents of all subfolders when a directory is selected.
Show file details	Displays all the details about the files you select.
Preview pane	Displays a preview pane at the bottom of the dialog box. The preview pane displays additional information about the items you select.
Beginning backup date	Determines the earliest date for which you want to display backup resources.
Ending backup date	Determines the latest date for which you want to display backup resources.
View by Resource	Lets you view selections as a list of resources.
View by Media	Lets you view selections as a list of media.
View Selection Details	Lets you view selections as a list of files and directories.

Device and media options for duplicate backup jobs in a non-CASO environment

You select the storage device and media set on which the duplicate backup job will run.

See [“Duplicating backed up data”](#) on page 369.

This dialog box includes the following options:

Table 6-15 Device and Media options for duplicate backup jobs in a non-CASO environment

Item	Description
Device	<p>Designates a device pool, a robotic library drive, a stand-alone drive, a backup-to-disk folder, a removable backup-to-disk folder, or other type of supported storage device to which you want to send backup data.</p> <p>See “About tape drives and robotic libraries” on page 451.</p> <p>See “About backup-to-disk folders” on page 496.</p> <p>See “About device pools” on page 515.</p> <p>See “About the All Virtual Disks device pool in the Storage Provisioning Option” on page 2029.</p> <p>See “About the Remote Media Agent for Linux Servers” on page 1966.</p> <p>See “About Symantec Online Storage folders” on page 2054.</p> <p>See “About vault stores in the Archiving Option” on page 1445.</p> <p>See “About OpenStorage devices” on page 1574.</p> <p>See “About deduplication storage folders” on page 1579.</p>
Use client-side deduplication	<p>Enables a remote computer that is configured as a Remote Agent for Deduplication to deduplicate data and then send only the unique data directly to the deduplication device that is selected in the Device field. This option is selected by default when you select an OpenStorage device or a deduplication storage folder in the Device field.</p> <p>Note: This option is available only if the Deduplication Option is installed and a deduplication device is selected in the Device field.</p> <p>See “About client-side deduplication” on page 1590.</p>

Table 6-15 Device and Media options for duplicate backup jobs in a non-CASO environment (*continued*)

Item	Description
<p>Use media server deduplication</p>	<p>Enables the media server to deduplicate data. Media server deduplication reduces the size of backups, which reduces storage requirements. If the media server does not support deduplication, the data is deduplicated on an intelligent disk device, such as Symantec PureDisk or a device from a third-party vendor.</p> <p>Note: This option is available only if the Deduplication Option is installed and a deduplication device is selected in the Device field.</p> <p>See “About the Deduplication Option” on page 1568.</p>
<p>Media set</p>	<p>Specifies the media set for the duplicate backup. If you select Overwrite, the media in the drive is overwritten if the media is scratch, or if its overwrite protection period has expired. If allocated or imported media are in the drive, they may also be overwritten depending on the Media Overwrite Protection Level that is set.</p> <p>If you selected one of the append options, the backup will be added to an appendable media (if one exists).</p>
<p>Overwrite media</p>	<p>Places this duplicate backup on an overwritable media. Make sure that appropriate media is in the stand-alone drive or drive pool you select in the Device field in this dialog box.</p> <p>The media in the drive is overwritten if the media is scratch or recyclable (its overwrite protection period has expired). If allocated or imported media are in the drive, they may also be overwritten depending on the Media Overwrite Protection Level that is set.</p> <p>Depending on your configuration, overwritable media is selected from scratch media or recyclable media.</p> <p>See “How Backup Exec searches for overwritable media” on page 233.</p> <p>If the media in the drive is not overwritable, an alert appears requesting that you insert overwritable media.</p>

Table 6-15 Device and Media options for duplicate backup jobs in a non-CASO environment (*continued*)

Item	Description
Append to media, overwrite if no appendable media is available	<p>Appends this duplicate backup to the media set listed in the Media Set field in this dialog box. The duplicate backup set is appended if an appendable media is available in the selected media set; if not, an overwritable media is used and added to the media set.</p> <p>If an append job fills a media, the job continues on another piece of overwritable media.</p> <p>If the media in the drive is not overwritable, an alert appears requesting that you insert overwritable media.</p>
Append to media, terminate job if no appendable media is available	<p>Appends this duplicate backup to the media set listed in the Media Set field in this dialog box. The duplicate backup set is appended if an appendable media is available in the selected media set; if not, the job is terminated.</p>
Eject media after job completes	<p>Ejects the media in the drive when the operation completes.</p>
Retension media before backup	<p>Runs the tape in the drive from beginning to end at a fast speed, which helps the tape wind evenly and run more smoothly past the tape drive heads. Retensioning is primarily for Mini Cartridge and quarter-inch cartridges and is not supported on most other types of tape drives.</p>
Use Write once, read many (WORM) media	<p>Specifies the use of WORM (write once, read many) media for this backup job. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent.</p> <p>See "About WORM media" on page 247.</p>
Enable DirectCopy to tape	<p>Enables data to be copied from a virtual tape library directly to a physical device. The Backup Exec media server records information about the data in the catalog. Because the information about the copied data is in the catalog, you can restore data from either the virtual device or the physical device.</p>

Device and media options for a duplicate backup job in a CASO environment

You select the storage device and media set on which the duplicate backup job will run.

See “[Duplicating backed up data](#)” on page 369.

This dialog box includes the following options:

Table 6-16 Device and media options for a duplicate backup job in a CASO environment

Item	Description
Media or Resource	Displays a list of media that is required for the duplicate job, or the name of the resource that you selected to duplicate.
Media Location	<p>Displays the location of the media. If the media is listed as Offline or Unknown, you must select a device in the Restore Device or Media Server column. Then place the media in a device that the managed media server can access.</p> <p>If the data that is selected resides in a media vault, then Offline appears.</p> <p>If the data that is selected for duplication resides in an unknown media location, then Unknown appears. The media cannot be found in any compatible storage devices that are candidates to run the job.</p>
Device	<p>Displays the names of the devices that match the following criteria:</p> <ul style="list-style-type: none"> ■ They are compatible with the media that you want to duplicate. ■ They are possible candidates to process the job. <p>Backup Exec creates a separate selection list and a separate duplicate job for every device.</p>

Table 6-16 Device and media options for a duplicate backup job in a CASO environment (*continued*)

Item	Description
Media set	<p>Specifies the media set for the duplicate backup. If you select Overwrite, the media in the drive is overwritten if the media is scratch, or if its overwrite protection period has expired. If allocated or imported media are in the drive, they may also be overwritten depending on the Media Overwrite Protection Level that is set.</p> <p>If you selected one of the append options, the backup will be added to an appendable media (if one exists).</p>
Overwrite media	<p>Places this duplicate backup on an overwritable media. Make sure that appropriate media is in the stand-alone drive or drive pool you select in the Device field in this dialog box.</p> <p>The media in the drive is overwritten if the media is scratch or recyclable (its overwrite protection period has expired). If allocated or imported media are in the drive, they may also be overwritten depending on the Media Overwrite Protection Level that is set.</p> <p>Depending on your configuration, overwritable media is selected from scratch media or recyclable media.</p> <p>See “How Backup Exec searches for overwritable media” on page 233.</p> <p>If the media in the drive is not overwritable, an alert appears requesting that you insert overwritable media.</p>

Table 6-16 Device and media options for a duplicate backup job in a CASO environment (*continued*)

Item	Description
Append to media, overwrite if no appendable media is available	<p>Appends this duplicate backup to the media set listed in the Media Set field in this dialog box. The duplicate backup set is appended if an appendable media is available in the selected media set; if not, an overwritable media is used and added to the media set.</p> <p>If an append job fills a media, the job continues on another piece of overwritable media.</p> <p>If the media in the drive is not overwritable, an alert appears requesting that you insert overwritable media.</p>
Append to media, terminate job if no appendable media is available	<p>Appends this duplicate backup to the media set listed in the Media Set field in this dialog box. The duplicate backup set is appended if an appendable media is available in the selected media set; if not, the job is terminated.</p>
Eject media after job completes	<p>Ejects the media in the drive when the operation completes.</p>
Retension media before backup	<p>Runs the tape in the drive from beginning to end at a fast speed, which helps the tape wind evenly and run more smoothly past the tape drive heads. Retensioning is primarily for Mini Cartridge and quarter-inch cartridges and is not supported on most other types of tape drives.</p>
Use Write once, read many (WORM) media	<p>Specifies the use of WORM (write once, read many) media for this backup job. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent.</p> <p>See “About WORM media” on page 247.</p>

Table 6-16 Device and media options for a duplicate backup job in a CASO environment (*continued*)

Item	Description
Enable DirectCopy to tape	Enables data to be copied from a virtual tape library directly to a physical device. The Backup Exec media server records information about the data in the catalog. Because the information about the copied data is in the catalog, you can restore data from either the virtual device or the physical device.

General options for new duplicate backup set jobs

You can create a job to duplicate backup data. You can select either to duplicate existing backup sets or to duplicate backup sets immediately following a scheduled job.

See [“Duplicating backed up data”](#) on page 369.

Table 6-17 General options for new duplicate backup set jobs

Item	Description
Job name	Designates the name for this backup job.
Job priority	Displays the priority of the access to the devices for this job. See “About job priority” on page 195.
Backup set description	Designates a description of the information you back up.
Preferred source device	Designates the device that is used as the destination device for the original backup job.

Advanced options for new duplicate backup set jobs

You can create a job to duplicate backup data. You can select either to duplicate existing backup sets or to duplicate backup sets immediately following a scheduled job.

See [“Duplicating backed up data”](#) on page 369.

Table 6-18 Advanced options for new duplicate backup set jobs

Item	Description
Verify after job completes	Performs a verify operation automatically to make sure that the media can be read after the backup has been completed. Verifying all backups is recommended.
Compression type	<p>Lets you choose from the following compression types:</p> <ul style="list-style-type: none"> ■ None. This option copies the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage media space. Hardware data compression should not be used in environments where devices that support hardware compression are used interchangeably with devices that do not have that functionality. In this situation, hardware compression is automatically disabled. You can manually reenable hardware compression on the drives that support it, but this results in media inconsistency. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive. ■ Hardware [if available, otherwise none]. Select This option uses hardware data compression (if the storage device supports it). If the drive does not feature data compression the data is backed up uncompressed.

Network and Security options for duplicate backup set jobs

You can choose to encrypt a duplicate backup set job.

See [“About encryption”](#) on page 415.

Table 6-19 Network and Security options for duplicate backup set jobs

Item	Description
Encryption type	<p>Specifies the type of encryption you want to use, if any.</p> <p>If the source backup set is encrypted using software, the duplicate backup set is automatically encrypted using software as well. The duplicate backup set is encrypted even if you do not select an encryption type for it.</p>

Table 6-19 Network and Security options for duplicate backup set jobs
(continued)

Item	Description
Encryption key	Specifies the encryption key you want to use. If the source backup set is encrypted, the duplicate backup set automatically uses the same encryption key as the source backup set. The duplicate backup set uses the same encryption key as the source backup set even if you select a different encryption key for it.
Manage keys	Lets you create a new encryption key. You can also replace or delete an existing key.

How to copy data directly from a virtual tape library to a physical tape device

Backup Exec's **DirectCopy to tape** option enables data to be copied from a virtual tape library directly to a physical tape device during a duplicate backup job. The Backup Exec media server coordinates the copy job, but it does not copy the data. Instead, the virtual tape library copies the virtual tape image directly to the physical device. The Backup Exec media server records information about the data in the catalog. Because the information about the copied data is in the catalog, you can restore data from either the virtual tape library or the physical device. The job log for the duplicate backup job indicates that DirectCopy to tape is enabled.

See [“Copying data from a virtual tape library to a physical tape device”](#) on page 382.

To use DirectCopy, both the source device and the destination device must be NDMP-enabled. If the devices are not NDMP-enabled, then Backup Exec performs a regular duplicate backup job.

Note: If you select a backup-to-disk folder as the destination device for a duplicate job with **DirectCopy to tape** enabled, Backup Exec performs a regular duplicate job.

Both hardware encryption and software encryption are supported with DirectCopy. For software encryption, both the source backup set and the destination backup set must use software encryption.

Copying data from a virtual tape library to a physical tape device

You can create a duplicate backup job to copy data directly from a virtual tape library to a physical tape device.

Note: Both the source device and the destination device must be NDMP-enabled. If the devices are not NDMP-enabled, then Backup Exec performs a regular duplicate backup job.

See [“How to copy data directly from a virtual tape library to a physical tape device”](#) on page 381.

Table 6-20 How to use DirectCopy to copy data from a virtual tape library to a physical device

Step	Notes	For more information
Create a regular backup job.	In the Device and Media view, select a virtual tape library as the destination.	See “Creating a backup job by setting job properties” on page 332. See “Device and media options for backup jobs and templates” on page 339.
Create a duplicate backup job.	In the Device and Media view, do the following: <ul style="list-style-type: none"> ■ Select a physical tape device as the destination. ■ Select Enable DirectCopy to tape. 	See “Duplicating backed up data” on page 369. See “Device and media options for duplicate backup jobs in a non-CASO environment” on page 372.

Verifying a backup

In addition to the verification of files that is done when a backup job runs, you can submit verify jobs to test the integrity of the media.

If you perform a verify operation and files fail to verify, the media may be bad. Details about files that failed to verify are provided in the job log, which can be viewed from the Job Monitor.

See [“Setting default backup options”](#) on page 391.

See [“Duplicating backed up data”](#) on page 369.

To verify a backup

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Backup Tasks**, click **New job to verify backup data**.
- 3 Select the media you want to verify.
- 4 In the **Properties** pane, under **Destination**, click **Device**.
- 5 Select the device that contains the media you want to verify.
- 6 In the **Properties** pane, under **Settings**, click **General**.
- 7 Select the appropriate options.
 See “[General properties for verify jobs](#)” on page 384.
- 8 If you want Backup Exec to notify someone when the backup job completes, in the **Properties** pane, under **Settings**, click **Notification**.
 See “[Notification options for jobs](#)” on page 682.
- 9 If you want to run the job now, click **Run Now**. Otherwise, in the **Properties** pane, under **Frequency**, click **Schedule** to set the scheduling options you want to use.
 See “[Schedule options](#)” on page 356.
 After verification has completed, you can check the results in the job log.

Selections properties for verify jobs

You can submit verify jobs to test the integrity of the backup media.

See “[Verifying a backup](#)” on page 382.

Table 6-21 Selections properties for verify jobs

Item	Description
Selection list	Designates the selection list you want to use for the verify job.
Search Catalogs	Enables you to find files or other items that you want to verify.
Include/Exclude	Lets you include or exclude files based on file name attributes.
Include subdirectories	Selects the contents of all subfolders when a directory is selected.

Table 6-21 Selections properties for verify jobs (*continued*)

Item	Description
Show file details	Displays all the details about the files you select.
Preview pane	Displays a preview pane at the bottom of the dialog box. The preview pane displays additional information about the items you select.
Beginning backup date	Specifies the earliest date for which you want to search backups.
Ending backup date	Specifies the latest date for which you want to search backups.
View by Resource	Lets you view selections as a list of resources.
View by Media	Lets you view selections as a list of media.
View Selection Details	Lets you view selections as a list of files and directories.

Device properties for verify jobs

In addition to verifying files after a backup job runs, verify jobs also test the integrity of the media.

See [“Verifying a backup”](#) on page 382.

The **Device** field indicates which device contains the media you want to verify.

General properties for verify jobs

In addition to verifying files after a backup job runs, verify jobs also test the integrity of the media.

See [“Verifying a backup”](#) on page 382.

Table 6-22 General properties for verify jobs

Item	Description
Job name	Designates a name that describes the data you are verifying.

Table 6-22 General properties for verify jobs (*continued*)

Item	Description
Job priority	Displays the priority of the access to the devices for this job. See “About job priority” on page 195.

About test run jobs

The Backup Exec test run option determines if a scheduled backup will complete successfully. When you run a test job, you can monitor the job just as you would a normal backup job, but no data is backed up. During the test run, the tape capacity, credentials, and media are checked. If there is an error, the job will continue to run and the error will appear in the job log. Notification can also be sent to a designated recipient.

During a test run job, the following may cause a job to fail:

- The logon credentials are incorrect.
- Insufficient media is available.
- Media is not in the drive.
- There is no overwritable media for an overwrite job.
- There is no appendable media for an append job.

Test run jobs that are targeted to All Drives will fail the test if any of the devices in the All Drives drive pool cannot handle the job. For example, if one of the devices does not have any media.

A test run job checks media capacity available for the selected job. However, you can check if there is enough available media for multiple test run jobs in the Test Run Results report.

Before you create a test run job, Symantec recommends that you run backup jobs to your devices first. Backup Exec does not recognize the capacity of a backup device until an actual backup job is targeted to the device. If you create a test run job before any other jobs, Backup Exec cannot check that the device has sufficient capacity to perform the backup job. After at least one backup job has been targeted to a device, Backup Exec can determine the capacity.

See [“Creating a test run job”](#) on page 386.

See [“Test Run Results Report”](#) on page 767.

Creating a test run job

The Backup Exec test run option determines if a scheduled backup will complete successfully. When you run a test job, you can monitor the job just as you would a normal backup job, but no data is backed up.

See [“About test run jobs”](#) on page 385.

To create a test run job

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Jobs** pane, select the job for which you want to create a test run.
- 3 Under **General Tasks**, click **Test run**.
- 4 Select the appropriate **General** options.
See [“General properties for test run jobs”](#) on page 386.
- 5 If you want Backup Exec to notify a recipient when the backup job completes, in the **Properties** pane, under **Settings**, click **Notification**.
See [“Notification options for jobs”](#) on page 682.
- 6 If you want to run the job now, click **Run Now**.

Otherwise, in the **Properties** pane, under **Frequency**, click **Schedule** to set the scheduling options you want to use.

See [“Scheduling jobs”](#) on page 356.

General properties for test run jobs

The Backup Exec test run option determines if a scheduled backup will complete successfully. When you run a test job, you can monitor the job just as you would a normal backup job, but no data is backed up.

See [“About test run jobs”](#) on page 385.

See [“Creating a test run job”](#) on page 386.

Table 6-23 General properties for test run jobs

Item	Description
Job name	Designates a name for the test run job.
Credentials check	Verifies that the Backup Exec logon account is correct for the resources being backed up.

Table 6-23 General properties for test run jobs (*continued*)

Item	Description
Media capacity check to complete individual job	<p>Tests if there is enough available capacity on the media to complete the job.</p> <p>During the test run job, the number of scheduled jobs in the queue is not checked; therefore, jobs that are scheduled before the test run job may use the media that was available when the test run job was performed.</p>
Media check	Tests whether the media is online and overwritable.
Use previous job history, if available	Uses past job histories to determine whether there is enough media available to run the scheduled backup job. Checking the previous job history is faster than performing a pre-scan.
Perform Pre-scan	Enables Backup Exec to scan the scheduled backup job to determine whether there is enough media available to run the job. This is the most accurate method of determining media capacity and should be selected if there is not an existing job history.
Upon any failure, place the scheduled job on hold	Places the scheduled job on hold if any failures are detected during the test run.
Run at priority	<p>Designates a priority level for the test job. If another job is scheduled to run at the same time as this job, the priority you set determines which job runs first.</p> <p>You can choose the following priorities:</p> <ul style="list-style-type: none"> ■ Highest ■ High ■ Medium ■ Low ■ Lowest

Setting test run default options

You can set up test run jobs to check the following items:

- Whether credentials are correct
- Whether there is enough available capacity on the media
- Whether the media is online and overwritable

To set test run default options

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Test Run**.
- 3 Select the appropriate options.
See “[Test Run default options](#)” on page 388.
- 4 Click **OK**.

Test Run default options

You can set up test run jobs to check the following items:

- Whether credentials are correct
- Whether there is enough available capacity on the media
- Whether the media is online and overwritable

See “[Setting test run default options](#)” on page 387.

Table 6-24 Test Run default options

Item	Description
Check credentials	Verifies that the Backup Exec logon account is correct for the resources being backed up.
Check media capacity to complete job	Tests if there is enough available capacity on the media to complete the job. During the test run job, the number of scheduled jobs in the queue is not checked; therefore, jobs that are scheduled before the test run job may use the media that was available when the test run job was performed.
Check media availability	Tests whether the media is online and overwritable.
Use previous job history, if available	Uses past job histories to determine whether there is enough media available to run the scheduled backup job. Checking the previous job history is faster than performing a pre-scan.
Perform Pre-scan	Enables Backup Exec to scan the scheduled backup job to determine whether there is enough media available to run the job. This is the most accurate method of determining media capacity and should be selected if there is not an existing job history.

Table 6-24 Test Run default options (*continued*)

Item	Description
Place the scheduled job on hold if any failure occurs during the Test Run job	Places the scheduled job on hold if any failures are detected during the test run.

Customizing backup options

This chapter includes the following topics:

- [Setting default backup options](#)
- [About pre/post commands](#)
- [About specifying backup networks](#)
- [About using Backup Exec with Symantec Endpoint Protection](#)
- [About using Backup Exec with firewalls](#)
- [About encryption](#)
- [Encryption keys](#)
- [About configuring DBA-initiated job settings](#)
- [Editing DBA-initiated jobs](#)
- [Deleting a job template for DBA-initiated jobs](#)
- [About preferred server configurations](#)

Setting default backup options

You can set up Backup Exec with the settings that you want to use for most backup operations, such as the backup method and compression type. If the default options are not appropriate for a particular backup job, you can override the default options when you set up a backup job.

To set default backup options

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Backup**.
- 3 Select the appropriate options.

See “[Default Backup options](#)” on page 392.

Default Backup options

You can set up Backup Exec with the settings that you want to use for most backup operations, such as the backup method and compression type. If the default options are not appropriate for a particular backup job, you can override the default options when you set up a backup job.

See “[Setting default backup options](#)” on page 391.

Table 7-1 Default Backup options

Item	Description
Backup method for files	Designates the default Backup method. Full - Back up files - Using archive bit (reset archive bit) is the typical selection for this field. See “ About backup methods ” on page 274.
Files accessed in x days	Specifies the number of days for which to include accessed files when the Working Set backup method is selected.
Use the Microsoft Change Journal if available	Allows you to use Windows’ NTFS Change Journal to determine which files have been modified since the last full backup. This option can only be used with NTFS volumes and only when the backup method selected is Full - Back up files - Using modified time , Differential - Using modified time , or Incremental - Using modified time .

Table 7-1 Default Backup options (*continued*)

Item	Description
<p>Collect additional information for synthetic backup and for true image restore</p>	<p>Collects additional information for synthetic backup jobs and true image restore jobs. This option displays only for templates.</p> <p>Select this option if you want Backup Exec to do the following:</p> <ul style="list-style-type: none"> ■ Collect the information that is required to detect files and directories that have been moved, renamed, or newly installed since the last backup ■ Include those files and directories in the backup jobs. <p>If you do not select this option, Backup Exec skips the files and directories that have unchanged archive bits. When you select this option, Backup Exec compares path names, file names, modified times, and other attributes with those from the previous full and incremental backups. If any of these attributes are new or changed, then the file or directory is backed up.</p> <p>Backup jobs that have this option selected require more disk space, and take more time to run, than backups that do not.</p> <p>You must select this option for the baseline and incremental backup template in a synthetic backup policy.</p> <p>See “About the synthetic backup feature” on page 899.</p>

Table 7-1 Default Backup options (*continued*)

Item	Description
<p>Media overwrite protection</p>	<p>Provides the following media overwrite options:</p> <ul style="list-style-type: none"> <p>■ Overwrite media</p> <p>Places this backup on an overwritable media. Make sure that appropriate media is in the stand-alone drive or drive pool you select in the Device field in this dialog box. The media in the drive is overwritten if the media is scratch or recyclable (its overwrite protection period has expired). If allocated or imported media are in the drive, they may also be overwritten depending on the Media Overwrite Protection Level that is set.</p> <p>See “Media overwrite protection levels” on page 232.</p> <p>If the media in the drive is not overwritable, a message is displayed requesting that you insert overwritable media.</p> <p>■ Append to media, overwrite if no appendable media is available</p> <p>Adds this backup to the media set listed in the Media Set field in the General applications dialog box.</p> <p>See “Changing application settings for Backup Exec preferences” on page 196.</p> <p>The backup set is appended if an appendable media is available in the selected media set; if not, an overwritable media is used and added to the media set.</p> <p>If an append job fills a media, the job continues on another piece of overwritable media.</p> <p>Depending on your configuration, overwritable media is selected from Scratch media or Recyclable media.</p> <p>See “About media overwrite protection” on page 223.</p> <p>If the media in the drive is not overwritable, a message is displayed requesting that you insert overwritable media.</p> <p>■ Append to media, terminate job if no appendable media is available</p> <p>Adds this backup to the media set listed in the Media Set field in the General applications dialog box.</p> <p>See “Changing application settings for Backup Exec preferences” on page 196.</p> <p>The backup set is appended if an appendable media is available in the selected media set; if not, the job is terminated.</p>

Table 7-1 Default Backup options (*continued*)

Item	Description
Compression type	<p>Provides the following compression types:</p> <ul style="list-style-type: none"> ■ None Copies the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage media space. ■ Software Uses STAC software data compression, which compresses the data before it is sent to the storage device. ■ Hardware [if available, otherwise none] Uses hardware data compression (if the storage device supports it). If the drive does not feature data compression, the data is backed up uncompressed. ■ Hardware [if available, otherwise software] Uses hardware data compression (if the storage device supports it). If the drive does not feature hardware data compression, STAC software compression is used.
Verify after backup	<p>Verifies backups after they are completed. Verify operations make sure the media can be read once the backup has been completed. Verifying all backups is recommended.</p>

Table 7-1 Default Backup options (*continued*)

Item	Description
<p>Back up files and directories by following junction points</p>	<p>Backs up the information for the junction points and the files and directories to which they are linked. If this check box is not selected, then only the information for the junction points is backed up; the files and directories to which they are linked are not backed up.</p> <p>Backup Exec does not follow junction points automatically created by Microsoft Windows Vista/Server 2008 because it can cause the data to be backed up repeatedly. You can find information about junction points at the following URL: http://entsupport.symantec.com/umi/V-269-9</p> <p>If you use junction points created by linkd.exe (or a similar tool) to span volumes, then Advanced Open File Option (AOFO) backups and Change Journal incremental backups will not follow the junction points properly. To perform AOFO and Change Journal incremental backups of volumes with junction points, clear this option. Junction points created by Disk Manager or mountvol.exe are supported.</p> <p>Since Mounted Drives that do not have a drive letter assigned to them cannot be selected, the files and directories to which they are linked are backed up regardless of whether this option is selected.</p> <p>If this option is selected and the actual files and directories to which the junction points are linked are also included in the backup selections, then the files and directories are backed up twice; once during the full file and directory backup, and again via the junction point.</p> <p>Warning: If a junction point is linked to a location that encompasses it, then recursion (a situation where data is backed up repeatedly) will occur, resulting in an error and job failure. For example, if c:\junctionpoint is linked to c:\, recursion will occur when attempting to back up c:\junctionpoint, and the backup job will fail.</p>

Table 7-1 Default Backup options (*continued*)

Item	Description
<p>Back up files and directories by following symbolic links</p>	<p>Backs up the information for the symbolic links and the files and directories to which they are linked.</p> <p>If you do not select this option, only the information for the symbolic links is backed up. The files and directories to which they are linked are not backed up.</p> <p>If the symbolic link points to files and directories on a remote computer, the files and directories on the remote computer are not backed up.</p>
<p>Back up data in Remote Storage</p>	<p>Backs up data that has been migrated from primary storage to secondary storage. The data will not be recalled to its original location; it will be backed up directly to the backup media.</p> <p>If this option is selected, you should not run a backup of your entire system because Backup Exec will have to load the data that has been migrated to secondary storage and additional time will be required for any set that includes migrated data.</p> <p>If this check box is cleared, only the placeholder that stores the location of the data on secondary storage will be backed up, not the data itself.</p> <p>This option should not be selected if the device used for secondary storage and backups contains only one drive because Remote Storage and Backup Exec will compete for use of the drive.</p>

Table 7-1 Default Backup options (*continued*)

Item	Description
<p>Enable single instance backup for NTFS</p>	<p>Enables single instance backup for NTFS volumes. This option is only available if you use the Microsoft Windows Single Instance Store (SIS) feature. Single instance backup checks the NTFS volume for identical files. If Backup Exec finds multiple copies of a file, it only backs up one instance of that file, regardless of how many SIS links reference it.</p> <p>Single instance backup can considerably reduce the storage space that is required for your backups. Many applications automatically generate files that have identical content. The actual amount of space you save depends on the number of duplicate files on the volume.</p> <p>If the backup job does not run to completion, the file data may not be included in the backup set. Rerun the backup until it is successfully completed. If the incremental backup method was used, running the job again will not back up the same files. You must run a full or copy backup to ensure that all files are backed up completely. If the incremental - using modified time backup method was used, running the same backup job to completion will back up the files correctly.</p>
<p>Never</p>	<p>Skips open files if they are encountered during the backup operation. A listing of skipped files appears in the job log for the backup.</p>
<p>If closed within x seconds</p>	<p>Waits the specified time interval for files to close before skipping the open file and continuing the backup operation.</p> <p>If the file does not close during the specified interval, it is skipped. A listing of skipped files appears in the job log for the backup.</p> <p>If multiple files are open, Backup Exec waits the specified time interval for each file; depending on the number of open files, this may significantly increase the backup time.</p>
<p>With a lock</p>	<p>Attempts to open files that are in use. If Backup Exec is able to open a file, the file is locked while it is being backed up to prevent other processes from writing to it. Backing up open files is not as effective as closing applications and allowing the files to be backed up in a consistent state.</p>

Table 7-1 Default Backup options (*continued*)

Item	Description
Without a lock	<p>Attempts to open files that are in use. If Backup Exec is able to open the file, the file is NOT locked while it is being backed up. This allows other applications to write data to the file during the backup operation.</p> <p>Warning: This option allows files to be backed up that contain inconsistent data and possibly corrupt data.</p> <p>See “About the Advanced Open File Option” on page 937.</p>
If Backup Exec Granular Recovery Technology (GRT) is enabled for backup, enter a path on the NTFS volume of the local media server where Backup Exec can stage temporary data	<p>Designates a location where Backup Exec can stage temporary data during GRT-enabled jobs. Ensure that the default location of C:\temp is an NTFS volume, and that it is not a system volume. If C:\temp does not meet these requirements, type another path to an NTFS volume on the local media server where Backup Exec can stage temporary data.</p> <p>Backup Exec deletes the data when the backup completes.</p> <p>At least 1 GB of disk space is required.</p>

About pre/post commands

You can set defaults for the commands you want to run before or after all backup and restore jobs. If the default options are not appropriate for a particular job, you can override the default options when you create the job.

Conditions that you can set for these commands include the following:

- Run the backup or restore job only if the pre-command is successful
- Run the post-command only if the pre-command is successful
- Run the post-command even if the backup or restore job fails
- Allow Backup Exec to check the return codes (or exit codes) of the pre- and post-commands to determine if the commands completed successfully. An exit code of zero returned to the operating system by the pre- or post-command is interpreted by Backup Exec to mean the job completed successfully. A non-zero exit code is interpreted by Backup Exec to mean the job ended with an error.

If it is critical that the job not run if the pre-command fails, then configure Backup Exec to check the return codes of the pre- and post-commands to determine if the pre-command failed or completed successfully.

For example, if a pre-command that shuts down a database before a backup is run fails, the database could be corrupted when the backup is run. In this situation, it is critical that the backup job not run if the pre-command fails.

Additionally, if Backup Exec is configured to check the return codes of the pre- and post-commands, and the post-command returns a non-zero code, the job log reports that the post-command failed. If you also selected to run the job only if the pre-command is successful, and both the pre-command and the job ran successfully, Backup Exec will mark the job as failed if the post-command fails.

For example, if the pre-command runs successfully and shuts down the database and the backup job also runs successfully, but the post-command cannot restart the database, Backup Exec marks the job and the post-command as failed.

If you select the option **On each server backed up**, the pre- and post-commands are run and completed for each server before processing begins on the next selected server.

See [“Setting default pre/post commands”](#) on page 400.

See [“Pre/post commands for backup or restore jobs”](#) on page 352.

See [“Running pre and post commands for restore jobs”](#) on page 618.

Setting default pre/post commands

You can set defaults for the commands you want to run before or after all backup and restore jobs. If the default options are not appropriate for a particular job, you can override the default options when you create the job.

See [“About pre/post commands”](#) on page 399.

See [“Pre/post commands for backup or restore jobs”](#) on page 352.

To set default pre/post commands

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Pre/Post Commands**.
- 3 Select the appropriate options.

See [“Default Pre/Post Commands options”](#) on page 400.

- 4 Click **OK**.

Default Pre/Post Commands options

You can set defaults for the commands you want to run before or after all backup and restore jobs. If the default options are not appropriate for a particular job, you can override the default options when you create the job.

See “[Setting default pre/post commands](#)” on page 400.

Table 7-2 Default Pre/Post Commands options

Item	Description
<p>Allow pre- and post-commands to be successful only if completed with a return code of zero</p>	<p>Allows Backup Exec to check the return codes of the pre- and post-commands to determine if they completed successfully.</p> <p>An exit code of zero returned to the operating system by the pre- or post-command is interpreted by Backup Exec to mean the job completed successfully. A non-zero exit code is interpreted by Backup Exec to mean the job ended with an error.</p> <p>After checking the return codes, Backup Exec continues processing according to selections you made for running the pre- and post-commands.</p> <p>If this option is not selected, the success of the pre- and post-commands is not determined based on the return code.</p>
<p>Run job only if pre-command is successful</p>	<p>Runs the backup or restore job only if the pre-command is successful. If the pre-command fails, the job does not run, and is marked as failed.</p> <p>If it is critical that the job not run if the pre-command fails, then select Allow pre- and post-commands to be successful only if completed with a return code of zero. If a non-zero code is returned, it is interpreted by Backup Exec to mean that the pre-command did not run successfully. The job is not run and the job status is marked as Failed.</p>
<p>Run post-command only if pre-command is successful</p>	<p>Runs the post-command only if the pre-command is successful.</p> <p>If it is critical that the post-command fail if the pre-command fails, then select Allow pre- and post job commands to be successful only if completed with a return code of zero. If a non-zero code is returned for the pre-command, it is interpreted by Backup Exec to mean that the pre-command did not run successfully. The post-command is not run.</p> <p>If you also select Run job only if pre-command is successful, and both the pre-command and the job are successful, but the post-command returns a non-zero code, the job log reports both the job and the post-command as failed.</p>
<p>Run post-command even if job fails</p>	<p>Runs the post-command whether the job is successful or not.</p> <p>If you also select Allow pre- and post job commands to be successful only if completed with a return code of zero and the post-command returns a non-zero code, the job log reports the post-command as failed.</p>

Table 7-2 Default Pre/Post Commands options (*continued*)

Item	Description
Run post-command after job verification completes	Runs the post-command after the verification completes if you selected the Verify after backup completes option on the General backup properties dialog box.
Cancel command if not completed within x minutes	Designates the number of minutes Backup Exec should wait before canceling a pre-job or post-command that did not complete. The default time-out is 30 minutes.
On this media server	Runs the pre- and post-commands on this media server only.
On each server backed up or restored to	Runs the pre- and post-commands one time on each server backed up or restored to. The pre- and post-command selections apply to each server independently. If you select this option, the pre- and post-commands are run and completed for each server before processing begins on the next selected server.

About specifying backup networks


The Backup Network feature allows you to direct primary backup traffic generated by Backup Exec to a specific local network. Directing backup jobs to a specified local network isolates the backup data so that other connected networks are not affected when backup operations are performed. You also can use a backup network when restoring data. The feature is enabled on the media server and allows you to protect all the remote computers that reside on the specified local network.

When the feature is enabled and a backup job is submitted, Backup Exec verifies that the remote computer is on the same subnet as the selected interface on the media server. If the remote computer is on the selected subnet, then the backup operation is performed.

If the remote computer is not on the selected subnet, then the job fails. However, you can set up Backup Exec to use any available network to back up remote computers.

The following diagram shows an example of a basic backup network configuration.

Figure 7-1 Example of backup network



In the example, the database server and mail server are connected to both the backup network and the corporate network.

When backup operations are performed by the Backup Exec media server, the backup data will use either the backup network or the corporate network to back up the database server. If the backup data goes through the corporate network, the amount of time it takes to back up the database server will increase because the network route between the two computers is longer. This may cause users to experience network latencies when accessing the mail server since there is an increase in network traffic.

In contrast, if the Specified Backup Network feature is enabled and you back up the database server, the backup data traffic is isolated to the backup network and users accessing the mail server are not affected. The backup network will be used to perform all backup operations, unless the remote computer is not connected to the backup network.

If you want to back up remote computers that are not connected to the backup network, such as the database user's computer, then choose to use any available network route. This allows you to back up the remote computer even though it does not reside on the backup network.

See [“About using Backup Exec with firewalls”](#) on page 409.

See [“Browsing systems through a firewall”](#) on page 414.

About using IPv4 and IPv6 in Backup Exec

Backup Exec supports versions 4 and 6 of the Internet Protocol (IP), which are commonly referred to as IPv4 and IPv6. You can use IPv4 and IPv6 in backup and restore networks. Support for IPv6 is dependent upon operating system support for the protocol, as well as proper network configuration.

You can use Backup Exec in a mixed IPv4/IPv6 environment or an IPv4-only environment.

Enter an IPv4 or IPv6 address for a computer anywhere that you can enter a computer name in Backup Exec, except in the following locations:

- User-defined selections.
- Clusters. Microsoft Windows does not support an IPv6 address as a clustered resource.
- The Connect to Media Server dialog box.

A Remote Agent that supports IPv6 can be backed up or restored using IPv6 only from a media server that is IPv6-compliant.

Setting default backup network and security options

You can specify a network to be used as the default for every Backup Exec job. Before configuring the feature, test for network connectivity between the media server and the remote computers.

Note: The remote computers that you want to back up must have the most current version of Backup Exec Remote Agent installed.

You can also set default security options for Backup Exec jobs. You can select a default encryption type or key for your backup jobs. If you use Symantec Endpoint Protection 11.0 or later, you can configure it to prompt Backup Exec to back up data when global threats arise.

The backup settings you select are set as the default for all new backup jobs and templates you create. You can manually change these settings when you create specific jobs or templates.

To set the default backup network and security options

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Network and Security**.
- 3 Select the appropriate options.

See “[Default Network and Security options](#)” on page 405.

Default Network and Security options

You can select default network and security options for all new backup jobs and templates you create. You can manually change these settings when you create specific jobs or templates.

See “[Setting default backup network and security options](#)” on page 404.

Table 7-3 Default Network and Security options

Item	Description
Enable selection of user shares	Lets you include user-defined shares in jobs. If you do not select this option, you cannot select user-defined shares when you create jobs.
Enable remote agent TCP dynamic port range	Allows remote agents to use a range of ports for communications. You enter the port range. If the first port that Backup Exec attempts to use is not available, communications will be attempted through one of the other ports in the range. If none of the ports in the range are available, Backup Exec will use any available dynamic port. Default port ranges are 1025 to 65535. Symantec recommends using a range of 25 allocated ports for the remote systems if using Backup Exec with a firewall. See “ About using Backup Exec with firewalls ” on page 409.
Interface	Indicates the name of the network interface card that connects the media server to the network you want to use for the backup network. The list includes all available network interface cards on the media server.
Interface details	Displays the Media Access Control (MAC) address, Adapter type, Description, IP addresses, and subnet prefixes of the network interface you selected for the backup network.
Protocol	Allows you to choose from the following protocol options: <ul style="list-style-type: none"> ■ Use any available protocol ■ Use IPv4 ■ Use IPv6

Table 7-3 Default Network and Security options (*continued*)

Item	Description
Subnet	Displays the 32-bit number that determines the subnet to which the network interface card belongs.
Allow use of any available network interface, protocol, or subnet for remote agents not bound to the above network interface, protocol, or subnet	<p>Ensures that the data from the remote system is backed up or restored over any available network if the remote system that you selected for backup or restore is not part of the specified backup network.</p> <p>If you do not select this check box and you selected data from a remote system that is not part of the specified backup network, the job fails because Backup Exec cannot back up or restore the data from the remote system.</p>
Use a custom port to receive operation requests from the remote system	<p>Specifies the port used for communications between this computer and the remote computer for both DBA- and media server-initiated operations. By default, port 5633 is used.</p> <p>If you change the port number on the remote Windows or Linux computer, you must also change it on the media server, and then restart the Backup Exec Job Engine service on the media server.</p> <p>See “About Oracle instance information changes” on page 1333.</p>
Use FIPS 140-2 compliant software encryption	<p>Enables software encryption that complies with FIPS 140-2 standards. If you select this option, you must use a 256-bit AES encryption key. This option is only available for Windows computers.</p> <p>You must stop and restart the Backup Exec services for this change to take effect.</p>
Encryption type	<p>Specifies the type of encryption you want to use, if any.</p> <p>See “About encryption” on page 415.</p>
Encryption key	Specifies the default encryption key you want to use.
Manage keys	Allows you to create a new encryption key or delete an existing encryption key.
Run backup immediately when an elevated Symantec ThreatCon level is reached	Runs automatic backups when the Symantec ThreatCon reaches the level you specify in the Symantec ThreatCon level field. You must have Symantec Endpoint Protection 11.0 or later installed on the same computer as Backup Exec to use this feature.

Table 7-3 Default Network and Security options (*continued*)

Item	Description
Symantec ThreatCon level	<p>Specifies the ThreatCon level at which you want automatic backups to run.</p> <p>You can find more information about Symantec ThreatCon levels at the following URL:</p> <p>http://www.symantec.com</p>

Network and Security backup options

When you are setting up a new backup job, you can change the backup network for that job. When you change the backup network for a job, you also can set that backup network as the new default backup network for all future backup jobs.

See “[About specifying backup networks](#)” on page 402.

You can choose to encrypt a backup job. If you use Symantec Endpoint Protection 11.0 or later, you can configure the job to run automatically when global threats arise.

See “[About encryption](#)” on page 415.

See “[About using Backup Exec with Symantec Endpoint Protection](#)” on page 408.

Table 7-4 Network and Security backup options

Item	Description
Network Interface	<p>Specifies the name of the network interface card that connects the media server to the network you want to use for the backup network for this backup job. The list includes all available network interfaces on the media server.</p> <p>If you are using the Central Admin Server Option (CASO), select the Use the default network interface for the managed media server option if you want CASO delegated backup jobs to be processed using the network interface card configured as the default in the managed media server.</p>
Protocol	<p>Specifies the protocol you want to use for this backup job.</p> <p>The options are as follows:</p> <ul style="list-style-type: none"> ■ Use any available protocol ■ Use IPv4 ■ Use IPv6

Table 7-4 Network and Security backup options (*continued*)

Item	Description
Subnet	Displays the 32-bit number that determines the subnet to which the network interface card belongs.
Allow use of any available network interface, subnet, or protocol for remote agents not bound to the above network interface, subnet, or protocol	<p>Ensures that the data from the remote system is backed up or restored over any available network if the remote system that you selected for backup or restore is not part of the specified backup network.</p> <p>If you do not select this check box and you selected data from a remote system that is not part of the specified backup network, the job fails because Backup Exec cannot back up or restore the data from the remote system.</p>
Interface details	Displays the Media Access Control (MAC) address, Adapter type, Description, IP addresses, and subnet prefixes of the network interface that you selected for the backup network.
Encryption type	<p>Specifies the type of encryption you want to use, if any.</p> <p>See “About encryption” on page 415.</p>
Encryption key	Specifies the encryption key you want to use.
Manage keys	Allows you to create a new encryption key or delete an existing encryption key.
Run this backup job immediately when an elevated Symantec ThreatCon level is reached	Runs this backup automatically when the Symantec ThreatCon reaches the level you specify in the Symantec ThreatCon level field. You must have Symantec Endpoint Protection 11.0 or later installed on the same computer as Backup Exec to use this feature.
Symantec ThreatCon level	<p>Specifies the ThreatCon level at which you want this backup to run.</p> <p>You can find more information about Symantec ThreatCon levels at the following URL:</p> <p>http://www.symantec.com</p>

About using Backup Exec with Symantec Endpoint Protection

You can use Symantec Endpoint Protection version 11.0 or later with Backup Exec to provide extra security when the threat of viruses or malware is high. You can

also view Symantec Endpoint Protection security information from the Backup Exec Security Summary. You must install the Symantec Endpoint Protection Manager component to use the Security Summary.

Symantec Endpoint Protection uses a ThreatCon level to provide an overall view of global Internet security. Symantec's ThreatCon levels are based on a 1-4 rating system, with level 4 being the highest threat level.

You can find more information about Symantec ThreatCon levels at the following URL:

<http://www.symantec.com>

You can configure Backup Exec to automatically run a backup job when the ThreatCon reaches a level that you specify. You may want to configure special jobs for your most crucial data, for example. This strategy helps make sure that your vital data is safely backed up as soon as global threats are detected.

You should consider the types of jobs that you want to trigger automatically and the potential impact they can have on your system resources. The ThreatCon level is updated frequently and can be raised at any time without warning. If you configure large or resource-intensive jobs to launch automatically, they may interfere with your normal business operations.

The media server must be connected to the Internet to monitor the ThreatCon level. If the media server is not connected to the Internet, backup jobs are not triggered when the ThreatCon level elevates.

See the *Administrator's Guide for Symantec Endpoint Protection* for more information about Symantec Endpoint Protection.

See “[Network and Security backup options](#)” on page 407.

See “[Setting default backup network and security options](#)” on page 404.

See “[Viewing the Symantec Endpoint Protection Security Summary](#)” on page 590.

About using Backup Exec with firewalls

In firewall environments, Backup Exec provides the following advantages:

- The number of ports that are used for backup network connections is kept to a minimum.
- Open ports on the Backup Exec media server and remote systems are dynamic and offer high levels of flexibility during browsing, backup, and restore operations.

- You can set specific firewall port ranges and specify backup and restore networks within these ranges. You can use specific ranges to isolate data traffic and provide high levels of reliability.

Note: The Remote Agent for Windows Systems is required to perform remote backups and restores.

Firewalls affect system communication between a media server and any remote systems that reside outside the firewall environment. You should consider special port requirements for your firewall when you configure Backup Exec.

Symantec recommends that you open port 1000 and make sure that it is available on the Backup Exec media server and any remote systems. In addition, you must open the dynamic port ranges that Backup Exec uses for communications between the media server and remote agents.

See [“Backup Exec Ports”](#) on page 411.

When a media server connects to a remote system, it initially uses port 1000. The Remote Agent listens for connections on this predefined port. The media server is bound to an available port, but additional connections to the Remote Agent are initiated on any available port.

When you back up data, up to two ports may be required on the computer on which the Remote Agent is installed. To support simultaneous jobs, you must configure your firewall to allow a range of ports large enough to support the number of simultaneous operations desired.

If there is a conflict, you can change the default port to an alternate port number by modifying the `%systemroot%\System32\drivers\etc\services` file. You can use a text editor such as Notepad to modify your NDMP entry or add an NDMP entry with a new port number. You should format the entry as follows:

```
ndmp      10000/tcp      #Network Data Management Protocol
```

Note: If you change the default port, you must change it on the media server and all remote systems that are backed up through the firewall.

When you set up TCP dynamic port ranges, Symantec recommends that you use a range of 25 allocated ports for the remote computer. The number of ports that remote computers require depends on the number of devices you protect and the number of tape devices you use. You may need to increase these port ranges to maintain the highest level of performance.

Unless you specify a range, Backup Exec uses the full range of dynamic ports available. When performing remote backups through a firewall, you should select a specific range on the **Network and Firewall defaults** dialog box.

See [“Backup Exec Listening Ports”](#) on page 412.

See [“Backup Exec Desktop and Laptop Option ports”](#) on page 413.

Backup Exec Ports

You may have special port requirements for Backup Exec if you use a firewall. Firewalls sometimes affect system communications between a media server and remote systems that reside outside the firewall environment.

See [“About using Backup Exec with firewalls”](#) on page 409.

The following table provides more information about which ports Backup Exec and its agents and options use:

Table 7-5 Backup Exec Ports

Service or Process	Port	Port Type
Backup Exec Agent Browser (process=benetns.exe)	6101	TCP
Backup Exec Remote Agent for Windows Systems (process=beremote.exe)	10000	TCP
Backup Exec media server (process=beserver.exe)	3527, 6106	TCP
MSSQL\$BKUPEXEC (process=sqlservr.exe)	1125 1434 (ms-sql-m)	TCP UDP
Backup Exec Remote Agent for NetWare	10000 (Backup Exec 10.x), 6102 (Backup Exec 9.x)	TCP
Oracle Agent for Windows and Linux Servers	Random port unless configured otherwise	
DB2 Agent for Windows and Linux Servers	Random port unless configured otherwise	
Remote Agent for Linux or UNIX Servers (RALUS)	Default NDMP port, typically 10000	TCP
Kerberos	88	UDP

Table 7-5 Backup Exec Ports (*continued*)

Service or Process	Port	Port Type
NETBIOS	135	TCP, UDP
NETBIOS Name Service	137	UDP
NETBIOS Datagram Service	138	UDP
NETBIOS Session Service	139	TCP
NETBIOS (Windows 2000)	445	TCP
DCOM/RPC	3106	TCP
Backup Exec Remote Agent	6103	TCP
Push Install - Check for conflicts in message queue for CASO which is part of beserver.exe	103x	TCP
Push Install	441	TCP
SMTP email notification	25 outbound from media server	TCP
SNMP	162 outbound from media server	TCP

Backup Exec Listening Ports

You may have special port requirements for Backup Exec if you use a firewall. Firewalls sometimes affect system communications between a media server and remote systems that reside outside the firewall environment.

See [“About using Backup Exec with firewalls”](#) on page 409.

When Backup Exec is not running operations, it listens to ports for incoming communication from other services and agents. Backup Exec initially communicates with the Remote Agent using a static listening port to begin an operation. The agent and the media server then use dynamic ports to pass data back and forth.

Backup Exec uses the following listening ports:

Table 7-6 Backup Exec Listening Ports

Service	Port	Port Type
Backup Exec Agent Browser (benetns.exe)	6101	TCP
Backup Exec Remote Agent for Windows Server (beremote.exe)	10000	TCP
Backup Exec media server (beserver.exe)	3527, 6106	TCP
MSSQL\$BKUPEXEC (sqlsevr.exe)	1125 1434	TCP UDP
Backup Exec Remote Agent for NetWare	10000, 6102	TCP
Remote Agent for Linux and UNIX Servers (RALUS)	10000	TCP
DBA-initiated backups for Oracle and DB2	5633	TCP

Backup Exec Desktop and Laptop Option ports

You may have special port requirements for Backup Exec if you use a firewall. Firewalls sometimes affect system communications between a media server and remote systems that reside outside the firewall environment.

See [“About using Backup Exec with firewalls”](#) on page 409.

The Backup Exec Desktop and Laptop Option (DLO) uses the following ports:

Table 7-7 Backup Exec Desktop and Laptop Option Ports

Service or Process	Port	Port Type
Server Message Block (SMB) communication	135-139	TCP/UDP
Server Message Block (SMB) communication without NETBIOS	445	TCP/UDP
SQL	1434	TCP/UDP

Table 7-7 Backup Exec Desktop and Laptop Option Ports (*continued*)

Service or Process	Port	Port Type
DLOAdminSvcu.exe (DLO admin service)	3999 in listening mode	TCP/UDP

Browsing systems through a firewall

Because most firewalls do not allow a remote system to be displayed in the Microsoft Windows Network tree, you may need to take additional steps to select these remote systems in the Backup Exec administration console.

To browse systems through a firewall

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Network and Security**.
- 3 Verify that a dynamic range of ports has been set for the media server and remote agent and that the firewall is configured to pass these port ranges and the 10,000 port (which is used for the initial connection from the media server to the Remote Agent).

Port 6101 must be open to browse Windows systems in the backup selections tree.
- 4 Click **OK**.

About enabling a SQL instance behind a firewall

If you want to connect to a SQL instance behind a firewall, you must enable the SQL instance for communication. To enable the SQL instance for communication, you must make the SQL port static and configure the Windows Firewall.

The Backup Exec SQL instance is configured to use a dynamic port by default. Each time SQL Server is started, the port number can change.

See [“Changing the dynamic port on the SQL Express instance in CASO to a static port”](#) on page 1518.

See [“Opening a SQL port in CASO for a SQL 2005 or 2008 instance”](#) on page 1520.

You also must configure the Windows Firewall to allow connections to the SQL instance. There may be multiple ways to configure the Windows Firewall based on your system configuration. You can add sqlsvr.exe and sqlbrowser.exe to the Windows Firewall Exceptions list or you can open a port in the Windows Firewall for TCP access. Refer to the Microsoft Knowledge Base for more information or to determine which configuration is best for your network.

About encryption

Backup Exec provides you with the ability to encrypt data. When you encrypt data, you protect it from unauthorized access. Anyone that tries to access the data has to have an encryption key that you create. Backup Exec provides software encryption, but it also supports some devices that provide hardware encryption with the T10 standard.

Backup Exec supports two security levels of encryption: 128-bit Advanced Encryption Standard (AES) and 256-bit AES. The 256-bit AES encryption provides a stronger level of security because the key is longer for 256-bit AES than for 128-bit AES. However, 128-bit AES encryption enables backup jobs to process more quickly. Hardware encryption using the T10 standard requires 256-bit AES.

See [“About software encryption”](#) on page 415.

See [“About hardware encryption”](#) on page 416.

See [“Encryption keys”](#) on page 416.

About software encryption

When you install Backup Exec, the installation program installs the necessary encryption software on the media server and on remote computers that use the Remote Agent. Backup Exec can encrypt data at a computer that uses the Remote Agent, and then transfer the encrypted data to the media server. Backup Exec then writes the encrypted data on a set-by-set basis to tape or to a backup-to-disk folder.

Backup Exec encrypts the following types of data:

- User data, such as files and Microsoft Exchange databases.
- Metadata, such as file names, attributes, and operating system information.
- On-tape catalog file and directory information.

Backup Exec does not encrypt Backup Exec metadata or on-disk catalog file and directory information.

You can use software compression with encryption for a backup job. First Backup Exec compresses the files, and then encrypts them. However, backup jobs take longer to complete when you use both encryption and software compression.

Symantec recommends that you avoid using hardware compression with software encryption. Hardware compression is performed after encryption. Data becomes randomized during the encryption process. Compression does not work effectively on data that is randomized.

See [“About encryption”](#) on page 415.

About hardware encryption

Backup Exec supports hardware encryption for storage devices that use the T10 encryption standard. When you use hardware encryption, the data is transmitted from the host computer to the target device and then encrypted on the device. Backup Exec manages the encryption keys that are used to access the encrypted data.

Backup Exec only supports approved devices for T10 encryption.

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/v-269-2>

See “[About encryption](#)” on page 415.

Encryption keys

You can set a default encryption key to use for all backup jobs, templates, and duplicate backup set jobs. However, you can override the default key for a specific job. You can also use encryption in policies when you create Backup templates or Duplicate Backup Set templates. When you create a Duplicate Backup Set template or a duplicate backup sets job, backup sets that are already encrypted are not re-encrypted. However, you can encrypt any unencrypted backup sets.

If you use encryption in a synthetic backup policy, all the templates in the policy must use the same encryption key. You should not change the key after you create the policy. For the synthetic backup template, Backup Exec automatically uses the encryption key that you select for the other templates in the policy.

When you select encrypted data for restore, Backup Exec verifies that encryption keys for the data are available in the database. If any of the keys are not available, Backup Exec prompts you to recreate the missing keys. If you delete the key after you schedule the job to run, the job fails.

If Backup Exec cannot locate an encryption key while a catalog job is processing, Backup Exec sends an alert. You can then recreate the missing encryption key if you know the pass phrase.

If you use encryption keys with the Intelligent Disaster Recovery option, special considerations apply.

See “[About encrypted backup sets and the Intelligent Disaster Recovery Wizard](#)” on page 1831.

See “[About encryption](#)” on page 415.

See “[Setting default backup network and security options](#)” on page 404.

See “[About deleting an encryption key](#)” on page 421.

See [“Deleting an encryption key”](#) on page 422.

About restricted keys and common keys in encryption

Backup Exec has the following types of encryption keys:

Table 7-8 Types of encryption keys

Key type	Description
Common	Anyone can use the key to encrypt data during a backup job and to restore encrypted data.
Restricted	Anyone can use the key to encrypt data during a backup job. If a user other than the key owner tries to restore data that was encrypted with a restricted key, Backup Exec prompts the user for the key’s pass phrase. If the user cannot supply the correct pass phrase for the key, the user cannot restore the data.

About pass phrases in encryption

Encryption keys require a pass phrase, which is similar to a password. Pass phrases are usually longer than passwords and are comprised of several words or groups of text. A good pass phrase is between eight and 128 characters. The minimum number of characters for 128-bit AES encryption is eight. The minimum number of characters for 256-bit AES encryption is 16. Symantec recommends that you use more than the minimum number of characters.

Note: Hardware encryption that uses the T10 standard requires 256-bit AES. Backup Exec does not let you enable hardware encryption for a job unless it uses at least a 16-character pass phrase.

Also, a good pass phrase contains a combination of upper and lower case numbers, letters, and special characters. You should avoid using literary quotations in pass phrases.

A pass phrase can include only printable ASCII characters, which are characters 32 through 126. ASCII character 32 is the space character, which is entered using the space bar on the keyboard. ASCII characters 33 through 126 include the following:

! " # \$ % & ' * + , - . / 0 1 2 3 4 5 6 7 8 9 : ; < = > ? @ A B C D E F G H I J K L M N O P Q R S T U V W X Y Z [\] ^ _ ` a b c d e f g h i j k l m n o p q r s t u v w x y z { | } ~

See [“Creating an encryption key”](#) on page 420.

See [“Setting default backup network and security options”](#) on page 404.

About encryption key management

When a user creates an encryption key, Backup Exec marks that key with an identifier based on the logged-on user's security identifier. The person who creates the key becomes the owner of the key.

Backup Exec stores the keys in the Backup Exec database. However, Backup Exec does not store the pass phrases for the keys. The owner of each key is responsible for remembering the pass phrase for the key.

To protect your keys, Symantec recommends the following:

- Maintain a written log of the pass phrases. Keep the log in a safe place in a separate physical location from the encrypted backup sets.
- Back up the Backup Exec database. The database keeps a record of the keys.

Caution: If you do not have a backup of the Backup Exec database and do not remember your pass phrases, you cannot restore data from the encrypted media. In addition, Symantec cannot restore encrypted data in this situation.

A key that is created on a media server is specific to that media server. You cannot move keys between media servers. However, you can create new keys on a different media server by using existing pass phrases. A pass phrase always generates the same key. In addition, if you delete a key accidentally, you can recreate it by using the pass phrase.

If a Backup Exec database becomes corrupted on a media server and is replaced by a new database, you must manually recreate all of the encryption keys that were stored on the original database.

If you move a database from one media server to another media server, the encryption keys remain intact as long as the new media server meets the following criteria:

- Has the same user accounts as the original media server.
- Is in the same domain as the original media server.

See [“Encryption keys”](#) on page 416.

See [“About pass phrases in encryption”](#) on page 417.

See [“About deleting an encryption key”](#) on page 421.

See [“Replacing an encryption key”](#) on page 421.

See [“Deleting an encryption key”](#) on page 422.

Encryption Key Management options

From the **Encryption Key Management** dialog box, you can perform several encryption key management tasks.

See [“Creating an encryption key”](#) on page 420.

See [“Replacing an encryption key”](#) on page 421.

See [“Deleting an encryption key”](#) on page 422.

Table 7-9 Encryption Key Management options

Item	Description
Key name	Indicates the name of the encryption key.
Created by	Indicates who created the encryption key. When a user creates an encryption key, Backup Exec marks that key with an identifier based on the logged-on user’s security identifier. The person who creates the key becomes the owner of the key.
Restricted	Indicates if the key is a restricted key. If a key is restricted, anyone can use the key to back up data. But only the key owner or a user who knows the pass phrase can use the restricted key to restore the encrypted data.
Default	Indicates whether the key is configured as the default key for the jobs that are encrypted.
Encryption Type	Indicates the type of encryption that is associated with the encryption key.
Date Created	Indicates the date the encryption key was created.
Date Last Accessed	Indicates the date the encryption key was last accessed.
New	Lets you create a new encryption key.
Delete	Deletes the selected encryption key.
Replace	Replaces the selected encryption key with the key you select from the Replace Encryption Key dialog.

Creating an encryption key

When you create an encryption key, you select the type of encryption to use.

See [“About encryption key management”](#) on page 418.

To create an encryption key

- 1 On the **Tools** menu, click **Encryption Keys**.
- 2 Click **New**.
- 3 Complete the appropriate options.
See [“Add Encryption Key options”](#) on page 420.
- 4 Click **OK**.

Add Encryption Key options

You have several options when you create an encryption key.

See [“Creating an encryption key”](#) on page 420.

Table 7-10 Add Encryption Key options

Item	Description
Key name	Designates a unique name for this key. The name can include up to 256 characters.
Encryption type	Designates the encryption type to use for this key. Your choices are 128-bit AES or 256-bit AES. The default type is 256-bit AES. The 256-bit AES encryption provides a stronger level of security than 128-bit AES encryption. However, backup jobs may process more slowly with 256-bit AES encryption than with 128-bit AES encryption. Hardware encryption that uses the T10 standard requires 256-bit AES.
Pass phrase	Designates a pass phrase for this key. For 128-bit AES encryption, the pass phrase must be at least eight characters. For 256-bit AES encryption, the pass phrase must be at least 16 characters. Symantec recommends that you use more than the minimum number of characters. You can use only printable ASCII characters. See “About pass phrases in encryption” on page 417.
Confirm pass phrase	Confirms the pass phrase.
Common	Makes this key a common key. If a key is common, any user of this installation of Backup Exec can use the key to back up and restore data.

Table 7-10 Add Encryption Key options (*continued*)

Item	Description
Restricted	Makes the key a restricted key. If a key is restricted, anyone can use the key to back up data. But only the key owner or a user who knows the pass phrase can use the restricted key to restore the encrypted data.

Replacing an encryption key

You can replace one encryption key with another for all backup jobs, templates, and duplicate backup set jobs.

See [“About encryption key management”](#) on page 418.

To replace an encryption key

- 1 On the **Tools** menu, click **Encryption Keys**.
- 2 Select the key that you want to replace.
- 3 Click **Replace**.
- 4 In the **Select an encryption key to replace “key name”** box, do one of the following:

To use an existing key Select the key from the list.

To create a new key Click the arrow, and then click **<new encryption key>**.
See [“Add Encryption Key options”](#) on page 420.

- 5 Click **OK**.

About deleting an encryption key

You should be cautious when you delete encryption keys. When you delete an encryption key, you cannot restore the backup sets that you encrypted with that key unless you create a new key that uses the same encryption key and pass phrase as the original key.

See [“Deleting an encryption key”](#) on page 422.

You can delete encryption keys in the following situations:

- The encrypted data on the tape has expired or if the tape is retired.
- The encryption key is not the default key.

- The encryption key is not being used in a job or a template. If the key is being used, you must select a new key for the job or template.
- The encryption key is not being used in a selection list for restore jobs and for verify duplicate backup set jobs. If you delete a key that is being used in one of the listed job types, the selection list can no longer be used.

If you delete an encryption key that is being used in a scheduled restore job, you cannot replace the key. Therefore, any scheduled restore job in which you delete an encryption key fails.

See [“About encryption key management”](#) on page 418.

See [“Replacing an encryption key”](#) on page 421.

Deleting an encryption key

You should be cautious when you delete encryption keys. When you delete an encryption key, you cannot restore the backup sets that you encrypted with that key unless you create a new key that uses the same encryption key and pass phrase as the original key.

See [“About deleting an encryption key”](#) on page 421.

To delete an encryption key

- 1 On the **Tools** menu, click **Encryption Keys**.
- 2 Select the key that you want to delete.
- 3 Click **Delete**.
- 4 Click **Yes**.
- 5 If the key is used in a job or template, do the following:
 - In the **Select an encryption key to replace "key name"** box, select the new key for the jobs or templates listed.
 - Click **OK**.

About restoring encrypted data

Encrypted backup sets are identified in the restore selection list by an icon with a lock on it. When you select encrypted data to restore, Backup Exec automatically validates the encryption key for the data. If the encryption key that was used to back up the data is still in the Backup Exec database, then Backup Exec selects that encryption key automatically. However, if the encryption key cannot be located, Backup Exec prompts you to provide the pass phrase for the encryption

key that was used to back up the data. If you enter the correct pass phrase, Backup Exec recreates the key.

When restricted encryption keys are used to back up data, any users other than the key owner must enter the pass phrase to restore the data and to edit a restore job.

See [“About pass phrases in encryption”](#) on page 417.

See [“About encryption key management”](#) on page 418.

See [“Replacing an encryption key”](#) on page 421.

About cataloging media that contains encrypted backup sets

When you catalog media that contains encrypted backup sets, Backup Exec attempts to find valid encryption keys for the sets in the Backup Exec database. If Backup Exec does not find a valid key, it issues an alert that instructs you to create one. After you create a valid key, you can respond to the alert to retry cataloging the encrypted set. Alternatively, you can skip the encrypted set and continue to catalog the rest of the media, or cancel the catalog job.

See [“About encryption key management”](#) on page 418.

See [“Creating an encryption key”](#) on page 420.

About configuring DBA-initiated job settings

When you create a DBA-initiated backup operation, you can specify the default job template in Backup Exec. You can also specify a new job template that you create in Backup Exec. The job template contains the settings that Backup Exec applies to DBA-initiated jobs.

Make sure that the name of the job template that you want to use is also configured in the instance information on the Windows computer.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1318.

Note the following about DBA-initiated jobs:

- DBA-initiated jobs fail when the related job template is deleted. To stop DBA-initiated jobs from running, delete the related DBA-initiated job template. See [“Deleting a job template for DBA-initiated jobs”](#) on page 435.
- All DBA-initiated backup and restore jobs are deleted after the jobs have completed.
- You cannot set minimum device requirements for DBA-initiated jobs.

See [“About performing a DBA-initiated backup job for Oracle”](#) on page 1340.

See [“Creating a template for DBA-initiated jobs”](#) on page 424.

See [“Editing DBA-initiated jobs”](#) on page 434.

Creating a template for DBA-initiated jobs

You can create a new job template that Backup Exec applies to DBA-initiated jobs.

See [“About configuring DBA-initiated job settings”](#) on page 423.

See [“Troubleshooting the Oracle Agent”](#) on page 1353.

See [“Deleting a job template for DBA-initiated jobs”](#) on page 435.

To create a template for DBA-initiated jobs

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **DBA-initiated Job Settings**
- 3 Click **New**.
- 4 In the **Properties** pane, under **Backup Job Template**, click **Device and Media**, and then complete the options as appropriate.
See [“Device and media options for DBA-initiated jobs”](#) on page 425.
- 5 In the **Properties** pane, under **Backup Job Template**, click **General**, and then complete the options as appropriate.
See [“General options for DBA-initiated jobs”](#) on page 427.
- 6 In the **Properties** pane, under **Backup Job Template**, click **Network and Security**, and then complete the options as appropriate.
See [“Network and security options for DBA-initiated jobs”](#) on page 429.
- 7 In the **Properties** pane, under **Backup Job Template**, click **Migrator for Enterprise Vault**, and then complete the options as appropriate.
See [“Migrator for Enterprise Vault options”](#) on page 1053.
- 8 If you want Backup Exec to notify someone when the backup job completes, in the **Properties** pane, under **Backup Job Template**, click **Notification**, and then complete the options as appropriate.
See [“Sending a notification when a job completes”](#) on page 682.
- 9 In the **Properties** pane, under **Duplicate Job Template**, click **Settings**, and then complete the options as appropriate.
See [“Duplicate job template settings for DBA-initiated jobs”](#) on page 430.
- 10 Click **OK**.

Device and media options for DBA-initiated jobs

You can configure device and media settings for DBA-initiated jobs.

See [“About configuring DBA-initiated job settings”](#) on page 423.

Table 7-11 Device and media options for DBA-initiated jobs

Item	Description
Device	Indicates the device that you want to use as the default device for jobs.
Use client-side deduplication	<p>Enables a remote computer that is configured as a Remote Agent for Deduplication to deduplicate data and then send only the unique data directly to the deduplication device that is selected in the Device field. This option is selected by default when you select an OpenStorage device or a deduplication storage folder in the Device field.</p> <p>This option is available only if the Deduplication Option is installed and a deduplication device is selected in the Device field.</p> <p>See “About client-side deduplication” on page 1590.</p>
Use media server deduplication	<p>Enables the media server to deduplicate data. Media server deduplication reduces the size of backups, which reduces storage requirements. If the media server does not support deduplication, the data is deduplicated on an intelligent disk device, such as Symantec PureDisk or a device from a third-party vendor.</p> <p>This option is available only if the Deduplication Option is installed and a deduplication device is selected in the Device field.</p> <p>See “About the Deduplication Option” on page 1568.</p>
Media set	Indicates the media set that you want to use as the default media set for jobs.

Table 7-11 Device and media options for DBA-initiated jobs (*continued*)

Item	Description
<p>Overwrite media</p>	<p>Places this backup on an overwritable media. Make sure that appropriate media is in the stand-alone drive or drive pool you select in the Device field in this dialog box.</p> <p>The media in the drive is overwritten if the media is scratch or recyclable (its overwrite protection period has expired). If allocated media or imported media are in the drive, they may also be overwritten depending on the Media Overwrite Protection Level that is set.</p> <p>Depending on your configuration, overwritable media is selected from scratch media or recyclable media.</p> <p>See “How Backup Exec searches for overwritable media” on page 233.</p> <p>If the media in the drive is not overwritable, an alert appears requesting that you insert overwritable media.</p>
<p>Append to media, overwrite if no appendable media is available</p>	<p>Appends this backup to the media set listed in the Media setfield in this dialog box. The backup set is appended if an appendable media is available in the selected media set. If appendable media is not available, an overwritable media is used and added to the media set.</p> <p>If an append job fills a media, the job continues on another piece of overwritable media.</p> <p>If the media in the drive is not overwritable, an alert appears requesting that you insert overwritable media.</p>
<p>Append to media, terminate job if no appendable media is available</p>	<p>Appends this backup to the media set listed in the Media setfield in this dialog box. The backup set is appended if an appendable media is available in the selected media set; if not, the job is terminated.</p>

Table 7-11 Device and media options for DBA-initiated jobs (*continued*)

Item	Description
Eject media after job completes	Ejects the media in the drive when the operation completes.
Retension media before backup	Runs the tape in the drive from beginning to end at a fast speed. Retensioning helps the tape wind evenly and run more smoothly past the tape drive heads. Retensioning is primarily for Mini Cartridge and quarter-inch cartridges and is not supported on most other types of tape drives.
Use Write once, read many (WORM) media	Specifies the use of WORM (write once, read many) media for this backup job. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent. See “About WORM media” on page 247.

General options for DBA-initiated jobs

You can configure general options for DBA-initiated jobs.

See [“About configuring DBA-initiated job settings”](#) on page 423.

Table 7-12 General options for DBA-initiated jobs

Item	Description
Template name	Specifies the name for this backup template. You can accept the default name that appears or enter a name. The name must be unique.
Backup set description	Describes the information in the backup set for future reference.

Table 7-12 General options for DBA-initiated jobs (*continued*)

Item	Description
<p>Compression type</p>	<p>Provides the following compression options:</p> <ul style="list-style-type: none"> <p>■ None.</p> <p>This option copies the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage media space.</p> <p>Hardware data compression should not be used in environments where devices that support hardware compression are used interchangeably with devices that do not have that functionality.</p> <p>In this situation, hardware compression is automatically disabled. You can manually reenable hardware compression on the drives that support it, but this results in media inconsistency. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive.</p> <p>■ Software.</p> <p>This option uses STAC software data compression, which compresses the data before it is sent to the storage device.</p> <p>■ Hardware [if available, otherwise none].</p> <p>This option uses hardware data compression (if the storage device supports it). If the drive does not feature data compression the data is backed up uncompressed.</p> <p>■ Hardware [if available, otherwise software].</p> <p>This option uses hardware data compression (if the storage device supports it). If the drive does not feature hardware data compression, STAC software compression is used.</p>

Table 7-12 General options for DBA-initiated jobs (*continued*)

Item	Description
Verify after backup completes	Performs a verify operation automatically to make sure that the media can be read once the backup has been completed. Verifying all backups is recommended.

Network and security options for DBA-initiated jobs

You can configure network and security options for DBA-initiated jobs.

See [“About configuring DBA-initiated job settings”](#) on page 423.

Table 7-13 Network and security options for DBA-initiated jobs

Item	Description
Network interface	<p>Specifies the name of the network interface card that connects the media server to the network you want to use for the backup network for this backup job. The list includes all available network interfaces on the media server.</p> <p>If you are using the Central Admin Server Option (CASO), select the Use the default network interface for the managed media server option if you want CASO delegated backup jobs to be processed using the network interface card configured as the default in the managed media server.</p>
Protocol	<p>Specifies the protocol you want to use for this backup job.</p> <p>The options are as follows:</p> <ul style="list-style-type: none"> ■ Use any available protocol ■ Use IPv4 ■ Use IPv6
Subnet	Displays the 32-bit number that determines the subnet to which the network interface card belongs.

Table 7-13 Network and security options for DBA-initiated jobs (*continued*)

Item	Description
Allow use of any available network interface, subnet, or protocol for remote agents not bound to the above network interface, subnet, or protocol	<p>Ensures that the data from the remote system is backed up or restored over any available network if the remote system that you selected for backup or restore is not part of the specified backup network.</p> <p>If you do not select this check box and you selected data from a remote system that is not part of the specified backup network, the job fails because Backup Exec cannot back up or restore the data from the remote system.</p>
Interface details	<p>Displays the Media Access Control (MAC) address, Adapter type, Description, IP addresses, and subnet prefixes of the network interface that you selected for the backup network.</p>
Encryption type	<p>Specifies the encryption key you want to use.</p>
Encryption key	<p>Specifies the encryption key you want to use.</p>
Manage keys	<p>Lets you create a new encryption key or delete an existing encryption key.</p>

Duplicate job template settings for DBA-initiated jobs

You can configure duplicate job template settings for DBA-initiated jobs.

See [“About configuring DBA-initiated job settings”](#) on page 423.

Table 7-14 Duplicate job template settings for DBA-initiated jobs

Item	Description
Enable settings to duplicate backup sets for this job	<p>Enables the settings for a duplicate backup set template.</p>
Device	<p>Indicates the device that you want to use as the default device for jobs.</p>
Media set	<p>Indicates the media set that you want to use as the default media set for jobs.</p>

Table 7-14 Duplicate job template settings for DBA-initiated jobs (*continued*)

Item	Description
<p>Overwrite media</p>	<p>Places this backup on an overwritable media. Make sure that appropriate media is in the stand-alone drive or drive pool you select in the Device field in this dialog box.</p> <p>The media in the drive is overwritten if the media is scratch or recyclable (its overwrite protection period has expired). If allocated or imported media are in the drive, they may also be overwritten depending on the Media Overwrite Protection Level that is set.</p> <p>Depending on your configuration, overwritable media is selected from scratch media or recyclable media.</p> <p>See “How Backup Exec searches for overwritable media” on page 233.</p> <p>If the media in the drive is not overwritable, an alert appears requesting that you insert overwritable media.</p>
<p>Append to media, overwrite if no appendable media is available</p>	<p>Appends this backup to the media set listed in the Media set field in this dialog box. The backup set is appended if an appendable media is available in the selected media set; if not, an overwritable media is used and added to the media set.</p> <p>If an append job fills a media, the job continues on another piece of overwritable media.</p> <p>If the media in the drive is not overwritable, an alert appears requesting that you insert overwritable media.</p>
<p>Append to media, terminate job if no appendable media is available</p>	<p>Appends this backup to the media set listed in the Media set field in this dialog box. The backup set is appended if an appendable media is available in the selected media set; if not, the job is terminated.</p>
<p>Eject media after job completes</p>	<p>Ejects the media in the drive when the operation completes.</p>

Table 7-14 Duplicate job template settings for DBA-initiated jobs (*continued*)

Item	Description
Retension media before backup	Runs the tape in the drive from beginning to end at a fast speed, which helps the tape wind evenly and run more smoothly past the tape drive heads. Retensioning is primarily for Mini Cartridge and quarter-inch cartridges and is not supported on most other types of tape drives.
Use Write once, read many (WORM) media	Specifies the use of WORM (write once, read many) media for this backup job. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent. See “About WORM media” on page 247.
Enable DirectCopy to tape	Enables Backup Exec to coordinate the movement of data from a virtual device directly to a physical device. The Backup Exec media server records information about the data in the catalog. Therefore, you can restore data from either the virtual device or the physical device. See “How to copy data directly from a virtual tape library to a physical tape device” on page 381.
Encryption type	Specifies the type of encryption you want to use, if any. See “About encryption” on page 415.
Encryption key	Specifies the encryption key you want to use.
Manage keys	Allows you to create a new encryption key or delete an existing encryption key.
Preferred source device	Specifies the preferred source device that you want to use as the default device for jobs.

Table 7-14 Duplicate job template settings for DBA-initiated jobs (*continued*)

Item	Description
<p>Compression type</p>	<p>Provides the following compression options:</p> <ul style="list-style-type: none"> <p>■ None.</p> <p>This option copies the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage media space.</p> <p>Hardware data compression should not be used in environments where devices that support hardware compression are used interchangeably with devices that do not have that functionality.</p> <p>In this situation, hardware compression is automatically disabled. You can manually re-enable hardware compression on the drives that support it, but this results in media inconsistency. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive.</p> <p>■ Software.</p> <p>This option uses STAC software data compression, which compresses the data before it is sent to the storage device.</p> <p>■ Hardware [if available, otherwise none].</p> <p>This option uses hardware data compression (if the storage device supports it). If the drive does not feature data compression the data is backed up uncompressed.</p> <p>■ Hardware [if available, otherwise software].</p> <p>This option uses hardware data compression (if the storage device supports it). If the drive does not feature hardware data compression, STAC software compression is used.</p>

Table 7-14 Duplicate job template settings for DBA-initiated jobs (*continued*)

Item	Description
Verify after backup completes	Performs a verify operation automatically to make sure the media can be read once the backup has been completed. Verifying all backups is recommended.

Editing DBA-initiated jobs

You can edit the job template settings that Backup Exec applies to DBA-initiated jobs.

See [“About configuring DBA-initiated job settings”](#) on page 423.

To edit DBA-initiated job settings for Oracle

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **DBA-initiated Job Settings**.
- 3 Select the job template that you want to edit.
- 4 Click **Edit**.
- 5 In the **Properties** pane, under **Backup Job Template**, click **Device and Media**, and then edit the options as appropriate.
 See [“Device and media options for DBA-initiated jobs”](#) on page 425.
- 6 In the **Properties** pane, under **Backup Job Template**, click **General**, and then edit the options as appropriate.
 See [“General options for DBA-initiated jobs”](#) on page 427.
- 7 In the **Properties** pane, under **Backup Job Template**, click **Network and Security**, and then edit the options as appropriate.
 See [“Network and security options for DBA-initiated jobs”](#) on page 429.
- 8 In the **Properties** pane, under **Backup Job Template**, click **Migrator for Enterprise Vault**, and then edit the options as appropriate.
 See [“Migrator for Enterprise Vault options”](#) on page 1053.
- 9 In the **Properties** pane, under **Backup Job Template**, click **Notification**, and then edit the options as appropriate.
 See [“Sending a notification when a job completes”](#) on page 682.

- 10 In the **Properties** pane, under **Duplicate Job Template**, click **Settings**, and then edit the options as appropriate.
See [“Duplicate job template settings for DBA-initiated jobs”](#) on page 430.
- 11 Click **OK**.

Deleting a job template for DBA-initiated jobs

The job template contains the settings that Backup Exec applies to DBA-initiated jobs.

See [“About configuring DBA-initiated job settings”](#) on page 423.

To delete a job template for DBA-initiated jobs for Oracle

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **DBA-initiated Job Settings**
- 3 Select the job template that you want to delete.
- 4 Click **Delete**.
- 5 Click **OK**.

About preferred server configurations

Preferred server configurations are collections of one or more servers and sites that you select as preferred backup sources. Preferred server configurations take priority as backup sources in instances where database copies are replicated between multiple servers. You can create preferred server configurations for Microsoft Exchange Database Availability Groups (DAG).

You do not have to create a preferred server configuration to back up replicated database copies. You can let Backup Exec choose the best server from which to back up the replicated database copies. Designating a preferred server configuration gives you more control over your backup jobs. For example, you can select a local preferred server configuration to avoid having to back up replicated data over your WAN.

Backup Exec automatically includes the children of any site or DAG that you select as part of the preferred server configuration. So if you want to ensure that a backup is performed locally, you can select the local site as the preferred server configuration. Backup Exec selects from any of the local servers that belong to that site during the backup job. If you want to ensure that a specific server is used for the backup, select only that server as the preferred server configuration.

See [“Creating preferred server configurations”](#) on page 436.

See [“Deleting preferred server configurations”](#) on page 438.

See [“Editing settings for preferred server configurations”](#) on page 438.

See [“Designating a default preferred server configuration”](#) on page 438.

Creating preferred server configurations

You can create preferred server configurations for Microsoft Exchange Database Availability Groups. Preferred server configurations give you more control over your backup jobs by letting you specify a preferred server from which Backup Exec backs up replicated data.

See [“About preferred server configurations”](#) on page 435.

To create preferred server configurations

- 1 On the **Edit** menu, click **Manage Preferred Servers**.
- 2 Click **New**.
- 3 Complete the appropriate options.
See [“Preferred Servers backup options”](#) on page 437.
- 4 On the **Backup Preferred Server Group** dialog box, click **OK**.
- 5 On the **Manage Preferred Servers** dialog box, click **OK**.

Manage Preferred Servers options

You can manage the settings for preferred servers.

See [“About preferred server configurations”](#) on page 435.

Table 7-15 Manage Preferred Servers options

Item	Description
Name	Indicates the name of the preferred server configuration.
New	Lets you create a new preferred server configuration. See “Creating preferred server configurations” on page 436.
Delete	Deletes the selected preferred server configuration. See “Deleting preferred server configurations” on page 438.

Table 7-15 **Manage Preferred Servers options** (*continued*)

Item	Description
Edit	<p>Lets you change settings for the selected preferred server configuration.</p> <p>See “Editing settings for preferred server configurations” on page 438.</p>
Set as Default	<p>Lets you establish the selected preferred server configuration as the default.</p> <p>See “Designating a default preferred server configuration” on page 438.</p>
Remove Default	<p>Removes the default status for the selected preferred server configuration.</p> <p>See “Removing the default status for a preferred server configuration” on page 439.</p>

Preferred Servers backup options

You can configure settings for preferred servers for backup jobs.

See [“About preferred server configurations”](#) on page 435.

Table 7-16 **Preferred Servers backup options**

Item	Description
Preferred servers configuration	<p>Indicates the name of the preferred server configuration.</p>
New	<p>Lets you create a new preferred server configuration. This option enables the lists of available and selected servers from which you designate the preferred server.</p> <p>Note: The New option appears only if you create a preferred server configuration while you create a new backup job or selection list.</p>
Available Servers and Sites	<p>Lists any available servers and sites that can be used in the preferred server configuration.</p>
Selected Servers and Sites	<p>Lists the servers and sites that you have selected to use as part of the preferred server configuration.</p>

Deleting preferred server configurations

You can delete a preferred server configuration if you no longer need it.

See [“About preferred server configurations”](#) on page 435.

To delete preferred server configurations

- 1 On the **Edit** menu, click **Manage Preferred Servers**.
- 2 Select the preferred server configuration you want to delete.
- 3 Click **Delete**.
- 4 Click **OK**.

Editing settings for preferred server configurations

You can edit the settings for an existing preferred server configuration.

See [“About preferred server configurations”](#) on page 435.

To edit settings for preferred server configurations

- 1 On the **Edit** menu, click **Manage Preferred Servers**.
- 2 Select the preferred server configuration you want to edit.
- 3 Click **Edit**.
- 4 Complete the appropriate options.
See [“Preferred Servers backup options”](#) on page 437.
- 5 On the **Backup Preferred Server Group** dialog box, click **OK**.
- 6 On the **Manage Preferred Servers** dialog box, click **OK**.

Designating a default preferred server configuration

You can designate a default preferred server configuration for all of your backup jobs that contain the appropriate replication data. When you back up data from a Microsoft Exchange Database Availability Group, you can set up Backup Exec to use your default preferred server configuration. You can override the default preferred server configuration for specific jobs in the backup job or selection list properties.

Note: When you designate a default preferred server configuration, it is not applied to existing selection lists. It is considered the default preferred server configuration for any subsequent selection lists you create.

See [“About preferred server configurations”](#) on page 435.

See [“Creating a backup job by setting job properties”](#) on page 332.

If you no longer want the preferred server configuration to be the default, you can remove its default status.

See [“Removing the default status for a preferred server configuration”](#) on page 439.

To designate a default preferred server configuration

- 1 On the **Edit** menu, click **Manage Preferred Servers**.
- 2 Select the preferred server configuration you want to designate as the default.
- 3 Click **Set as Default**.
- 4 Click **OK**.

Removing the default status for a preferred server configuration

You can designate a default preferred server configuration for all of your backup jobs that contain the appropriate replication data.

See [“Designating a default preferred server configuration”](#) on page 438.

If you no longer want the preferred server configuration to be the default, you can remove its default status.

To remove the default status for a preferred server configuration

- 1 On the **Edit** menu, click **Manage Preferred Servers**.
- 2 Select the preferred server configuration from which you want to remove the default status.
- 3 Click **Remove Default**.
- 4 Click **OK**.

About devices

This chapter includes the following topics:

- [About storage devices](#)
- [About the Configure Devices Assistant](#)
- [About sharing storage](#)
- [Pausing a media server](#)
- [Resuming a media server](#)
- [Pausing storage devices](#)
- [Resuming storage devices](#)
- [Renaming storage devices](#)
- [About inventorying media](#)
- [Inventorying media in a device](#)
- [Erasing media](#)

About storage devices

Device management in Backup Exec simplifies how you organize and allocate the storage devices recognized by Backup Exec, including the following:

- Tape drives or robotic libraries physically attached to a media server.
- Virtual tape libraries, which Backup Exec treats as physical robotic libraries.
- Backup-to-disk folders, which are storage devices that you create.
- Shared devices used in a SAN or CASO environment.

- Removable storage devices shared by applications through the use of Microsoft's Removable Storage Feature.
- Simulated tape libraries that you create with the Symantec Tape Library Simulator Utility for the Remote Media Agent for Linux Servers.
- Storage arrays that you configure with the Backup Exec Storage Provisioning Option.
- Deduplication storage folders and OpenStorage devices that you can configure when you install the Deduplication Option.

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

You can use the storage devices as they are configured by Backup Exec at installation, without making any changes. The default device pool, **All Devices** is the default destination device when you create a backup job. By default, the **All Devices** device pool contains all locally attached devices. Devices and simulated tape libraries that are on any computers on which the Remote Media Agent for Linux Servers is installed are excluded from the **All Devices** device pool. Deduplication storage folders and OpenStorage devices are also excluded from the **All Devices** device pool. If you have installed Backup Exec for the first time, along with the Storage Provisioning Option, the **All Virtual Disks** device pool is the default destination device pool. The **All Virtual Disks** device pool contains all virtual disks on all storage arrays.

If you have installed the SAN Shared Storage Option, both locally attached and shared storage devices appear in **All Devices (Computer Name)**. If you have installed the Backup Exec NDMP Option, you can add an NDMP server as a storage device. If you have installed the Backup Exec Storage Provisioning Option, storage arrays and their components also appear.

In addition to device pools, Backup Exec provides other device management capabilities.

You can do the following:

- Identify and monitor the current status of all storage devices.
- Change physical tape devices without rebooting the Backup Exec server.
- Monitor device usage statistics and track hardware errors. Backup Exec keeps track of the device's age, hours of use, mounts, number of bytes processed (written and read), errors, when the device was last cleaned, and so on.
- Manage the physical devices attached to the media server and perform operations on these devices and the media contained in them.

Note: Most of the benefits derived from Backup Exec’s device management functionality are realized when using more than one storage device. However, users with only one device can still take full advantage of Backup Exec’s device monitoring to help make sure their devices are working properly.

See [“About the Configure Devices Assistant”](#) on page 443.

About the Configure Devices Assistant

Use the **Configure Devices Assistant** to configure devices.

See [“Configuring storage devices by using the Configure Devices Assistant”](#) on page 444.

Table 8-1 Configure Devices Assistant options

Device	More information
Device Pool	See “About device pools” on page 515.
Tape Devices	See “About tape drives and robotic libraries” on page 451.
Backup-To-Disk Folder	See “About backup-to-disk folders ” on page 496.
Removable Backup-To-Disk Folder	See “About backup-to-disk folders ” on page 496.
Deduplication Storage Folder	See “About deduplication storage folders” on page 1579.
OpenStorage	See “About OpenStorage devices” on page 1574.
Remote Agent Deduplication	See “Configuring a Remote Agent for Deduplication” on page 1592.
Symantec Protection Network	See “About Symantec Online Storage for Backup Exec” on page 2051.
NDMP Storage	See “About the NDMP Option” on page 1847.
Remote Media Agent Storage	See “About the Remote Media Agent for Linux Servers” on page 1966.
Storage Array	See “About the Storage Provisioning Option” on page 2018.

Table 8-1 Configure Devices Assistant options (*continued*)

Device	More information
Vault Store	See “About vault stores in the Archiving Option” on page 1445.

Configuring storage devices by using the Configure Devices Assistant

Use the following steps to configure devices by using the Configure Devices Assistant.

See [“About the Configure Devices Assistant”](#) on page 443.

To configure storages devices by using the Cofigure Devices Assistant

- 1 On the navigation bar, click **Devices**.
- 2 In the task pane, under **Devices Tasks**, click **Configure devices assistant**.
- 3 Click the type of device that you want to configure.
- 4 Set device options as appropriate for each type of device you configure.

About sharing storage

In environments that use the Backup Exec Central Admin Server Option (CASO) or the SAN Shared Storage Option (SSO), media servers can share storage. In an SSO environment or a CASO environment, Backup Exec maintains a database of the shared devices. Otherwise, the backup data that one server submits to the device can overwrite the data that another server submits. In a CASO environment, you can add a device to a central administration server, a managed media server, or both. Multiple media servers in a CASO environment can share a device.

Media servers can share the following types of storage:

- Devices that are attached to an NDMP server
- Deduplication storage folders
- OpenStorage devices
- Remote Media Agents
- The remote agents that are configured to send data directly to a storage device

Note: You can also share backup-to-disk devices. However, the process is different for backup-to-disk folders.

See [“Sharing an existing backup-to-disk folder”](#) on page 506.

When you add a storage device that supports sharing, you can select which media servers can access the device. The media server from which you added the storage device is enabled to share the device automatically. However, you can remove the sharing capability from that media server at any time. For example, if you add a storage device to a central administration server, then that server can use the device. However, if your environment does not allow the central administration server to operate as a managed media server, then you can remove the sharing capability from the central administration server.

If you have multiple media servers and storage devices in your environment, you can select a media server and manage the storage for it. You can enable and disable the storage devices that you want the media server to use.

See [“Managing shared storage”](#) on page 445.

See [“Sharing the devices on an NDMP server between multiple media servers”](#) on page 1850.

See [“Sharing a deduplication device between multiple media servers”](#) on page 1589.

See [“Sharing a Remote Media Agent between multiple media servers”](#) on page 1977.

Managing shared storage

You can set up a media server to access multiple storage devices.

See [“About sharing storage”](#) on page 444.

To manage shared storage

- 1 On the navigation bar, click **Devices**.
- 2 Right-click a media server.
- 3 Select **Manage Shared Storage**.
- 4 In **Media server**, select the media server for which you want to share storage.
- 5 Select each storage device that you want to use with the selected media server.
- 6 Click **OK**.

Manage Media Server Shared Storage options

You can set up a media server to access multiple storage devices.

See [“Managing shared storage”](#) on page 445.

Table 8-2 Manage Media Server Shared Storage options

Item	Description
Media server	Indicates the name of the media server for which storage is shared.
Storage	Indicates the name of the storage device.
Type	Indicates the type of storage device.

Pausing a media server

You can pause a media server to prevent scheduled and new jobs from running on its devices while maintenance activities are performed. Active jobs are not affected if they start before the media server is paused.

The status **Paused** appears next to the media server name in the **Devices** view if it is currently paused.

To pause a media server

- 1 On the navigation bar, click **Devices**.
- 2 Click the media server.
- 3 Under **General Tasks** in the task pane, select the **Pause** check box.

Resuming a media server

If a media server is paused, you can resume it.

The status **Paused** appears next to the device name if it is currently paused.

To resume a media server

- 1 On the navigation bar, click **Devices**.
- 2 Click the server icon of the server that is paused.
- 3 Under **General Tasks** in the task pane, select the **Pause** check box to uncheck it.

Pausing storage devices

You can pause a storage device to prevent scheduled and new jobs from running on that device while maintenance activities are performed. Active jobs are not affected if they start before the device is paused.

The status **Paused** appears next to the device name in the **Devices** view if it is currently paused.

To pause storage devices

- 1 On the navigation bar, click **Devices**.
- 2 Click the storage device icon.
- 3 Under **General Tasks** in the task pane, select the **Pause** check box.

Resuming storage devices

If a storage device is paused, you can resume it.

The status **Paused** appears next to the device name if it is currently paused.

To resume storage devices

- 1 On the navigation bar, click **Devices**.
- 2 Click the storage device icon of the device that is paused.
- 3 Under **General Tasks** in the task pane, select the **Pause** check box to uncheck it.

Renaming storage devices

You can rename the media server's storage devices.

Backup-to-disk folder names must not exceed 128 characters. The backup-to-disk path name, which includes the backup-to-disk folder name, must not exceed 512 characters. When you use the Backup Exec **Rename** option to rename a backup-to-disk folder, the name changes in Backup Exec, but not on the disk.

You can also change the name of the Windows folder in Windows Explorer.

See [“Changing the path of a backup-to-disk folder”](#) on page 506.

The default **All Devices** device pool cannot be renamed, but you can rename any user-created device pool using either the **Rename** option or the device pool's **Properties** dialog box.

To rename storage devices

- 1 On the navigation bar, click **Devices**.
- 2 Click the storage device that you want to rename.
- 3 Under **General Tasks** in the task pane, click **Rename**.
- 4 Type the new name, and then click **OK**.

About inventorying media

You should run an inventory operation when Backup Exec is started for the first time following a new installation or a product upgrade. When Backup Exec is exited and restarted, it saves information pertaining to the location and contents of all the media from the last Backup Exec session (provided the media in the devices hasn't changed). With this information, Backup Exec can immediately begin processing operations when it is restarted.

When media is changed in a robotic library, you can inventory all of the slots in the robotic library or select the slots to be inventoried. You are not required to re-inventory slots when adding media requested by Backup Exec.

For example, if you are performing a restore operation, and the data is contained on media that is not currently in the robotic library, you are prompted to insert the media for the restore operation. In this case, you are not required to re-inventory the slot where the restore source media is placed.

When media that is not requested by Backup Exec is added or removed from the magazine, you should perform an inventory operation on the changed slots. This updates the media databases so Backup Exec doesn't load and unload each media in the magazine searching for the correct media on which to process jobs. You can select specific slots to inventory. If you swap media often you may want Backup Exec to perform an inventory on the robotic library magazine each time Backup Exec services are started.

Inventorying media in a device

Run **Inventory** to mount the media in the device and read the media label, which is then displayed in the **Devices** view. If this is the first time that Backup Exec has encountered this media, the media label is also added to the **Media** view.

If you change the media in the robotic library or device, run **Inventory** so that the current media in the device is displayed in the views; otherwise, the previous media is still displayed as being in the device.

There may be a delay (up to several minutes for some drives) as the media is mounted and inventoried.

The inventory operation can be monitored or canceled through the **Job Monitor**.

To inventory media in a device

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.

- 3 If you want to inventory a drive or slot, select the drive or slot containing the media you want to inventory.
- 4 If you want to inventory a backup-to-disk file, do the following in the order listed:
 - Double-click the icon for the computer where the backup-to-disk folder is located.
 - Click the backup-to-disk folder that contains the file you want to inventory.
 - On the **Results** pane, select the file you want to inventory.
- 5 Under **Media Tasks**, in the task pane, select **Inventory**.
- 6 To specify a job name or a job priority, on the **Properties** pane, under **Settings**, click **General**.
See [“General options for utility jobs”](#) on page 482.
- 7 If you want a person or group to be notified when the job completes, on the **Properties** pane, under **Settings**, click **Notification** and select the options you want.
See [“Notification options for jobs”](#) on page 682.
- 8 If you want to run the job now, click **Run Now**. Otherwise, on the **Properties** pane, under **Frequency**, click **Schedule** to set the scheduling options you want to use.
See [“Schedule options”](#) on page 356.

Erasing media

You can erase media by using either Quick erase or Long erase. Not all devices support a long erase; those that do not can only perform a quick erase.

Quick erase writes an indicator at the beginning of the media that makes the data contained on the media inaccessible. For most uses, a Quick erase is sufficient.

Long erase instructs the drive to physically erase the entire media. If you have sensitive information on the media and you want to dispose of it, use Long erase. Running Long erase on media can take several minutes to several hours to complete (depending on the drive and the media capacity).

Quick and Long erase do not change the media label. To change a media label, use Label Media or Rename prior to the Erase operation.

You cannot cancel an Erase operation after it has started; however, you can use Cancel to stop a queued erase operation.

To erase media

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.
- 3 Select the drive or slot containing the media you want to erase.
- 4 Under **Media Tasks** in the task pane, select either **Erase media, quick** or **Erase media, long**.

If the drive does not support a long erase, Erase media, long will not be available.

The following warning is displayed:

"This operation will be performed on the current media in the drive or slot. If the media has been changed since the last inventory operation was performed, the media label in the next dialog may not match the media in the drive or slot selected."

- 5 Click **Yes** to continue.

The media displayed was read during the last inventory operation; the display does not change until another inventory operation occurs. Therefore, if you change media in the slot or drive but did not run Inventory, the media label displayed may not match the actual media in the slot or drive.
- 6 When prompted, click **Yes** to erase the media.
- 7 To specify a job name or a job priority, in the **Properties** pane, under **Settings**, click **General**.

See ["General options for utility jobs"](#) on page 482.
- 8 If you want a person or group to be notified when the job completes, in the **Properties** pane, under **Settings**, click **Notification** and select the options you want.

See ["Notification options for jobs"](#) on page 682.
- 9 If you want to run the job now, click **Run Now**. Otherwise, on the **Properties** pane, under **Frequency**, click **Schedule** to set the scheduling options you want to use.

See ["Schedule options"](#) on page 356.

Managing tape drives and robotic libraries

This chapter includes the following topics:

- [About tape drives and robotic libraries](#)
- [About configuring tape devices by using the Tape Device Configuration Wizard](#)
- [About adding or replacing devices by using the Hot-swappable Device Wizard](#)
- [About installing Symantec tape device drivers](#)
- [Changing the preferred block size, buffer size, buffer count, and high water count for devices](#)
- [Enabling hardware compression for devices](#)
- [Specifying read and write operations on types of media](#)
- [Viewing storage device properties](#)
- [About robotic libraries in Backup Exec](#)
- [About creating utility jobs to help manage devices and media](#)

About tape drives and robotic libraries

When you install Backup Exec, all stand-alone tape drives and robotic libraries that are connected to the media server are automatically recognized. Robotic libraries include virtual tape libraries and simulated tape libraries. A stand-alone drive is a single, locally attached tape drive. The view on the **Devices** tab displays how devices are organized logically into device pools, and how devices are arranged physically on servers.

If you group one or more robotic library slots into partitions, the partition drive pools appear in the **Devices** view under the **Device Pools** icon.

The **Configure Devices Assistant** is available to help you configure storage devices, storage folders, and online storage destinations.

The **Tape Device Configuration Wizard** is available to help you install Symantec tape device drivers, and to correct the robotic library drives that are inadvertently displayed as unknown devices in the Devices view.

When you install Backup Exec, support for the following items is included:

- The first robotic library drive per robotic library.
- Every single-drive virtual tape library.

Support for additional drives is available with the Library Expansion Option and the Virtual Tape Library Unlimited Drive Option.

See “[Configuring storage devices by using the Configure Devices Assistant](#)” on page 444.

See “[About the Library Expansion Option](#)” on page 453.

See “[About the Virtual Tape Library Unlimited Drive Option](#)” on page 452.

See “[About the Tape Library Simulator Utility](#)” on page 1979.

See “[About installing Symantec tape device drivers](#)” on page 455.

See “[About robotic libraries in Backup Exec](#)” on page 467.

See “[About creating utility jobs to help manage devices and media](#)” on page 480.

See “[About devices in the SAN Shared Storage Option](#)” on page 1995.

About the Virtual Tape Library Unlimited Drive Option

When you install Backup Exec, support for every single-drive virtual tape library is included. The Virtual Tape Library Unlimited Drive Option enables support for all additional drives in each virtual tape library.

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/v-269-2>

You can find license information for the Virtual Tape Library Unlimited Drive Option at the following URL:

<http://entsupport.symantec.com/umi/V-269-21>

To install the Virtual Tape Library Unlimited Drive Option, add a license key.

See “[Adding licenses](#)” on page 175.

See “[About Backup Exec’s standard features](#)” on page 112.

About the Library Expansion Option

When you install Backup Exec, support for the first drive in every robotic library is included. The Library Expansion Option enables support for each additional drive in a robotic library.

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/v-269-2>

You can find license information for the Library Expansion Option at the following URL:

<http://entsupport.symantec.com/umi/V-269-21>

To install the Library Expansion Option, add a license key.

See “[Adding licenses](#)” on page 175.

See “[About Backup Exec’s standard features](#)” on page 112.

About configuring tape devices by using the Tape Device Configuration Wizard

Use the **Tape Device Configuration Wizard** to do the following:

- Configure robotic library drives to correct where the drives display in the **Devices** view.
After you install Backup Exec, a stand-alone drive may be inadvertently displayed as an unknown device in the **Devices** view.
- Install Symantec tape device drivers by using the **Symantec Device Driver Installation Wizard**.

See “[Configuring storage devices by using the Configure Devices Assistant](#)” on page 444.

See “[About installing Symantec tape device drivers](#)” on page 455.

About adding or replacing devices by using the Hot-swappable Device Wizard

Use the **Hot-swappable Device Wizard** to replace or add a hot-swappable storage device on a Backup Exec media server without having to reboot the server.

If you remove and then reconnect Universal Serial Bus (USB) tape devices to the USB port, you must run the **Hot-swappable Device Wizard** to allow Backup Exec to rediscover the devices.

For iSCSI-attached devices, you must list the device as a **Persistent Target** in the iSCSI control panel applet, and then run the **Hot-swappable Device Wizard**. Listing the device as a **Persistent Target** lets Backup Exec rediscover the device whenever you restart the media server.

After you start the **Hot-swappable Device Wizard**, you are prompted to close the Backup Exec Administration Console. The **Hot-swappable Device Wizard** waits until any jobs that were processing are completed. The wizard pauses the media server and stops the Backup Exec services. You can then add or replace any storage devices. The wizard detects the new or replaced device, and adds information about the device to the Backup Exec Database. The wizard is then completed, and you can reopen the Backup Exec Administration Console.

Any new storage device is displayed in the **Devices** view, and usage statistics for the device begin accumulating. You can add the new device to a device pool.

Any replaced storage device is displayed in the **Devices** view with a status of **Offline**.

See [“Adding or replacing devices by using the Hot-swappable Device Wizard”](#) on page 454.

Adding or replacing devices by using the Hot-swappable Device Wizard

Use the **Hot-swappable Device Wizard** to add or replace a hot-swappable storage device on a Backup Exec media server. You do not need to restart the media server.

See [“About adding or replacing devices by using the Hot-swappable Device Wizard”](#) on page 453.

Note: Start the **Hot-swappable Wizard** before you add or replace storage devices.

Adding or replacing devices by using the Hot-swappable Device Wizard

1 Do one of the following:

- | | |
|--------------------------------------|--|
| For iSCSI-attached devices: | In the iSCSI control panel applet, add the device to the Persistent Targets list.
Continue with the next step. |
| For any other hot-swappable devices: | Continue with the next step. |

2 On the navigation bar, click **Devices**.

3 In the task pane, click **Wizards**> **Hot-swappable Device Wizard**.

4 Follow the on-screen prompts.

About installing Symantec tape device drivers

Use the **Symantec Device Driver Installation Wizard** to install Symantec tape device drivers.

Before you install Symantec tape device drivers, do the following:

- Ensure that the tape device is supported by Backup Exec.
You can find a list of compatible devices at the following URL:
<http://entsupport.symantec.com/umi/V-269-2>
- Run the Windows Device Manager to ensure that it lists the tape device.

See “Installing Symantec tape device drivers by running `tapeinst.exe`” on page 455.

See “Installing Symantec tape device drivers by using the Tape Device Configuration Wizard” on page 456.

Installing Symantec tape device drivers by running `tapeinst.exe`

You can install Symantec tape device drivers by running `tapeinst.exe`, located in the Backup Exec installation directory. Updates for `tapeinst.exe` are available in the **Device Driver Installer** package.

You can download the **Device Driver Installer** package from the following URL:
<http://go.symantec.com/support/BEWS-downloads-drivers>

Note: You must run `tapeinst.exe` locally at the media server where you want to install tape device drivers. You cannot use `tapeinst.exe` to push-install tape device drivers to remote media servers.

To install Symantec tape device drivers by running tapeinst.exe

- 1 From the Backup Exec installation directory, double-click the tapeinst.exe file.

The default installation directory is C:\Program Files\Symantec\Backup Exec.

- 2 On the **Symantec Device Driver Installation Wizard**, follow the on-screen prompts.

Installing Symantec tape device drivers by using the Tape Device Configuration Wizard

You can install Symantec tape device drivers by using the **Tape Device Configuration Wizard** to run the **Device Driver Installation Wizard**.

To configure storage devices by using the Tape Device Configuration Wizard

- 1 On the navigation bar, click **Devices**.
- 2 In the task pane, click **Wizards > Tape Device Configuration Wizard**.
- 3 On the **Welcome** panel, click **Next**.
- 4 On the **Review Backup Devices** panel, click **Next**.
- 5 On the **Create and Configure Backup Devices** panel, click **Install tape device drivers**, and then click **Next**.
- 6 On the **Symantec Device Driver Installation Wizard**, follow the on-screen prompts.

Changing the preferred block size, buffer size, buffer count, and high water count for devices

Caution: Use the preferred configuration settings for a device to tune the performance of backup and restore operations. Changing preferred configuration settings is not generally recommended and may have a negative effect on backup performance and system performance. Thoroughly test any changes before you put them into general use to ensure that system performance does not deteriorate.

To change the preferred block size, buffer size, buffer count, and high water count for devices

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.

- 3 Select the device for which you want to view properties.
- 4 Under **General Tasks** in the task pane, select **Properties**, and then on the **Drive Properties** dialog box, click **Configuration**.
- 5 Click the drop-down menu for the item that you want to change, and then select a new setting.
 See [“Configuration properties for devices”](#) on page 460.
- 6 Click **OK**.

Enabling hardware compression for devices

You can enable or disable hardware compression for a device if the device supports compression.

To enable hardware compression for devices

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.
- 3 Select the device for which you want to view properties.
- 4 Under **General Tasks** in the task pane, select **Properties**, and then on the **Drive Properties** dialog box, click **Configuration**.
- 5 Click the drop-down menu for the item that you want to change, and then select a new setting.
 See [“Configuration properties for devices”](#) on page 460.
- 6 Click **Enable compression**.
- 7 Click **OK**.

Specifying read and write operations on types of media

You can specify that a device is limited to performing read and write operations on specific media types. This information is then incorporated into the device and media database, allowing Backup Exec to exclude this media type when searching for media to be used for a job that requires writing to the media.

To specify read and write operations on types of media

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.

- 3 Select the device for which you want to specify read and write operations on types of media.
- 4 Under **General Tasks** in the task pane, select **Properties**, and then on the **Drive Properties** dialog box, click **Media Type**.
See “[Media type properties for devices](#)” on page 466.
- 5 Change the **Read** and **Write** check boxes as appropriate.
See “[Bar code rules in mixed media libraries](#)” on page 245.

Viewing storage device properties

Properties provide detailed information about storage devices, such as statistics, dates, and settings.

To view storage device properties

- 1 Select the item for which you want to view properties, and then in the task pane under **General Tasks**, click **Properties**.
- 2 To view general properties, click **General**.
See “[General properties for devices](#)” on page 458.
- 3 To view configuration properties, click **Configuration**.
See “[Configuration properties for devices](#)” on page 460.
- 4 To view SCSI information properties, click **SCSI Information**.
See “[SCSI information for devices](#)” on page 463.
- 5 To view statistics properties, click **Statistics**.
See “[Statistics properties for devices](#)” on page 463.
- 6 To view cleaning properties, click **Cleaning**.
See “[Cleaning properties for devices](#)” on page 465.
- 7 To view media type properties, click **Media Type**.
See “[Media type properties for devices](#)” on page 466.

General properties for devices

General properties for devices include information about the status, type, and vendor of the device.

See “[Viewing storage device properties](#)” on page 458.

Table 9-1 General properties for devices

Item	Description
Name	Displays the name of the device.
Status	<p>Displays any of the following statuses:</p> <ul style="list-style-type: none"> ■ Pause. Indicates if this device is paused. ■ Enable. Indicates if Backup Exec has exclusive use of this device. If the check box is clear, the device is disabled and cannot be used by Backup Exec. The device is available for other applications. ■ Online. Indicates that the device is online if a dimmed check box with a check mark appears. If the device is offline, a check mark does not appear. No operations are allowed on the device until it is online again. The device appears as offline if the following occurs: <ul style="list-style-type: none"> ■ The device was turned off after Backup Exec was started. ■ The device was being used by another application (such as a Windows backup utility) when Backup Exec was started. ■ The device is removed from the computer. ■ If a device reports a critical error. ■ The firmware of the device was updated; Backup Exec will behave as if the device with its old name or identity no longer exists. <p>See “Changing the status of a device to online” on page 508.</p>
Vendor	Displays the name of the vendor of the drive or robotic library.
Product ID	Displays the product ID from the SCSI Inquiry string.
Firmware	Displays the version of the firmware used in the device.

Table 9-1 General properties for devices (*continued*)

Item	Description
Library type	<p>Displays the default first slot of the robotic library. Virtual tape libraries are identified by the string VTL. The simulated tape libraries that the Tape Library Simulator Utility creates are identified by the string TLS.</p> <p>See “About the Tape Library Simulator Utility” on page 1979.</p>
Media type	<p>Displays the media type used in this drive type.</p>
Date in service	<p>Displays the date this device was first detected by this installation of Backup Exec.</p>
Serial number	<p>Displays the serial number of the drive.</p>
Encryption	<p>Displays whether the tape device is currently capable of hardware encryption. If the field says Yes and the job is configured to use hardware encryption, Backup Exec uses its included encryption key management to encrypt the data. If the field says No, Backup Exec does not encrypt the data. You can still encrypt the data on the tape, however, if you use hardware encryption from a third-party. Please consult your hardware vendor for encryption key management options and licensing.</p> <p>See “About hardware encryption” on page 416.</p>
WORM	<p>Displays whether the tape device is capable of Write once, read many (WORM) data storage. Backup Exec cannot erase or recycle WORM media.</p> <p>See “About WORM media” on page 247.</p>

Configuration properties for devices

Use the configuration properties for devices to do the following:

- Enable or disable hardware compression (if compression is supported by the drive).
- Change the preferred block size, buffer size, buffer count, and high water count.

Caution: Preferred configuration settings are used to tune the performance of backup and restore operations. Changing preferred configuration settings is not generally recommended and may have a negative effect on your backup and system performance. Any changes should be thoroughly tested to make sure that system performance does not deteriorate.

See “[Viewing storage device properties](#)” on page 458.

Table 9-2 Configuration properties for devices

Item	Description
<p>Enable compression</p>	<p>Indicates if hardware compression is enabled.</p> <p>If this option is available, this device is capable of supporting hardware compression.</p> <p>If a job is configured to use hardware compression, but is run on a device on which hardware compression is disabled (even though it is supported), hardware compression is considered unavailable and is not used.</p>
<p>Block size (per device)</p>	<p>Displays the size of the blocks of data that are written to new media in this device. The default is the preferred block size.</p> <p>Some devices (for example, LTO devices) provide better performance when larger block sizes are used. The preferred block size can range from 512 bytes to 64 kilobytes or larger. If you use a device that supports larger block sizes, you can change the device’s block size in the Device Configuration tab. However, if the option to change the block size is unavailable, you must configure the device to use a larger size.</p> <p>See the device manufacturer’s documentation for help in configuring the device.</p> <p>Backup Exec does not ensure that the requested block size is in fact supported by that device. You should check the device specifications to make sure that the block size is supported. If the device does not support a block size, it will default to its standard block size.</p> <p>If the device does not support block size configuration, this option is unavailable.</p>

Table 9-2 Configuration properties for devices (*continued*)

Item	Description
Buffer size (per device)	<p>Displays the amount of data sent to the device on each read or write request. The buffer size must be an even multiple of the block size.</p> <p>Depending on the amount of memory in your system, increasing this value may improve device performance. Each type of device requires a different buffer size to achieve maximum throughput.</p>
Buffer count	<p>Displays the number of buffers allocated for this device.</p> <p>Depending on the amount of memory in your system, increasing this value may improve device performance. Each type of device requires a different number of buffers to achieve maximum throughput.</p> <p>If you change the buffer count, you may need to adjust the high water count accordingly.</p>
High water count	<p>Displays the number of buffers to be filled before data is first sent to the device, and any time after that if the device underruns.</p> <p>The high water count cannot exceed the buffer count. A value of 0 disables the use of high water logic; that is, each buffer is sent to the device as it is filled.</p> <p>The default setting provides satisfactory performance in most instances; in some configurations, throughput performance may be increased when other values are specified in this field. If you increase or decrease the buffer count, the high water count should be adjusted accordingly. If a device has a high water count default of 0, it should be left at 0.</p>
Default Settings	<p>Returns all of the preferred configuration settings to their defaults.</p>
Read single block mode	<p>Indicates if this device reads only one block of data at a time, regardless of the size of the buffer block.</p>
Write single block mode	<p>Indicates if this device writes only one block of data at a time. This option provides greater control over the handling of data write errors.</p> <p>Symantec recommends this option if the device is a shared storage device.</p>

Table 9-2 Configuration properties for devices (*continued*)

Item	Description
Read SCSI pass-through mode	Indicates if this device reads data without going through a Microsoft tape device API. This option allows the data to pass directly through the device and allows more detailed information if device errors occur.
Write SCSI pass-through mode	Indicates if this device writes data without going through the Microsoft tape device API. This option allows data to pass directly through the device driver and allows more detailed information if device errors occur. Symantec recommends this option if the device is a shared storage device.

SCSI information for devices

SCSI information for a device provides properties of the Small Computer System Interface (SCSI).

See [“Viewing storage device properties”](#) on page 458.

Table 9-3 SCSI information for a device

Item	Description
Inquiry	Displays device information read from the device firmware.
Port	Displays the identifying number of the port on the server to which the device is attached.
Bus	Displays the identifying number of the bus to which the device is attached.
Target ID	Displays the unique SCSI ID number (physical unit number).
LUN	Displays the Logical Unit Number of the device.

Statistics properties for devices

Statistics include the date the device was last mounted, device totals such as the total number of bytes written and read, and device errors. Error rates are affected by media, head cleaning, and head wear. Information includes only the statistics that are gathered after Backup Exec first discovers the device.

The documentation included with your device should list the acceptable limits for hard and soft errors; if not, check with the hardware manufacturer.

See [“Viewing storage device properties”](#) on page 458.

Table 9-4 Statistics properties for devices

Item	Description
Last mount date	Displays the last date that media was mounted by this device.
Total bytes written	Displays the number of bytes that have been written by this device.
Total bytes read	Displays the number of bytes that have been read by this device.
Total mounts	Displays the number of times media has been mounted by this device.
Total seeks	Displays the total number of seek operations (performed when a specific piece of information is being located) that have been performed by this device.
Total hours in use	Displays the total number of hours that this device has been in use (performing read, write, mount, and seek operations).
Seek error	Displays the number of errors encountered while trying to locate data.
Soft read errors	Displays the number of recoverable read errors encountered. If you receive soft errors, it may indicate the beginning of a problem. If you receive excessive errors for your environment, check the device and perform maintenance on it, and check the media for damage.
Hard read errors	Displays the number of unrecoverable read errors encountered. If you receive hard errors, check the device and perform maintenance on it, and check the media for damage.
Soft write errors	Displays the number of recoverable write errors encountered. If you receive soft errors, it may indicate the beginning of a problem. If you receive excessive errors for your environment, check the device and perform maintenance on it, and check the media for damage.
Hard write errors	Displays the number of unrecoverable write errors encountered. If you receive hard errors, check the device and perform maintenance on it, and check the media for damage.

Cleaning properties for devices

Cleaning properties for devices provide statistics on totals and errors since the last cleaning. The documentation included with your device should list the acceptable limits for hard and soft errors; if not, check with the hardware manufacturer.

For robotic library drives, the statistics on the **Cleaning** tab are automatically updated when a cleaning job successfully completes.

If you want to maintain accurate cleaning statistics for your stand-alone drives, you can reset the cleaning statistics after the drive has been manually cleaned.

See [“Viewing storage device properties”](#) on page 458.

See [“Creating a cleaning job”](#) on page 488.

Table 9-5 Cleaning properties for devices

Item	Description
Last cleaning date	Displays the last date a cleaning operation was performed on the device.
Hours since last cleaning	Displays the number of hours that the device has been in use since the last cleaning.
Reset Cleaning Statistics	Resets all cleaning statistics to zero (stand-alone drives only). You cannot undo this operation.
Bytes written	Displays the number of bytes that have been written by this device since the last cleaning.
Bytes read	Displays the number of bytes that have been read by this device since the last cleaning.
Total mounts	Displays the number of times media has been mounted by this device since the last cleaning.
Total seeks	Displays the total number of seek operations that have been performed by this device since the last cleaning. Seek operations are run to locate a specific piece of information .

Table 9-5 Cleaning properties for devices (*continued*)

Item	Description
Hours in use	Displays the total number of hours that this device has been in use since the last cleaning.
Seek errors	Displays the number of seek errors encountered since the last cleaning.
Soft read errors	Displays the number of recoverable read errors encountered since the last cleaning. Soft errors may indicate the beginning of a problem. If excessive errors are reported for your environment, check the device and perform maintenance on it, and check the media for damage.
Hard read errors	Displays the number of unrecoverable read errors encountered since the last cleaning. If you receive hard errors, check the device and perform maintenance on it, and check the media for damage.
Soft write errors	Displays the number of recoverable write errors encountered since the last cleaning. Soft errors may indicate the beginning of a problem. If excessive errors are reported for your environment, check the device and perform maintenance on it, and check the media for damage.
Hard write errors	Displays the number of unrecoverable write errors encountered since the last cleaning. If you receive hard errors, check the device and perform maintenance on it, and check the media for damage.

Media type properties for devices

Media types properties display the types of media that you can place in the device. You can specify the media types to use for read or write operations. By default, all media type categories are listed, and are allowed for use by both read and write operations.

See [“Viewing storage device properties”](#) on page 458.

Backup Exec’s device and media database maintains a list of media types, such as 4mm, and then further defines the subcategories of a media type. For example, a 4mm media type can include 4mm DDS-1 with a length of 60m and the storage capacity of 1.3 GB. Another 4mm tape might also be a 4mm DDS-1 but have a length of 90m and a storage capacity of 2GB.

If you have bar code support for a robotic library that uses different types of drives, you can create a bar code rule so that Backup Exec can identify which media type to use in a drive.

See [“Bar code rules in mixed media libraries”](#) on page 245.

Table 9-6 Media type properties for devices

Item	Description
Media Types	Displays types of media, such as 4mm, and any defined category of this media type, such as CLN for cleaning tape. Media types that have numbers appearing in brackets (for example, 4mm [6]) can be used to define specific bar code rules.
Read	Displays Yes if this media type can be read by the device.
Write	Displays Yes if this media type can be written to by the device.

About robotic libraries in Backup Exec

Backup Exec’s Advanced Device and Media Management (ADAMM) feature provides powerful functionality for robotic libraries. With typical robotic library modules, you divide slots in the robotic library into defined groups, and then target backups to those groups. This arrangement works as long as there is enough media in the group to process the jobs targeted there. Problems occur when the data exceeds the available media in the group, because operations cannot continue until overwritable media is physically added, and you create an import media job to insert media into your robotic library. This situation can take place even though slots in the robotic library assigned to other groups contain usable media.

Backup Exec’s Device and Media Management feature solves the problems associated with typical robotic library modules. Rather than targeting a backup job to a specific group of slots with a finite number of media, Backup Exec accesses all of the media in the robotic library and uses media that belongs to the job’s targeted media set. If the backup job exceeds the capacity of one piece of media, Backup Exec searches all media contained in the robotic library, finds a suitable media, and uses it for the job.

For example, an operator has a robotic library with six slots. The operator inserts six blank tapes and targets backup jobs to various media sets within the robotic library. Depending on whether the backups are overwrite or append jobs, Backup Exec automatically allocates available tapes in the robotic library. If a job exceeds the capacity of one tape and another overwritable tape is available in the robotic library, the job will automatically continue on that tape. When Backup Exec runs out of tapes, it prompts the operator to import overwritable media.

In a robotic library, Backup Exec selects the oldest recyclable media in the library to use first. If more than one media meeting the requirements is found, Backup Exec then selects the media in the lowest-numbered slot; for example, media in slot 2 would be selected before equivalent media in slot 4.

See “[About the Library Expansion Option](#)” on page 453.

See “[About the Virtual Tape Library Unlimited Drive Option](#)” on page 452.

See “[Utility jobs for virtual tape libraries and simulated tape libraries](#)” on page 482.

See “[Requirements for setting up robotic library hardware](#)” on page 468.

See “[Importing media to a robotic library](#)” on page 489.

Requirements for setting up robotic library hardware

You can configure Backup Exec to work with robotic library drives by making associations between the robotic library’s drives, robotic arm, and Backup Exec. Backup Exec supports serialized drives. Manual configuration of serialized drives is not required.

You can find a list of supported devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

Ensure that the robotic library hardware is configured as follows:

- Ensure that the robotic arm is set to Random mode. Refer to your robotic library documentation for more information.
- Ensure the following for a multi-LUN robotic library:
 - The controller card is set to support multiple LUNs (if supported).
 - The target LUN for the tape drive is lower than the target LUN for the changer.
- Determine which drive is the first drive in the robotic library (Storage Device 0), and then arrange the SCSI IDs to match the sequence of the drive element addresses. Refer to your robotic library documentation to determine the drive element address for each storage device.

- Ensure that the SCSI ID of the robotic arm precedes the SCSI IDs of the drives in the robotic library. Do not use 0 or 1 because these SCSI IDs are typically reserved for boot devices.

In the following example, if your robotic library has two drives, the drive with the lowest drive element address should be assigned the lower SCSI ID.

Table 9-7 Example configuration for a multi-drive robotic library

Data Transfer Element (Storage Devices)	SCSI ID	Drive Element Address
Robotic Arm	4	N/A
Storage Device 0	5	00008000
Storage Device 1	6	00008001

See [“Troubleshooting the display of robotic library devices”](#) on page 469.

Troubleshooting the display of robotic library devices

If the robotic library devices are not correctly displayed in the **Devices** view, try the following:

Table 9-8 Troubleshooting the display of robotic library devices

Issue	Action
If a robotic library appears in the backup devices list as a stand-alone drive:	<p>Run the Tape Device Configuration Wizard. In the Configure Library Drives panel, correct the association by clicking and dragging the drive to the appropriate robotic library.</p> <p>See “Configuring storage devices by using the Configure Devices Assistant” on page 444.</p>
If the robotic arm is not shown:	<p>Ensure that you have enabled robotic library support.</p> <p>See “About Backup Exec’s standard features” on page 112.</p>

If you make any changes, run an Inventory operation to update Backup Exec’s media database.

See [“About inventorying media”](#) on page 448.

See [“Requirements for setting up robotic library hardware”](#) on page 468.

Initializing robotic libraries when the Backup Exec service starts

You can enable Backup Exec to initialize a robotic library whenever the Backup Exec service starts.

During startup, if there is media in the storage devices in the robotic library, Backup Exec attempts to return the media to its original magazine slot. If the media cannot be returned to the slot, an error message appears requesting that the media be ejected from the storage device.

You can also create a job to initialize a robotic library.

See [“Creating a job to initialize a robotic library”](#) on page 484.

To initialize robotic libraries when the Backup Exec service starts

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.
- 3 Select a robotic library.
- 4 Under **Robotic Library Tasks** in the task pane, select **Properties**.
- 5 On the **Configuration** tab, click **Enable startup initialization**.
See [“Configuration properties for robotic libraries”](#) on page 471.
- 6 Click **OK**.

Enabling bar code rules for robotic libraries

After you create a bar code rule to specify the types of media that Backup Exec should use in a robotic library drive, you must enable bar code rules for the library.

See [“Bar code labeling”](#) on page 244.

To enable bar code rules for robotic libraries

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.
- 3 Select a robotic library.
- 4 Under **Robotic Library Tasks** in the task pane, select **Properties**.
- 5 On the **Configuration** tab, click **Enable bar code rules**.
- 6 Click **OK**.

Defining a cleaning slot

Before submitting a cleaning job, you must define a cleaning slot that contains the cleaning tape.

Make sure that the cleaning tape is located in the slot that you defined as the cleaning slot. After defining the cleaning slot, you can set up a cleaning job for the robotic library drive.

See [“Creating a cleaning job”](#) on page 488.

Defined cleaning slots are not inventoried when an inventory job runs.

To define a cleaning slot

- 1 On the navigation bar, click **Devices**.
- 2 Click the drive or robotic library for which you are setting up the cleaning.
- 3 Select the slot that contains the cleaning tape.
- 4 Under **General Tasks** in the task bar, select **Properties**.
- 5 Check **Cleaning slot**, and then click **OK**.

Configuration properties for robotic libraries

Configuration properties let you enable a robotic library to be initialized when Backup Exec starts, enable bar code rules, and specify slot base numbering.

See [“Viewing storage device properties”](#) on page 458.

Table 9-9 Configuration properties for robotic libraries

Item	Description
Enable startup initialization	<p>Indicates if Backup Exec initializes the robotic library when the Backup Exec service is started. Depending upon the type of robotic library, initialization can determine which slots have media and can read all bar code labels on media.</p> <p>The default setting is off.</p> <p>You may want to enable this option if the library does not initialize itself when it starts. However, if the library is shared by multiple servers, you should not enable this option since each server must initialize the library. Other servers cannot access the library until all of the initialization processes are complete.</p> <p>If you do not want to initialize the library at startup, you can run an initialization job at any time.</p> <p>See “Creating a job to initialize a robotic library” on page 484.</p>

Table 9-9 Configuration properties for robotic libraries (*continued*)

Item	Description
Enable bar code rules	<p>Indicates if bar code rules are enabled for the robotic library. If you create a bar code rule to specify the type of media that Backup Exec should use in a robotic library drive, you must enable bar code rules for that library before the rules are used.</p> <p>The default setting is off.</p> <p>See “Bar code rules in mixed media libraries” on page 245.</p>
Slot base	<p>Depicts the starting slot for this robotic library. Backup Exec determines what the starting slot should be for this type of library. Some robotic libraries have slots that start at 0. Other libraries start at 1. You can change the starting slot if necessary.</p> <p>See “Reassigning a slot base number for robotic libraries” on page 477.</p>

Statistics properties for robotic libraries

You can view statistics for a robotic library.

See [“Viewing storage device properties”](#) on page 458.

Table 9-10 Statistics properties for robotic libraries

Item	Description
Slot count	Displays the number of slots in the robotic library.
Drive element count	Displays the number of drive elements contained in the robotic library.
Total mounts	Displays the number of times media has been mounted by this device.
Mount errors	Displays the number of errors encountered while mounting media in a drive.

Properties for robotic library slots

You can view information about a slot in the robotic library, and about any media that is in the slot.

See [“Viewing storage device properties”](#) on page 458.

Table 9-11 Properties for robotic library slots

Item	Description
Slot number	Displays the number of the slot.
Bar code	Displays the label obtained from a bar code reader. Bar code information only appears if the robotic library has a bar code reader and a bar code label is on the media.
Cleaning slot	Indicates if this slot has been defined as a cleaning slot. See “Defining a cleaning slot” on page 471.

Table 9-11 Properties for robotic library slots (*continued*)

Item	Description
<p>Media label</p>	<p>Displays the media label as one of the following:</p> <ul style="list-style-type: none"> ■ A label that Backup Exec automatically assigns. ■ A label that the administrator assigns. ■ A pre-assigned bar code label. <p>You can edit the media label, which is limited to 32 characters. Editing the label changes the name of the media in the display, but does not write the new label to the media until an overwrite operation occurs. When you edit a media label, try to make it a concise identifier that will remain constant even when the media is reused. You should write this media label on a label fixed to the outside of the physical media.</p> <p>Duplicate labels can be automatically generated. For example, reinstalling Backup Exec or bringing media from another Backup Exec installation could cause duplication in labels. Duplicate labels are allowed, but not recommended.</p> <p>If a bar code is available, and a bar code-equipped device is used, then the media label automatically defaults to that bar code.</p>
<p>Description</p>	<p>Displays the original media label if the media is imported media.</p> <p>You can edit the media description, which is limited to 128 characters, to make it a more descriptive label.</p>

Table 9-11 Properties for robotic library slots (*continued*)

Item	Description
Media type	Displays the media type and subtype (if available). Click the button next to the field to change the media type or subtype.
Export pending	Displays Yes when a job runs that has an associated Export Media template to export this media. See “ About export media templates ” on page 536.
Media set	Displays the name of the media set this media belongs to.
Media location	Displays the name of the device or vault where this media is located.
Creation date	Displays the date and time when the media was first entered into Backup Exec.
Allocated date	Displays the date and time when the media was added to a media set as a result of an overwrite operation.
Modified date	Displays the date and time when data was last written to the media.
Overwrite protection until	Displays the date and time after which the media can be overwritten.
Appendable until	Displays the date and time after which the media can no longer be appended to.
Supports HW encryption	Displays Yes if this media supports hardware encryption. See “ About hardware encryption ” on page 416.

About robotic library partitions



You can group one or more robotic library slots into partitions. Partitioning robotic library slots provides more control over which media is used for backup jobs.

When you set up robotic library partitions, Backup Exec creates a device pool for each partition. Jobs targeted to a partition device pool run on the media located in the partition's slots. For example, if you set up a partition that contains slots 1 and 2 and you want to run a weekly backup only on the media in these slots, you would submit the job to the partition device pool containing slots 1 and 2.

A partition divider lists the range of slots included in the partition.

Figure 9-1 Configure Partitions dialog box

Partition
Divider



For example, if you want to create two 5-slot partitions on a robotic library with 10 slots, click Slot 1 and Slot 6. In this example, Slots 1-5 will be included in the first partition and Slots 6-10 will be included in the second.

Partitions can include any number of robotic library slots; however, the first partition cannot be moved or deleted when other partitions exist.

Depending upon the robotic library configuration, the first slot could be numbered 1 or 0. If the robotic library uses a zero-based slot configuration and you assign the first partition to begin with slot 1, the Partition Utility will actually use slot 0 as the first slot for partition 1 and adjust the starting slot accordingly for all other partitions.

The partition device pools appear in the **Devices** view under the **Device Pools** icon. If the robotic library is partitioned, Backup Exec searches for the oldest recyclable media in the targeted partition only. If more than one media meeting the requirements is found, Backup Exec then selects the media in the lowest-numbered slot; for example, media in slot 2 would be selected before equivalent media in slot 4.

In order to fully benefit from Backup Exec's partition management feature, Symantec recommends that you create a partitioning scheme that best matches the manner in which you want to control your backups. For example, some administrators may feel that network backups are best managed by allowing access to partitions based on users and groups, while others may want to base their partitions on operation types.

See [“Creating robotic library partitions”](#) on page 477.

See [“Removing robotic library partitions”](#) on page 480.

See [“About redefining robotic library partitions”](#) on page 478.

Creating robotic library partitions

You can create partitions for robotic library slots to control which media is used for backup jobs. After you create the partitions, you can submit jobs to those partitions' device pools.

The partition device pools appear under **Robotic Libraries** for the robotic library on which they were created. All partition device pools for a robotic library have the same name and display the slot ranges for the partition in parentheses within the name.

See [“About robotic library partitions”](#) on page 475.

To create robotic library partitions

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.
- 3 Select the robotic library containing the slots that you want to partition.
- 4 Under **Robotic Library Tasks** in the task pane, select **Configure partitions**.
- 5 Select the robotic library slots to include in each partition by clicking the slot on which each new partition should begin.
- 6 Click **OK** after configuring the partitions.
- 7 Click **Yes** to accept the partitions.

Reassigning a slot base number for robotic libraries

Backup Exec automatically assigns slot base numbers for robotic libraries. If necessary, you can reassign how robotic library slots are displayed in Backup Exec. Slot base numbers in some robotic libraries start at 0, while slots in other robotic libraries start at 1. If the robotic library uses a zero-based slot configuration, you can reassign how the slots are displayed.

To reassign a slot base number for robotic libraries

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.
- 3 Select the robotic library for which you want to reassign the slot base number.
- 4 In the task pane, under **General Tasks**, click **Properties**.
- 5 On the **Configuration** tab, in the **Slot base** field, type the appropriate number.
- 6 Click **OK**.

Configure Partitions options

You can configure new robotic library partitions, remove partitions, or rearrange partitions.

See [“Creating robotic library partitions”](#) on page 477.

See [“Removing robotic library partitions”](#) on page 480.

Table 9-12 Configure Partitions options

Item	Description
Select starting slots for partitions below	Displays the available slots that you can designate as a starting slot for a robotic library partition.
Remove Partition	Removes the selected partition. The slots contained in the partition you are removing are added to the partition preceding it.
Move Up	Moves the selected partition divider up to increase the number of slots in the partition. (The number of slots in the preceding partition is decreased.)
Move Down	Moves the selected partition divider down to decrease the number of slots in the partition. (The number of slots in the preceding partition is increased.)
Remove All	Removes all partition settings.

About redefining robotic library partitions

You can reassign slots to different partitions or even create or delete partitions from a partition drive pool by providing different beginning slot parameters. For example, if your current set-up is a 6-slot robotic library with two partitions (partition 1 = slots 1-3 and partition 2 = slots 4-6), but you want to have three

partitions with slots 1-2 in partition 1, slots 3-5 in partition 2, and slot 6 in partition 3, you would select slots 1, 3, and 6.

Because the first two partition drive pools maintain the same identity, even though the slots have been reassigned, jobs submitted to those partition drive pools will not have to be retargeted.

However, if you change from three partitions to two partitions, any jobs submitted to the third partition must be retargeted since that third partition no longer exists. Also, if you create a new partition that completely contains two or more of the old partitions, jobs submitted to the old partition must be retargeted.

For example, if a robotic library that had been partitioned with the following:

Table 9-13 Robotic library partition example

Partition	Slot
Partition 1	Slots 1 - 2
Partition 2	Slots 3 - 4
Partition 3	Slots 5 - 10

The library is repartitioned as follows:

Table 9-14 Robotic library repartition example

Partition	Slot
Partition 1	Slots 1 - 4
Partition 2	Slots 5 - 6
Partition 3	Slots 7 - 10

Then any jobs targeted to the old partition 2 (slots 3-4) must be retargeted.

Note: If a job is targeted to a particular robotic library drive (or a device pool that is not a partition drive pool), the job defaults to the first partition in the robotic library.

See [“Creating robotic library partitions”](#) on page 477.

See [“Retarget Job options”](#) on page 520.

Removing robotic library partitions

You can remove one or all partitions in a robotic library.

To remove robotic library partitions

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.
- 3 Select the robotic library that contains the partitions that you want to remove.
- 4 Under **Robotic Library Tasks** in the task pane, select **Configure partitions**.
- 5 Do one of the following:

To remove one partition:

Select the partition that you want to remove, and then click **Remove Partition**.

To remove all partitions:

Click **Remove all**.

- 6 Click **OK** after you configure the partitions.
- 7 Click **Yes** to accept the new configuration.

About creating utility jobs to help manage devices and media

Backup Exec includes utility jobs that aid in managing devices and media. You can specify a job priority and a recipient for notification when these jobs run. Utility jobs, which are similar to backup and restore jobs, generate job history records and an XML job log when they run.

Jobs that can be scheduled as recurring jobs are listed in the following table:

Table 9-15 Utility jobs that can be scheduled

Utility job that can be scheduled	For more information
Vault media job	See “Scheduling a job to move media” on page 255.
Catalog media	See “Creating a new catalog” on page 248.
Restore data	See “Creating a restore job while reviewing media or devices” on page 250.
Inventory robotic library or device	See “About inventorying media” on page 448.

Table 9-15 Utility jobs that can be scheduled (*continued*)

Utility job that can be scheduled	For more information
Erase media in a robotic library or device	See “Erasing media” on page 449.
Import media	See “Importing media to a robotic library” on page 489.
Export media	See “Exporting media from a robotic library” on page 490.
Export expired media (robotic libraries only)	See “Exporting expired media from a robotic library” on page 492.
Lock robotic library	See “Locking the robotic library’s front panel” on page 493.
Unlock robotic library	See “Unlocking the robotic library’s front panel” on page 494.
Clean drive	See “Defining a cleaning slot ” on page 471.

Utility jobs that can be created only as run-once jobs, which are jobs that are scheduled to run now or to run once at a specified date and time, are listed in the following table:

Table 9-16 Utility jobs that can run once

Utility jobs that can run once	For more information
Label media	See “Labeling media” on page 486.
Format media, including WORM media	See “Formatting media in a drive” on page 485.
Retension media	See “Retensioning a tape” on page 484.
Eject media	See “Ejecting media from a drive” on page 487.
Initialize robotic library	See “Creating a job to initialize a robotic library” on page 484.

Utility jobs for virtual tape libraries and simulated tape libraries

Backup Exec treats virtual tape libraries and simulated tape libraries as physical robotic libraries. You can identify virtual tape libraries by the label VTL that displays on a library's properties pages. You can identify simulated tape libraries by the label TLS (Tape Library Simulator Utility).

See [“General properties for devices”](#) on page 458.

The virtual tape libraries and simulated tape libraries do not support all of the utility jobs that are available for physical robotic libraries.

The following table describes the tasks that are available for these libraries.

Table 9-17 Utility jobs for virtual tape libraries and simulated tape libraries

Utility job	Available for virtual tape libraries	Available for simulated tape libraries
Lock	No	No
Unlock	No	No
Export	Yes	No
Import	Yes	No
Label media	Yes	No
Export expired media	No	No
Cleaning slot	No	No
Bar code rules	Yes	No
Clean drive	No	No

General options for utility jobs

General options for utility jobs provide the name of the utility job, and the priority of the access to the devices for the utility job.

Table 9-18 General options for utility jobs

Item	Description
Job name	Displays the name for the job.

Table 9-18 General options for utility jobs (*continued*)

Item	Description
Job priority	Displays the priority of the access to the devices for this job. See “About job priority” on page 195.

See [“About creating utility jobs to help manage devices and media”](#) on page 480.

See [“Inventorying media in a device”](#) on page 448.

See [“Creating a new catalog”](#) on page 248.

See [“Erasing media”](#) on page 449.

See [“Retensioning a tape”](#) on page 484.

See [“Formatting media in a drive”](#) on page 485.

See [“Labeling media”](#) on page 486.

See [“Ejecting media from a drive”](#) on page 487.

See [“Creating a cleaning job”](#) on page 488.

See [“Locking the robotic library’s front panel”](#) on page 493.

See [“Exporting media from a robotic library”](#) on page 490.

See [“Unlocking the robotic library’s front panel”](#) on page 494.

See [“Creating a job to initialize a robotic library”](#) on page 484.

See [“Exporting expired media from a robotic library”](#) on page 492.

Inventorying robotic libraries when Backup Exec services start

You can set a default so that all robotic libraries are included in the inventory job whenever Backup Exec services are started. Symantec recommends that you enable this default if media is often moved between robotic libraries. Backup Exec may take longer to start.

To inventory robotic libraries when Backup Exec services starts

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Settings**, click **Preferences**.

See [“Application settings for Backup Exec preferences”](#) on page 197.

- 3 Click **Include robotic libraries in inventory job when Backup Exec services start up**.
- 4 Click **OK**.

Creating a job to initialize a robotic library

You can create a job to initialize the robotic library. You can monitor this job on the **Job Monitor**.

You can also enable initialization whenever the Backup Exec service is started.

See [“Initializing robotic libraries when the Backup Exec service starts”](#) on page 470.

To create a job to initialize a robotic library

- 1 On the navigation bar, click **Devices**.
- 2 Select the robotic library.
- 3 Under **Robotic Library Tasks** in the task pane, select **Initialize**.
- 4 To specify a job name or a job priority, in the **Properties** pane, under **Settings**, click **General**.

See [“General options for utility jobs”](#) on page 482.

- 5 If you want Backup Exec to notify someone when the job completes, in the **Properties** pane, under **Settings**, click **Notification**.

See [“Sending a notification when a job completes”](#) on page 682.

- 6 Click **Run Now**.

Retensioning a tape

Use **Retension media** to run the tape in the tape drive from beginning to end at a fast speed so that the tape winds evenly and runs more smoothly past the tape drive heads. Refer to the documentation that came with your tape drive to see how often this utility should be performed.

Retensioning is primarily for Mini Cartridge and quarter-inch cartridges and is not supported on most other types of tape drives.

You cannot cancel a Retension operation after it has started; however, you can use **Cancel** to stop a queued retension operation.

The job will be submitted as a Run now job, unless you submitted the job on hold. You can monitor the Retension operation from the **Job Monitor**.

To retension a tape

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.
- 3 Select the drive or slot containing the media you want to retension.
- 4 Under **Media Tasks** in the task pane, select **Retension media**.
- 5 To specify a job name or a job priority, on the **Properties** pane, under **Settings**, click **General**.
See [“General options for utility jobs”](#) on page 482.
- 6 If you want a person or group to be notified when the job completes, in the **Properties** pane, under **Settings**, click **Notification**, and select the options you want.
See [“Sending a notification when a job completes”](#) on page 682.
- 7 Click **Run now**.

Formatting media in a drive

Use **Format media** to format the media currently in the drive. Most devices do not support formatting. If formatting is not supported, the option is not available.

If you use Format on a DC2000 tape, the formatting may take two or more hours to complete.

Caution: Formatting erases the media. All data on the media is lost.

You cannot cancel a Format operation after it has started; however, you can use Cancel to stop a queued Format operation.

The job will be submitted as a Run now job, unless you submitted the job on hold. You can monitor the Format operation from the **Job Monitor**.

To format media in a drive

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.
- 3 Select the drive or slot containing the media you want to format.

- 4 Under **Media Tasks** in the task pane, select **Format media**.

The media label that is displayed was read during the last inventory operation. The media label displayed does not change until another inventory operation occurs. Therefore, if you changed media in the slot or drive but did not run Inventory, the media label displayed may not match the actual media in the slot or drive.

- 5 To format the media that is displayed, click **Yes**.
- 6 To specify a job name or a job priority, on the **Properties** pane, under **Settings**, click **General**.

See “[General options for utility jobs](#)” on page 482.

- 7 If you want a person or group to be notified when the job completes, in the **Properties** pane, under **Settings**, click **Notification**, and select the options you want.

See “[Sending a notification when a job completes](#)” on page 682.

- 8 Click **Run now**.

Labeling media

Use **Label media** to immediately write a new media label on the media in the selected drive. This operation destroys any data on the media. To change the media label without destroying the data on the media (until an overwrite operation occurs), use **Rename**.

Note: Media that use bar code labels cannot be renamed. When you run label media jobs against the pieces of media that use bar code labels, the job logs report successfully completed jobs. However, the media label names do not change.

You cannot cancel a Label media operation after it has started; however, you can use **Cancel** to stop a queued Label media operation.

To label media

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon.
- 3 Select the drive or slot containing the media you want to label.

- 4 Under **Media Tasks** on the task pane, select **Label media**.
The following warning is displayed:
"This operation will be performed on the current media in the drive or slot. If the media has been changed since the last inventory operation was performed, the media label in the next dialog may not match the media in the drive or slot selected."
- 5 Click **OK**.
- 6 Type the name you want to use as the recorded media label for this media.
- 7 Click **OK** to erase all data on the media and re-label the media.
- 8 To specify a job name or a job priority, in the **Properties** pane, under **Settings**, click **General**.
See "[General options for utility jobs](#)" on page 482.
- 9 If you want a person or group to be notified when the job completes, on the **Properties** pane, under **Settings**, click **Notification** and select the options you want.
See "[Sending a notification when a job completes](#)" on page 682.
- 10 Click **Run now**.
- 11 Write this same media label on an external label fixed to the outside of the physical media.

Ejecting media from a drive

Use **Eject media** to eject the media currently in the standalone drive.

Some drives do not support a software-driven media eject. If the media is a tape, the tape is rewound and you may be instructed to manually remove it.

The job will be submitted as a Run now job, unless you submitted the job on hold.

To eject media from a drive

- 1 On the navigation bar, click **Devices**.
- 2 Expand the server icon, and then select the drive.
- 3 Under **Media Tasks** in the task pane, select **Eject media**.
- 4 To specify a job name or a job priority, in the **Properties** pane, under **Settings**, click **General**.

See "[General options for utility jobs](#)" on page 482.

- 5 If you want a person or group to be notified when the job completes, on the **Properties** pane, under Settings, click **Notification**, and select the options you want.
See [“Sending a notification when a job completes”](#) on page 682.
- 6 Click **Run Now**.

Creating a cleaning job

You can create and schedule a cleaning job for a robotic library drive. Additionally, Backup Exec automatically cleans a robotic library drive when the drive issues a tape alert that it requires cleaning.

Before submitting a cleaning job, you must define a cleaning slot that contains the cleaning tape.

See [“Defining a cleaning slot ”](#) on page 471.

You can view cleaning statistics for the drive.

See [“Cleaning properties for devices”](#) on page 465.

To run a cleaning job

- 1 On the navigation bar, click **Devices**.
- 2 Click the drive or robotic library containing the drive, and then select the drive.
- 3 Under **Drive Tasks** in the task pane, select **Clean**.
- 4 To specify a job name or a job priority, in the **Properties** pane, under **Settings**, click **General**.
See [“General options for utility jobs ”](#) on page 482.
- 5 If you want a person or group to be notified when the job completes, in the **Properties** pane, under **Settings**, click **Notification** and select the options you want.
See [“Sending a notification when a job completes”](#) on page 682.
- 6 If you want to run the job now, click **Run Now**. Otherwise, on the **Properties** pane, under **Frequency**, click **Schedule** to set the scheduling options you want to use.
See [“Scheduling jobs”](#) on page 356.

About importing media to a robotic library

An import media job inserts media into the robotic library so that the Backup Exec database is updated.

Before you create an import media job, note the following:

- If the media does not have a bar code, when you create the import job, you must select the option Auto-inventory after import is completed.
- If your robotic library uses a media magazine, make sure no jobs are currently running and that all media are ejected from the drive and are back in the magazine slots before swapping the magazine.

You can select any number of slots to import.

The Backup Exec import media job fully supports robotic libraries with portals. When this job is run, the slots you selected are checked for media. If media is found, it is exported to the portals. After all the media has been exported, you are prompted to insert new media into the portal so it can be imported. This process continues until all of the requested media has been imported into the robotic library.

See [“Importing media to a robotic library”](#) on page 489.

Importing media to a robotic library

To insert media into a robotic library, you must create an import media job so that the Backup Exec database is updated.

Before you create an import media job, note the following:

- If your robotic library uses a media magazine, make sure no jobs are currently running and that all media are ejected from the drive and are back in the magazine slots before swapping the magazine.

You can monitor this job on the Job Monitor.

To import media to a robotic library

- 1 On the navigation bar, click **Devices**.
- 2 Select the robotic library.
- 3 Click **Slots**.
- 4 On the Results pane, select the slots you want to import media to.
- 5 Under **Media Tasks** in the task pane, select **Import media**.
- 6 To specify a job name or a job priority, in the **Properties** pane, under **Settings**, click **General**.

See [“General options for utility jobs ”](#) on page 482.

- 7 If the media does not have a bar code, or if you want Backup Exec to automatically create an inventory job to run after the import job completes, under **Settings**, click **Options**.
See [“Imported Job Properties Options”](#) on page 490.
- 8 If you want Backup Exec to notify someone when the job completes, in the **Properties** pane, under **Settings**, click **Notification**.
See [“Setting up notification for alerts”](#) on page 661.
- 9 Do one of the following:
 - To run the job now, click **Run Now**.
 - To set scheduling options, in the **Properties** pane, under **Frequency**, click **Schedule**.
See [“Scheduling jobs”](#) on page 356.

Imported Job Properties Options

If you create an import media job to insert media into the robotic library, the **Auto-inventory after import is completed** option allows Backup Exec to create an inventory job. The inventory job automatically runs after the import job completes, and updates the Backup Exec database with information about the media.

See [“Importing media to a robotic library”](#) on page 489.

Exporting media from a robotic library

When you want to export media from a robotic library, you must create a job that updates the Backup Exec database.

The Backup Exec export media job fully supports robotic libraries with portals. When this job is run on one or more robotic library slots, the exported media is placed in the portals. If you select more media than there are portals, the robotic library will fill as many slots as possible, and then you are prompted to remove the media from the portal. This process continues until all of the selected media has been removed from the robotic library. You can also export expired media from a robotic library.

See [“Exporting expired media from a robotic library”](#) on page 492.

You can select a media vault that you want the exported media moved to after the export job has successfully completed.

You can monitor this job on the Job Monitor.

To export media from a robotic library

- 1 On the navigation bar, click **Devices**.
- 2 Select the robotic library.
- 3 Click **Slots**.
- 4 On the Results pane, select the slots you want to export media from.
- 5 Under **Media Tasks** in the task pane, select **Export media**.
- 6 To specify a job name or a job priority, in the **Properties** pane, under **Settings**, click **General**.
See [“General options for utility jobs”](#) on page 482.
- 7 If you want to move the media to a media vault after the export job is complete, in the **Properties** pane, under **Settings**, click **Options**.
See [“Export Media Job Properties Options”](#) on page 491.
- 8 If you want Backup Exec to notify someone when the job completes, in the **Properties** pane, under **Settings**, click **Notification**.
See [“Sending a notification when a job completes”](#) on page 682.
- 9 Do one of the following:
 - To run the job now, click **Run Now**.
 - To set scheduling options, on the **Properties** pane, under **Frequency**, click **Schedule**.
See [“Scheduling jobs”](#) on page 356.

Export Media Job Properties Options

When you select the option **Upon successful export, move the media to media vault**, the export media job updates the Backup Exec database with information about the media's location. You must physically move the media to an actual location that is represented by the vault name.

See [“Exporting media from a robotic library”](#) on page 490.

See [“Media locations and vaults”](#) on page 250.

About exporting expired media from a robotic library

The **export expired media** job lets you automate media handling in robotic libraries. This job removes the media that cannot be written to. You can then add scratch media to the robotic library to prepare for the next backup window.

After you export the expired media from the robotic library, the expired media appears in the **offline media location**. If the media is in a media set that has an applicable vault media rule, then the media appears in the vault location.

You can export cleaning media with the **export expired media** job. You can include all cleaning media, or all cleaning media that has been used more than a specified number of times.

You can choose to be reminded to import new media after an **export expired media** job completes successfully.

See [“Exporting expired media from a robotic library”](#) on page 492.

Exporting expired media from a robotic library

The **export expired media** job removes media that cannot be written to.

You can monitor the **export expired media** job through the Job Monitor.

To export expired media from a robotic library

- 1 On the navigation bar, click **Devices**.
- 2 Select the robotic library.
- 3 Under **Robotic Library Tasks** in the task pane, click **Export expired media**.
- 4 To specify a job name or a job priority, in the **Properties** pane, under **Settings**, click **General**.

See [“General options for utility jobs”](#) on page 482.

- 5 To set options for the **export expired media** job, in the **Properties** pane, under **Settings**, click **Options**.

- 6 Select the appropriate options.

See [“Options to export expired media”](#) on page 493.

- 7 If you want Backup Exec to notify someone when the job completes, in the **Properties** pane, under **Settings**, click **Notification**.

See [“Sending a notification when a job completes”](#) on page 682.

- 8 Do one of the following:

- To run the job now, click **Run Now**.
- To set scheduling options, on the **Properties** pane, under **Frequency**, click **Schedule**.

See [“Scheduling jobs”](#) on page 356.

Options to export expired media

Options for the **export expired media** job let you automate media handling in robotic libraries by removing the media that Backup Exec cannot write to.

See [“Exporting expired media from a robotic library”](#) on page 492.

Table 9-19 Options to export expired media

Item	Description
Include cleaning media in export	Includes cleaning media in the export of expired media.
Export cleaning media used more than x times	Displays the number of times that the cleaning media can be used before it is exported by this job.
After export, automatically prompt for new media to be imported	Displays a prompt to import new media to the slot after the export expired media operation has completed.
Upon successful export, move the media to media vault	Displays a vault to logically move the media to after the export job successfully completes. See “Creating media vaults” on page 251.

Locking the robotic library’s front panel

You can create a job to lock the robotic library’s front panel. You can monitor this job on the Job Monitor.

To lock the robotic library’s front panel

- 1 On the navigation bar, click **Devices**.
- 2 Select the robotic library.
- 3 Under **Robotic Library Tasks** in the task pane, select **Lock**.
- 4 To specify a job name or a job priority, in the **Properties** pane, under **Settings**, click **General**
See [“General options for utility jobs ”](#) on page 482.
- 5 If you want Backup Exec to notify someone when the job completes, in the **Properties** pane, under **Settings**, click **Notification**.
See [“Sending a notification when a job completes”](#) on page 682.
- 6 Do one of the following:
 - To run the job now, click **Run Now**.

- To set scheduling options, in the **Properties** pane, under **Frequency**, click **Schedule**.
See “[Scheduling jobs](#)” on page 356.

Unlocking the robotic library’s front panel

You must create a job to unlock the robotic library’s front panel. You can monitor this job on the **Job Monitor**.

To unlock the robotic library’s front panel

- 1 On the navigation bar, click **Devices**.
- 2 Select the robotic library.
- 3 Under **Robotic Library Tasks** in the task pane, click **Unlock**.
- 4 To specify a job name or a job priority, in the **Properties** pane, under **Settings**, click **General**.
See “[General options for utility jobs](#)” on page 482.
- 5 If you want Backup Exec to notify someone when the job completes, in the **Properties** pane, under **Settings**, click **Notification**.
See “[Sending a notification when a job completes](#)” on page 682.
- 6 Do one of the following:
 - To run the job now, click **Run Now**.
 - To set scheduling options, on the **Properties** pane, under **Frequency**, click **Schedule**.
See “[Scheduling jobs](#)” on page 356.

Managing backup-to-disk folders

This chapter includes the following topics:

- [About backup-to-disk folders](#)
- [About sharing backup-to-disk folders](#)
- [Changing the path of a backup-to-disk folder](#)
- [Deleting a backup-to-disk folder](#)
- [Recreating a backup-to-disk folder and its contents](#)
- [Changing the status of a device to online](#)
- [Renaming a backup-to-disk file](#)
- [Deleting a backup-to-disk file](#)
- [Recreating a deleted backup-to-disk file](#)
- [Erasing backup-to-disk files](#)
- [Recommendations for using backup-to-disk folders with backup jobs that use Granular Recovery Technology](#)
- [How to reclaim disk space for backup jobs that use Granular Recovery Technology](#)

About backup-to-disk folders

The backup-to-disk feature enables you to back up data to a folder on a hard disk. You can also use it as part of a backup process where you back up data to disk first and then transfer the data to a tape when more time is available.

On devices that have non-removable media, create a backup-to-disk folder. On devices that have removable media such as a zip drive, create a removable backup-to-disk folder. Backup-to-disk folders that are created on devices that have non-removable media support concurrent jobs from one or more media servers.

When you create a new backup-to-disk folder, Backup Exec automatically assigns the name Backup-to-Disk Folder x, where x is a number that is incremented by one each time a new backup-to-disk folder is created. You can rename the backup-to-disk folder at any time. You can also set defaults for backup-to-disk folders that will apply to every new backup-to-disk folder that is created. If you have the Central Admin Server Option (CASO) or the SAN Shared Storage Option installed, you can share backup-to-disk folders between computers.

When you back up to disk, Backup Exec places the data in a backup-to-disk file in the backup-to-disk folder you specify. Backup-to-disk files are virtual media where backed up data is stored. Backup-to-disk files are like any other type of media, so you can inventory, catalog, erase, and restore them.

Since Backup Exec recognizes the backup-to-disk folders as devices, you can view them by selecting Devices on the navigation bar. You can view the backup-to-disk files from both the Devices view and the Media view.

In Windows Explorer, the backup-to-disk folders display in the path you specified when you added the folders. The backup-to-disk files display with a .bkf file extension. Each backup-to-disk folder also contains a file named changer.cfg and a file named folder.cfg, which store information about the backup-to-disk files.

Note: Do not delete or edit the changer.cfg or folder.cfg files.

A subfolder with a prefix of IMG in the name may display under a backup-to-disk folder.

Backup Exec creates this subfolder when the following conditions are met in a backup job:

- The option to enable Granular Recovery Technology (GRT) is selected.
- A backup-to-disk folder is selected as the backup device.

Disaster recovery from backup-to-disk folders must be done by remote Intelligent Disaster Recovery using a media server with access to the backup-to-disk folders.

See [“Requirements for creating a backup-to-disk folder”](#) on page 497.

See [“Requirements for creating a removable backup-to-disk folder”](#) on page 498.

See [“About the Virtual Tape Library Unlimited Drive Option ”](#) on page 452.

Requirements for creating a backup-to-disk folder

You can create a backup-to-disk folder in any of the following locations to which you can write a file:

- NTFS partitions (local or remote)

The backup-to-disk folder must exist on an NTFS partition for backup jobs in which the Granular Recovery Technology (GRT) option is selected. This option is available for Microsoft Exchange databases and storage groups, Microsoft Active Directory, and Microsoft SharePoint content database and Team database.

See [“Recommendations for using backup-to-disk folders with backup jobs that use Granular Recovery Technology”](#) on page 511.

- DFS shares

- FAT/FAT32 partitions (local or remote)

- Veritas™ Volume Manager partitions

- RAID drives with any configuration

- NFS volumes

- Network Attached Storage (NAS) devices

If a NAS device is emulating a Windows operating system, contact the NAS manufacturer for assistance before creating backup-to-disk folders on the NAS device. Symantec does not certify NAS devices. If the operating system is a proprietary operating system and not a true Windows operating system, Symantec cannot properly troubleshoot the device.

You should create a backup-to-disk folder on a different physical disk than the disk you want to back up. For example, if the Backup Exec Advanced Open File Option (AOFO) is used to snap volumes during a backup, and if the destination device is a backup-to-disk folder, the backup-to-disk folder should be on a separate volume that is not being snapped.

Similarly, when making selections for backups that you are targeting to a backup-to-disk folder, avoid including that folder in the selections for the job. For example, if you create a new backup-to-disk folder in C:\Backup Folders and then select the entire C:\ volume for backup, make sure that you exclude C:\Backup Folders from the selection list.

See [“Requirements for creating a removable backup-to-disk folder”](#) on page 498.

See [“Creating a backup-to-disk folder by using the Backup-to-Disk Wizard”](#) on page 498.

See [“Creating a backup-to-disk folder by setting properties”](#) on page 499.

See [“Recommendations for using backup-to-disk folders with backup jobs that use Granular Recovery Technology”](#) on page 511.

Requirements for creating a removable backup-to-disk folder

A removable backup-to-disk folder works as follows:

- Supports spanning of backup sets from one piece of media to another.
- Does not support concurrent operations.

You should not share removable backup-to-disk folders between media servers.

You can create removable backup-to-disk folders on any device that has removable media, provided the device appears as a drive letter and is formatted with a file system.

Devices with removable media can include the following:

- CDR-RW
- DVD-RW
- ZIP
- REV
- Removable hard disk

Note: You must use Microsoft's Live File System to format new CDR-RW and DVD-RW media before you can create backup-to-disk folders on the media. you can also use erased CDR-RW and DVD-RW media that were previously formatted with Live File System. For more information about Live File System, see your Microsoft documentation.

Creating a backup-to-disk folder by using the Backup-to-Disk Wizard

If you are new to Backup Exec or are uncertain about how to set up a backup-to-disk folder, you can use the Backup-to-Disk Wizard. The wizard guides you through the process of creating a backup-to-disk folder or editing an existing one. While the wizard prompts you to select some options, most of the settings are based on the default settings.

To create a backup-to-disk folder by using the Backup-to-Disk Wizard

- 1 On the menu bar, click **Tools> Wizards> Backup-to-Disk Wizard**.
- 2 Follow the on-screen prompts.
 See [“Requirements for creating a backup-to-disk folder”](#) on page 497.
 See [“Creating a backup-to-disk folder by setting properties”](#) on page 499.

Creating a backup-to-disk folder by setting properties

You must create at least one backup-to-disk folder or removable backup-to-disk folder before you can use the backup-to-disk feature.

To create a backup-to-disk folder by setting properties

- 1 On the navigation bar, click **Devices**.
- 2 Right-click the server for which you want to create a backup-to-disk folder.
- 3 On the shortcut menu, click one of the following:

To create a backup-to-disk folder on a hard drive or on a network drive	Click New Backup-to-Disk Folder .
---	--

To create a backup-to-disk folder on a removable device	Click New Removable Backup-to-Disk Folder .
---	--

- 4 On the **General** tab, enter the appropriate information.
 See [“General properties for backup-to-disk folders”](#) on page 502.
- 5 On the **Advanced** tab, enter the appropriate information.
 See [“Advanced properties for backup-to-disk folders”](#) on page 501.
- 6 Click **OK**.

Application settings for new backup-to-disk folders

Application settings apply to new backup-to-disk folders that you create.

See [“Editing application settings for new backup-to-disk folders”](#) on page 505.

Table 10-1 Application settings for new backup-to-disk folders

Item	Description
Maximum number of backup sets per backup-to-disk file	<p>Displays the maximum number of backup sets to be written to each backup-to-disk file in this folder. The maximum number can range from 1 to 8192. The default is 100.</p> <p>Fewer backup sets in a backup-to-disk file may allow the overwrite protection period to expire sooner, and disk space to be reclaimed faster.</p>
Maximum size for backup-to-disk files	<p>Displays the maximum size for each backup-to-disk file contained in this folder. You can select either MB or GB as the unit of size. The file size can be from 1 MB to 4096 GB. The default is 4 GB.</p> <p>If you create smaller but more numerous backup-to-disk files, performance may be slower. If large backup-to-disk files are created, file system limitations could cause memory allocation problems or network issues, especially if the backup-to-disk files are stored across a network.</p> <p>This option works with the option Maximum number of backup sets per backup-to-disk file.</p>
Allocate the maximum size when creating the backup-to-disk file	<p>Creates the backup-to-disk file at the maximum size to reduce disk fragmentation.</p> <p>You may want to increase the append period. However, increased append periods can cause an increase in the overall overwrite protection period since the overwrite protection period starts at the end of the last append job. This can result in fewer backup jobs being targeted to this backup-to-disk folder. To avoid this, set the maximum size for backup-to-disk files to an appropriate size.</p> <p>When the backup-to-disk file is initially created at the maximum size, the backup job may be delayed while Backup Exec creates the file. The backup job remains in a running state until the backup-to-disk file is created and data can be written to it.</p> <p>When you select this option, Backup Exec hides the option Maximum number of backup sets per backup-to-disk file. As a result, all of the space that is allocated to the backup-to-disk file is used.</p> <p>This option is not available for removable backup-to-disk folders.</p>
Allow x concurrent jobs for this backup-to-disk folder	<p>Displays the number of concurrent operations that you want to allow to this folder. This number can range from 1 to 16.</p> <p>This option is not available for removable backup-to-disk folders.</p>

Table 10-1 Application settings for new backup-to-disk folders (*continued*)

Item	Description
Low disk space threshold	<p>Indicates if backup operations to the backup-to-disk folder are suspended when the amount of free space on the disk reaches a specific level.</p> <p>When the disk's free space reaches this threshold, Backup Exec places current jobs on hold until disk space is available. The low disk space threshold prevents disk-full errors and provides early warning. This threshold prevents jobs from being submitted to a backup-to-disk folder that does not have enough disk space to allow the job to complete. Backup Exec can instead submit the jobs to backup-to-disk folders that do have enough disk space. The backup-to-disk status displays Low Disk Space. You must free some disk space to allow job submission to resume.</p>
Backup-to-disk default folder location	Displays the default path for new backup-to-disk folders.

Advanced properties for backup-to-disk folders

Advanced properties for backup-to-disk folders provide information about disk space management and device settings.

See “[Creating a backup-to-disk folder by setting properties](#)” on page 499.

Table 10-2 Advanced properties for backup-to-disk folders

Item	Description
Low disk space threshold (at which backup operations are suspended)	<p>Indicates if backup operations to the backup-to-disk folder are suspended when the amount of free space on the disk reaches a specific level.</p> <p>When the disk's free space reaches this threshold, Backup Exec places current jobs on hold until disk space is available. The low disk space threshold prevents disk-full errors and provides early warning. This threshold prevents jobs from being submitted to a backup-to-disk folder that does not have enough disk space to allow the job to complete. Backup Exec can instead submit the jobs to backup-to-disk folders that do have enough disk space. The backup-to-disk status displays Low Disk Space. You must free some disk space to allow job submission to resume.</p>

Table 10-2 Advanced properties for backup-to-disk folders (*continued*)

Item	Description
Auto detect settings	Indicates if Backup Exec automatically detects the preferred settings for this device.
Buffered reads	<p>Indicates the following:</p> <ul style="list-style-type: none"> ■ You do not want Backup Exec to automatically detect settings for this device ■ You want this device to allow buffered reads, which is the reading of large blocks of data. <p>Enabling buffered reads may provide increased performance.</p>
Buffered writes	<p>Indicates the following:</p> <ul style="list-style-type: none"> ■ You do not want Backup Exec to automatically detect settings for this device ■ You want this device to allow buffered writes, which is the writing of large blocks of data.

See [“General properties for backup-to-disk folders”](#) on page 502.

General properties for backup-to-disk folders

General properties for backup-to-disk folders provide information about the folders.

See [“Creating a backup-to-disk folder by setting properties”](#) on page 499.

Table 10-3 General properties for backup-to-disk folders

Item	Description
Name	<p>Displays the name of the backup-to-disk folder. Backup-to-disk folder names must not exceed 128 characters.</p> <p>See “Renaming storage devices” on page 447.</p>

Table 10-3 General properties for backup-to-disk folders (*continued*)

Item	Description
Path	<p>Displays the path where the backup-to-disk folder is to reside. For a removable backup-to-disk folder, the drive on which the folder is located appears.</p> <p>The backup-to-disk path name, which includes the backup-to-disk folder name, must not exceed 512 characters.</p> <p>A browse button next to the Path field lets you browse to other paths.</p> <p>See “Changing the path of a backup-to-disk folder” on page 506.</p>
Pause	<p>Indicates if the backup-to-disk folder is paused.</p>
Enable	<p>Indicates if Backup Exec has exclusive use of this backup-to-disk folder. If the check box is clear, the device is disabled and cannot be used by Backup Exec. The device is available for other applications.</p>
Online	<p>Indicates that the backup-to-disk folder is online if a dimmed check box with a check mark appears. If the folder is offline, a check mark does not appear. No operations are allowed on the folder until it is online again.</p> <p>The folder appears as offline if the following occurs:</p> <ul style="list-style-type: none"> ■ The backup-to-disk folder is on a remote computer and connectivity is not available. ■ The access rights to the folder or to the remote computer are incorrect. ■ The backup-to-disk folder is write-protected. <p>See “Changing the status of a device to online” on page 508.</p>

Table 10-3 General properties for backup-to-disk folders (*continued*)

Item	Description
<p>Maximum size for backup-to-disk files</p>	<p>Displays the maximum size for each backup-to-disk file contained in this folder. The file size can be from 1 MB to 4096 GB. The default is 4 GB.</p> <p>Backup-to-disk folders that were created with earlier versions of Backup Exec continue to use the default file size of 1 GB.</p> <p>If you create smaller but more numerous backup-to-disk files, performance may be slower. If large backup-to-disk files are created, file system limitations could cause memory allocation problems or network issues, especially if the backup-to-disk files are stored across a network.</p> <p>This option works with the option Maximum number of backup sets per backup-to-disk file.</p>
<p>Allocate the maximum size for backup-to-disk files</p>	<p>Creates the backup-to-disk file at the maximum size to reduce disk fragmentation.</p> <p>You may want to increase the append period. However, increased append periods can cause an increase in the overall overwrite protection period since the overwrite protection period starts at the end of the last append job. This can result in fewer backup jobs being targeted to this backup-to-disk folder. To avoid this, set the maximum size for backup-to-disk files to an appropriate size.</p> <p>When the backup-to-disk file is initially created at the maximum size, the backup job may be delayed while Backup Exec creates the file. The backup job remains in a running state until the backup-to-disk file is created and data can be written to it.</p> <p>When you select this option, Backup Exec hides the option Maximum number of backup sets per backup-to-disk file. As a result, all of the space that is allocated to the backup-to-disk file is used.</p> <p>This option is not available for a removable backup-to-disk folder.</p>

Table 10-3 General properties for backup-to-disk folders (*continued*)

Item	Description
Maximum number of backup sets per backup-to-disk file	<p>Displays the maximum number of backup sets to be written to each backup-to-disk file in this folder. The maximum number can range from 1 to 8192. The default is 100.</p> <p>Fewer backup sets in a backup-to-disk file may allow the overwrite protection period to expire sooner, and disk space to be reclaimed faster.</p>
Allow x concurrent jobs for this backup-to-disk folder	<p>Displays the number of concurrent operations that you want to allow to this folder. This number can range from 1 to 16.</p> <p>This option is not available for a removable backup-to-disk folder.</p>

See [“Editing application settings for new backup-to-disk folders”](#) on page 505.

See [“Application settings for new backup-to-disk folders”](#) on page 499.

Editing application settings for new backup-to-disk folders

You can edit the application settings that apply to new backup-to-disk folders that you create.

To edit application settings for new backup-to-disk folders

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Settings**, click **Backup-to-Disk**.
- 3 Edit the application settings as appropriate.
 See [“Application settings for new backup-to-disk folders”](#) on page 499.
- 4 Click **OK**.

About sharing backup-to-disk folders

You can share backup-to-disk folders between computers if the Central Admin Server Option (CASO) or the SAN Shared Storage Option is installed.

Note: You cannot share a removable storage device.

In the **Devices** view, shared backup-to-disk folders are listed under each computer that has access to that backup-to-disk folder. All of the logical groupings of the backup-to-disk folders are displayed in the **Devices** view, under **Device Pools**.

To stop sharing a backup-to-disk folder, delete it from the computer that you don't want to share it with.

A backup-to-disk folder that was created by a previous installation of Backup Exec cannot be shared, and is considered unknown by Backup Exec. If Backup Exec finds an unknown backup-to-disk folder at the specified path, you are prompted to create a new backup-to-disk folder at that path.

See [“Sharing an existing backup-to-disk folder”](#) on page 506.

Sharing an existing backup-to-disk folder

If you have the Central Admin Server Option (CASO) or the SAN Shared Storage Option installed, you can share backup-to-disk folders between computers.

See [“About sharing backup-to-disk folders”](#) on page 505.

To share an existing backup-to-disk folder

- 1 On the computer on which you want to add the folder for sharing, on the navigation bar, click **Devices**.
- 2 Right-click the server on which you want to add the folder for sharing.
- 3 On the shortcut menu, click **Add Shared Backup-to-Disk folder**.
- 4 Type or browse to the path of the shared backup-to-disk folder that you want to add to this computer.
- 5 Click **OK**.

Changing the path of a backup-to-disk folder

To change the path of a backup-to-disk folder, you must first create a new backup-to-disk folder, and then move the backup-to-disk files from the original backup-to-disk folder to the new backup-to-disk folder.

To change the path of a backup-to-disk folder

- 1 Add a new backup-to-disk folder with a name and path that is different than the original backup-to-disk folder.
- 2 In Windows Explorer, copy and paste the backup-to-disk files from the original backup-to-disk folder to the new folder.
- 3 On the Backup Exec navigation bar, click **Devices**.

- 4 Right-click the new backup-to-disk folder, and then click **Scan** on the shortcut menu, or select the new folder and press <F5>.
- 5 Click the new backup-to-disk folder and verify that the backup-to-disk files appear in the results pane.
- 6 To rename the new backup-to-disk folder to match the name of the original folder, delete the original backup-to-disk folder.
See [“Deleting a backup-to-disk file”](#) on page 509.
- 7 Rename the new folder.
See [“Renaming storage devices”](#) on page 447.

Deleting a backup-to-disk folder

When you use the Backup Exec **Delete** option to delete a backup-to-disk folder, the folder is removed from Backup Exec, but the backup-to-disk folder and the files in it remain on the disk so you can recreate them later. If you also want to delete the folder from the disk, use the Windows Delete option. However, you cannot recreate the backup-to-disk folder or files after you delete them from the disk.

Note: If you want to delete the folder from the disk, use Windows Explorer to navigate to the folder and delete it. When the folder is removed from the disk using Windows Explorer, you cannot recreate the folder or the files in Backup Exec.

To delete a backup-to-disk folder

- 1 On the navigation bar, click **Devices**.
- 2 Expand the icon for the computer where the backup-to-disk folder is located.
- 3 Select the backup-to-disk folder that you want to remove.
- 4 Under **General Tasks** in the task pane, select **Delete**.
- 5 Click **Yes**.

Recreating a backup-to-disk folder and its contents

If you have deleted a backup-to-disk folder from Backup Exec, but have not deleted it from the disk, you can recreate the backup-to-disk folder and the files in it. You must know the name and path of the original backup-to-disk folder in order to

recreate it. If you deleted a backup-to-disk folder from the disk, you cannot recreate it.

To recreate a backup-to-disk folder and its contents

- 1 Add a new folder to Backup Exec using the same name and path as the deleted folder.
- 2 When you are prompted, click **Yes** to recreate the folder at the specified path.
- 3 On the navigation bar, click **Devices**.
- 4 Expand the icon for the computer where the backup-to-disk folder is located.
- 5 Select the new folder.
- 6 Under **Media Tasks** in the task pane, select **Inventory** and create and run an inventory job.

See [“Requirements for creating a backup-to-disk folder”](#) on page 497.

See [“About inventoring media”](#) on page 448.

Changing the status of a device to online

If a device goes offline, you can change the status to online after you correct the problem.

To change the status of a device to online

- 1 On the navigation bar, click **Devices**.
- 2 Expand the icon for the computer where the device is located.
- 3 Select the device that is offline.
- 4 Under **General Tasks** in the task pane, click **Online**.

See [“Troubleshooting hardware-related issues”](#) on page 789.

Renaming a backup-to-disk file

When you rename a backup-to-disk file, the name changes in Backup Exec, on the disk, and on the media label.

To rename a backup-to-disk file

- 1 On the navigation bar, click **Devices**.
- 2 Expand the icon for the computer where the backup-to-disk folder is located.
- 3 Select the folder that contains the file you want to rename.
- 4 On the results pane, select the file you want to rename.

- 5 Under **General Tasks** in the task pane, select **Rename**.
- 6 Type a new name for the file, and then click **OK**.

Deleting a backup-to-disk file

You must move backup-to-disk files to the **Retired Media** set before you can delete them. When you delete a backup-to-disk file from the **Media** tab in Backup Exec, it is deleted from Backup Exec but the Windows folder and files still exist in Windows Explorer. You can recreate the deleted backup-to-disk files.

If you want to delete the file from the disk, use Windows Explorer to navigate to the file and delete it. When the file is removed from the disk using Windows Explorer, you cannot restore it in Backup Exec.

To delete a backup-to-disk file

- 1 On the navigation bar, click **Media**.
- 2 Click the media set that contains the backup-to-disk file.
- 3 Use the Windows drag-and-drop feature to move the backup-to-disk file into the **Retired Media** set.
- 4 On the results pane, select the backup-to-disk file you want to delete.
- 5 Under **General Tasks** in the task pane, select **Delete**.
- 6 When prompted to delete the backup-to-disk file, click **Yes**, or if you selected multiple backup-to-disk files, click **Yes to All**.

Recreating a deleted backup-to-disk file

If you deleted a backup-to-disk file from Backup Exec, but did not use Windows Explorer to delete the file from the disk, you can recreate it by running **Inventory**.

To recreate a backup-to-disk file

- 1 On the navigation bar, click **Devices**.
- 2 Expand the icon for the computer where the backup-to-disk folder is located.
- 3 Select the folder where the backup-to-disk file was located before you deleted it.
- 4 Under **Media Tasks** in the task pane, select **Inventory** and create and run an inventory job.

See “[About inventorying media](#)” on page 448.

Erasing backup-to-disk files

Erasing backup-to-disk files removes the data from both the backup-to-disk folder and the disk, and removes the file references from the catalog. However, the backup-to-disk file remains for use with future backup jobs. You cannot restore the data after you erase it. If you want to remove data from Backup Exec and restore it later, delete the file from the **Media** view.

See [“Deleting a backup-to-disk file”](#) on page 509.

Unlike other types of devices, when you erase a file from a backup-to-disk folder you cannot choose whether to perform a quick erase or a long erase. Backup Exec performs only a quick erase on backup-to-disk files in backup-to-disk folders.

Caution: You cannot restore the data that you erase. Before you erase files, be sure that you no longer need them.

To erase a backup-to-disk file

- 1 On the navigation bar, click **Devices**.
- 2 Expand the icon for the computer where the backup-to-disk folder is located.
- 3 Under **Backup-to-Disk Folders**, click the backup-to-disk folder that contains the file you want to erase.
- 4 On the Results pane, select the file you want to erase.
- 5 Under **Media Tasks** in the task pane, select **Erase media, quick**.
- 6 Click **OK** to continue.
- 7 Click **Yes**, or if more than one file was selected, click **Yes to All**.
- 8 To specify a job name or a job priority, on the **Properties** pane, under **Settings**, click **General**.
See [“General options for utility jobs”](#) on page 482.
- 9 If you want a person or group to be notified when the job completes, on the **Properties** pane, under **Settings**, click **Notification** and select the options you want.
See [“Setting up notification for alerts”](#) on page 661.
- 10 If you want to run the job now, click **Run Now**. Otherwise, on the **Properties** pane, under **Frequency**, click **Schedule** to set the scheduling options you want to use.
See [“Scheduling jobs”](#) on page 356.

Recommendations for using backup-to-disk folders with backup jobs that use Granular Recovery Technology

The following recommendations help ensure that you do not run out of space for the backups that use Granular Recovery Technology (GRT):

Table 10-4 Recommendations for using backup-to-disk folders with GRT operations

Recommendation	Description
Create a separate backup-to-disk folder specifically for all GRT-enabled backup jobs	You must manage the media that is created by GRT-enabled jobs differently than other backup-to-disk media because of the IMG files. For best results, you should create a separate backup-to-disk folder specifically for all GRT-enabled backup jobs.
Select the specific backup-to-disk folder you want to use for GRT-enabled backup jobs	You should specifically select the backup-to-disk folder that you want to use when you create GRT-enabled jobs. If you do not change the default device setting of All Devices you could accidentally send your GRT job to a tape.
Do not allocate the maximum size for backup-to-disk files	If you select the Allocate the maximum size when creating the backup-to-disk file option, Backup Exec creates a backup-to-disk file that is as large as the size that you specified. Since GRT information is stored in IMG media, the backup-to-disk file does not hold backup data. The extra space that the backup-to-disk file occupies can often lead to failed jobs because of low disk space. See “Creating a backup-to-disk folder by setting properties” on page 499.

Table 10-4 Recommendations for using backup-to-disk folders with GRT operations (*continued*)

Recommendation	Description
Calculate your disk space requirements carefully before you assign a low disk space threshold	<p>The low disk space threshold is the amount of free space on the drive at which Backup Exec suspends backup operations to the backup-to-disk folder. If you assign a low disk space threshold to the backup-to-disk folder, you may avoid using all of the available disk space.</p> <p>You should be careful with low disk space thresholds. The amount you specify as a low disk space threshold is unavailable to Backup Exec for backup-to-disk backups. If you create a large threshold for low disk space, you may quickly run out of disk space. Be aware of your low disk space threshold before you run backup jobs.</p> <p>You should consider the low disk space threshold when you calculate the total amount of space available to a backup-to-disk folder. Remember to also consider any other data that exists on the volume. As the amount of other data that is contained on the volume increases, the amount of space available to the backup-to-disk folder decreases.</p> <p>See “Advanced properties for backup-to-disk folders” on page 501.</p>
Do not fill up a drive that hosts a backup-to-disk folder that is used for GRT operations	<p>When you calculate the total amount of available space on a volume, remember to consider the size of any other data that resides on it. This amount can include other backup-to-disk files or data from other applications.</p> <p>If the drive fills up or if the low disk space threshold is met, you have to reclaim disk space to run backup jobs.</p>

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“How to reclaim disk space for backup jobs that use Granular Recovery Technology”](#) on page 513.

See [“Requirements for creating a backup-to-disk folder”](#) on page 497.

How to reclaim disk space for backup jobs that use Granular Recovery Technology

When a job that uses Granular Recovery Technology (GRT) creates a backup set, Backup Exec erases an IMG media that has an expired protection period. Backup Exec automatically erases the expired media to reclaim disk space for the new media that the backup job creates. Backup Exec erases the oldest IMG media for each backup set in a GRT-enabled job. For example, if a GRT-enabled backup job creates three backup sets, Backup Exec erases three IMG media that have expired overwrite protection periods.

If you want to erase more than one IMG media per backup set, see the following Symantec knowledge base article:

<http://entsupport.symantec.com/umi/V-269-8>

If Backup Exec runs out of disk space during a GRT-enabled backup, it deletes any expired media and continues the job. If no expired media is available, the job is queued and the backup-to-disk folder is paused. To resume the job, you must reclaim disk space or wait for media to expire. Backup Exec automatically checks the amount of available space periodically. When enough space is available, Backup Exec automatically resumes the job.

The Job Monitor provides information about the GRT-enabled backup jobs that cannot run due to low disk space. The Job Monitor displays "Queued" as the job state and "Ready; No idle devices are available" as the job status. When Backup Exec checks for available space, the job state changes to "Mounting Media."

You can reclaim disk space by using any of the following methods:

Table 10-5 How to reclaim disk space for GRT backup operations

Method	Description
Erase IMG media or backup-to-disk files to provide adequate disk space	Delete any IMG media or backup-to-disk files you no longer need. See "Erasing backup-to-disk files" on page 510. See "Erasing media" on page 449.
Use Windows Explorer to remove data that is not related to Backup Exec	The volume may contain some data that is not related to Backup Exec. You can use Windows Explorer to delete this data. You should never use Windows Explorer to delete Backup Exec data.

Table 10-5 How to reclaim disk space for GRT backup operations (*continued*)

Method	Description
Wait for IMG media or backup-to-disk files to expire according to media set rules	Media set rules include the append and overwrite protection periods you set for media. You can wait until these rules expire to allow Backup Exec to reclaim disk space. See “About media in Backup Exec” on page 219.
Remove the low disk space threshold setting for the backup-to-disk folder for the current operation and take appropriate actions after the job has completed	If you set a low disk space threshold for the backup-to-disk folder, it may prohibit the job from completing due to low disk space. You can temporarily remove the low disk space threshold to allow Backup Exec access to that reserved space. When the job is complete, you can reinstate the low disk space threshold with a smaller reserve setting. See “Advanced properties for backup-to-disk folders” on page 501.

See [“Pausing storage devices”](#) on page 446.

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“Recommendations for using backup-to-disk folders with backup jobs that use Granular Recovery Technology”](#) on page 511.

Managing device pools

This chapter includes the following topics:

- [About device pools](#)
- [Creating device pools](#)
- [Device pool options](#)
- [Adding devices to a device pool](#)

About device pools

A device pool is a group of devices that enables load-balancing of Backup Exec jobs sent to the same device pool for processing. The jobs are spread among the devices in a pool so that the workload is shared among the devices.

When you submit a backup job to a device pool, the job is automatically sent to the first available device in that device pool. As other jobs are created and started, they can run concurrently on other devices in the device pool. By dynamically allocating devices as jobs are submitted, Backup Exec processes jobs quickly and efficiently.

Devices can belong to more than one device pool. Device pools can contain different types of devices, including specific devices in multi-device robotic libraries.

You can assign priorities to devices in a device pool so that a specific device is used before other devices in the device pool. The priority assigned to a device in one device pool is unrelated to that device's priority in any other device pool. For example, if Device 1 is placed in both Device Pool A and Device Pool B, you can assign different priorities to it in each device pool. Device 1 can have a high priority in Device Pool A and a low priority in Device Pool B.

You can send backup jobs to a specific device or you send jobs to a device pool. However, if that device is busy, the job must wait until the device becomes

available. When a specific device is assigned, the job cannot be automatically routed to the next available device.

Device pools also provide automatic job rescheduling. For example, if a device pool contains four stand-alone drives and the first device fails because of a hardware error, the job that was running on the failed device is resubmitted and placed on hold, and the other jobs are automatically routed to the working devices in the device pool.

In a device pool, Backup Exec selects the oldest recyclable media in all of the devices in the device pool to use first. If more than one media that meets the requirements is found, Backup Exec then searches the devices in a device pool according to device priority and uses the oldest recyclable media in the device that has the highest priority.

All Devices (Server Name) is the default device pool, created when Backup Exec is installed. All devices recognized by Backup Exec at startup are automatically assigned to **All Devices (Server Name)**. Devices and simulated tape libraries that are on any computers on which the Remote Media Agent for Linux Servers is installed are excluded from the **All Devices (Server Name)** device pool. Deduplication storage folders and OpenStorage devices are also excluded from the **All Devices** device pool.

Note: For a new installation of Backup Exec that includes the Storage Provisioning Option, the default device pool is the **All Virtual Disks** device pool.

You can create other device pools to meet your particular requirements, and assign and reassign devices to these pools. For example, you may want to separate high-performance devices from lower performance devices in a separate device pool in order to send high-priority jobs to the fast device pool for quicker completion.

See [“About the All Virtual Disks device pool in the Storage Provisioning Option”](#) on page 2029.

See [“About creating device pools for devices attached to the Remote Media Agent for Linux Servers ”](#) on page 1977.

Creating device pools

Device pools can consist of stand-alone drives, drives in single or multiple drive robotic libraries, and backup-to-disk folders.

To create a device pool

- 1 On the navigation bar, click **Devices**.
- 2 Click **Device Pools**.
- 3 Under **Device Tasks** in the task pane, click **Configure devices assistant**.
- 4 Click **Device Pool**.
- 5 Type or select the appropriate information, and then click **OK**.

See “[Device pool options](#)” on page 517.

Device pool options

Device pool options let you create a new device pool, or add devices to an existing pool.

See “[Creating device pools](#)” on page 516.

Table 11-1 Device pool options

Item	Description
Device pool name	Displays the name of the device pool.
Description	Displays the description of the device pool.
Device type	Displays a list of device types that you can use to filter the list of devices available for the new device pool. Only devices of this type are displayed for selection.
Device sub-type:	Displays a list of device sub-types if any are available. A device subtype lets you increase the filter on the list of devices available for the new device pool. Only devices of this type and sub-type are displayed for selection.
Select the devices to be included in this device pool	Displays the devices that you can include in the new device pool.

See “[About device pools](#)” on page 515.

Adding devices to a device pool

You can add a device to an existing device pool.

To add a device to a device pool

- 1 On the navigation bar, click **Devices**.
- 2 Select the device pool.
- 3 Under **Device Pool Tasks** in the task pane, select **Add device**.
- 4 Select the appropriate options, and then click **OK**.

See “[Device pool options](#)” on page 517.

See “[About the All Virtual Disks device pool in the Storage Provisioning Option](#)” on page 2029.

Setting priorities for devices in a device pool

You can set a priority that determines the order in which the devices in a device pool are used. The default priority is 10 so all devices have the same priority initially. The device to which you assign the lowest number is the first device to be used in the device pool; for example, a device with a priority of 1 is used before a device with a priority of 5. You can set a priority of 1 to 99.

Note: Overwrite and append periods for media take precedence over device priority.

The Priority option is only displayed when you are viewing device properties under a device pool icon. Drives that are displayed under the **Stand-alone Drives** icon or the **Robotic Libraries** icon do not display a Priority option because the drive may belong to multiple device pools and have a different priority in each device pool.

To set priorities for devices in a device pool

- 1 On the navigation bar, click **Devices**.
- 2 Under **Device Pools**, select the device pool containing the device for which you want to set a priority.
- 3 Select the device.
- 4 Under **General Tasks** in the task pane, select **Properties**.
- 5 Click **General**.
- 6 In **Priority**, type a number from 1 to 99, with 1 designating this device as the first device to be used in the device pool, and then click **OK**.

See “[Viewing storage device properties](#)” on page 458.

Removing devices from a device pool

You can remove a device from a device pool. The device will still be in the Backup Exec device database and will still be available for use in other device pools. It is not necessary to remove devices from a device pool before you remove that pool; the devices are automatically removed when the pool is deleted.

To remove a device from a device pool

- 1 On the navigation bar, click **Devices**.
- 2 Under **Device Pools**, select the device pool from which you want to remove a device.
- 3 Select the device that you want to remove from the device pool. You can select multiple devices to remove.

Make sure you select a device under **Device Pools** and not under the server icon; if you remove a device under the server icon, the device is deleted from the database, not just from the device pool.

- 4 Under **Device Pool Tasks** in the task pane, select **Remove device**.
- 5 When prompted, click **Yes** to remove the device from the device pool.

Deleting device pools

It is not necessary to delete devices from a device pool before you delete that pool; the devices are automatically removed when the pool is deleted.

You cannot delete the **All Devices** device pool, but you can delete all of the devices in it.

If scheduled jobs are assigned to the deleted device pool, you are prompted to redirect them to another device pool.

To delete a device pool

- 1 On the navigation bar, click **Devices**.
- 2 Under **Device Pools**, select the device pool or pools that you want to delete.
- 3 Under **General Tasks** on the task pane, select **Delete**.
- 4 When prompted, click **Yes** to delete the device pool.
- 5 If scheduled jobs are assigned to the deleted device pool, you are prompted to redirect the jobs to another device pool or stand-alone drive.

See [“Removing devices from a device pool”](#) on page 519.

See [“Retarget Job options ”](#) on page 520.

Retarget Job options

If you delete a device pool or a media set, and scheduled jobs are assigned to that device pool or media set, you are prompted to redirect the jobs to another device or device pool, or to another media set. The **Destination** field displays available devices or media sets to which you can reassign the scheduled job.

If you choose to not reassign the job, the job fails. To reassign the job to another device later, you must edit the job.

Device Pool Properties

Properties for device pools provide a name and description of the device pool, and when it was created.

See [“Viewing properties”](#) on page 217.

Table 11-2 Device Pool Properties

Item	Description
Name	Displays the name of the device pool. See “Renaming storage devices” on page 447.
Description	Displays a description of the device pool.
Creation date	Displays the date and time that this device pool was created.

See [“Creating device pools”](#) on page 516.

See [“Adding devices to a device pool”](#) on page 517.

Policies and templates

This chapter includes the following topics:

- [About policies and templates](#)
- [About template rules](#)
- [Setting template rules](#)
- [About creating jobs using policies and selection lists](#)
- [About duplicate backup set templates](#)

About policies and templates

Policies provide a method for managing backup jobs and strategies. Policies contain job templates, which are job attributes that define how and when Backup Exec processes a job. Templates specify the devices, settings, and schedule for a job, but do not include the selections to be backed up. To create jobs, combine a policy with a selection list.

Policies are useful in a number of situations.

For example, you can set up policies for the following:

- **Rotating media.** If you use the policy wizard to create a policy, you can use the Monthly full backup with weekly and daily backups option to create a Grandfather, Father, Son media rotation scheme.
- **Creating duplicate copies of backup sets.** Set up a policy that contains a backup template and a duplicate backup set template. The duplicate backup set template initiates a job that copies the backup set created by the backup job.
- **Verifying backup sets.** Set up a policy that contains a backup template and a verify backup sets template. The verify backup sets template verifies the integrity of the backup data after the backup is completed. You can schedule

the verify operation to run at any convenient time inside or outside of the backup window.

- **Setting relationships between jobs.** When a policy contains more than one template, you can use template rules to establish relationships between the templates. For example, you can set a template rule so that when one backup job completes, Backup Exec automatically starts another backup job.
- **Exporting media.** Set up a policy that contains an export media template and at least one backup template. Then, set up a template rule that schedules a media export job after the backup completes. You can also select a vault to move the media to after it is exported from the robotic library slots.
- **Creating a synthetic backup.** If you have purchased and installed the Advanced Disk-based Backup Option (ADBO), then you can set up a policy that contains the necessary job templates for creating a synthetic backup.
- **Enabling true image restore of backup sets.** If you have purchased and installed the Advanced Disk-based Backup Option (ADBO), then you can set up a policy that contains the necessary job templates for enabling true image restore of backup sets.

After you combine a selection list with a policy, Backup Exec automatically creates a job for each template in the policy. For example, if a policy contains three templates, Backup Exec will create one job for each template, for a total of three jobs. Policies are reusable, so you can create a single policy and combine it with several different selection lists.

See [“Creating a new policy”](#) on page 522.

See [“Creating a new policy using the Policy Wizard”](#) on page 523.

See [“Editing a policy”](#) on page 525.

See [“Deleting a policy”](#) on page 526.

See [“About the synthetic backup feature”](#) on page 899.

See [“About true image restore”](#) on page 912.

See [“Policy Jobs Summary Report”](#) on page 757.

See [“Backup Set Details by Resource Report”](#) on page 734.

See [“Policy Protected Resources”](#) on page 759.

See [“Resource Backup Policy Performance Report”](#) on page 761.

Creating a new policy

Creating a new policy involves choosing a name and description for the policy, adding templates to the policy, and setting up relationships between templates

(if necessary). After you set up all of the templates for a policy, you should combine the policy with a selection list to create jobs.

See [“Adding a backup template to a policy”](#) on page 530.

See [“Adding an export media template to a policy”](#) on page 537.

See [“Adding a duplicate backup template to a policy”](#) on page 550.

See [“Importing a template into a policy”](#) on page 538.

You can set up the policy manually or use the policy wizard.

See [“Creating a new policy using the Policy Wizard”](#) on page 523.

To create a new policy manually

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Policy Tasks**, click **New policy**.
- 3 Complete the appropriate fields.
See [“Policy properties”](#) on page 524.
- 4 Do one of the following:

To create a new template: ■ Click **New Template**.
■ Select the type of template that you want to add.

To import an existing template: ■ Click **Import Template**.
■ Select the templates that you want to import.

- 5 Click **OK** to start editing the template that you selected.

Creating a new policy using the Policy Wizard

Creating a new policy involves choosing a name and description for the policy, adding templates to the policy, and setting up relationships between templates (if necessary). After you set up all of the templates for a policy, you should combine the policy with a selection list to create jobs.

You can set up the policy manually or use the policy wizard.

See [“Creating a new policy”](#) on page 522.

To create a policy using the Policy Wizard

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Policy Tasks**, click **New policy using wizard**.
- 3 Follow the instructions in the wizard to create a policy.

Policy properties

When you create a new policy, you should give it a name and description that helps you to remember its purpose.

See [“Creating a new policy”](#) on page 522.

Table 12-1 Policy properties

Item	Description
Policy name	Designates a unique name for this policy. You can use a name that describes the type of backup or the resources that this policy protects, such as "Monthly full backup policy" or "My documents daily backup". The policy name that you enter here displays in the Policies list on the Job Setup view.
Policy description	Designates a description of this policy.
New Template	Creates a new template for this policy.
Edit Template	Edits an existing policy template.
Delete Template	Deletes an existing policy template.
Import Template	Imports an existing template to this policy.
New Rule	Creates a new template rule.
Edit Rule	Lets you edit an existing rule.
Delete Rule	Lets you delete an existing rule.

Template Selection options

You can choose a template to add to a new policy.

See [“About using templates in policies”](#) on page 529.

Table 12-2 Template Selection options

Item	Description
Backup Template	Adds a job template for a backup method. See “Adding a backup template to a policy” on page 530.

Table 12-2 Template Selection options (*continued*)

Item	Description
Duplicate Backup Sets Template	Adds a job template for creating duplicate copies of backup sets. See “Adding a duplicate backup template to a policy” on page 550.
Verify Backup Sets Template	Adds a job template for a verify operation to test the integrity of backup data. You can schedule the verify operation to run at any time after a backup job is completed. See “Adding a verify backup sets template to a policy” on page 534.
Export Media Template	Adds a job template for removing media from robotic library slots automatically when the backup job is completed. See “Adding an export media template to a policy” on page 537.
Synthetic Backup Template	Adds a job template for combining data from a baseline backup job and subsequent incremental backup jobs. See “About the synthetic backup feature” on page 899.

Editing a policy

You can change the settings for a policy at any time.

You can also edit any templates that belong to the policy.

See [“Editing a template in a policy”](#) on page 539.

To edit a policy

- 1 On the navigation bar, click **Job Setup**.
- 2 Double-click the policy.
- 3 Edit the policy as necessary.
See [“Policy properties”](#) on page 524.

Deleting a policy

If you no longer need a policy, you can delete it. Before you delete a policy, be certain that you no longer need the jobs that are associated with the policy. Before you can delete a policy, you must remove the association between the selection lists and the policy. When you remove the association between selection lists and policies, any active jobs that were associated with the policy will complete and then will be deleted.

To delete a policy

- 1 On the navigation bar, click **Job Setup**.
- 2 Select the policy that you want to delete.
- 3 Do one of the following:

If selection lists are associated with the policy

Do the following in the order listed:

- In the task pane, under **Policy Tasks**, click **Delete jobs created by policy**.
- Check the check boxes next to the selection list names to delete all of the jobs created by this policy.
- Click **OK**.
- Click **Yes** when prompted to continue.
- Select the policy again that you want to delete, and then in the task pane, under **General Tasks**, click **Delete**.
- When prompted to continue, click **Yes**.

If no selection lists are associated with the policy

Under **General Tasks** in the task pane, click **Delete**.

- 4 Click **Yes** to confirm that you want to delete this policy.

Using an example policy

Backup Exec provides example policies that contain standard settings for the following policy types: media rotation, duplicate backup, synthetic backup, verify, and differential backups for virtual machines. You can copy the example policies and then customize them to meet your needs.

To use an example policy

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Policies** pane, right-click the example policy you want to use, and then click **Copy**.

The Synthetic Backup example policy displays only if you have purchased and installed the Advanced Disk-based Backup Option.

See [“About creating a synthetic backup by copying the example policy”](#) on page 906.

- 3 Do one of the following:

To copy the example policy to this media server: Click **Copy to this media server**.

To copy the example policy to other media servers: Do the following in the order listed:

- Click **Copy to other media servers**.
- Select the media server you want to copy to.
- If the media server doesn't appear on the list, click **Add**, and then enter the media server name.

- 4 To overwrite an existing policy with the same name, check the **Overwrite policies with identical names that already exist on the destination media server** check box.

See [“Copying jobs, selection lists, or policies”](#) on page 554.

Backup Exec creates a new policy named "Copy of <example policy>" and places it in the list of policies on the **Job Setup** view of the media servers that you copied it to.

- 5 Customize the policy and templates as necessary.
For example, you may want to give the policy a new name and description, and change when the templates are scheduled to run.
See [“Adding a backup template to a policy”](#) on page 530.
- 6 When finished, click **OK**.
- 7 Create jobs using this policy and a selection list.
See [“About creating jobs using policies and selection lists”](#) on page 544.

Copy Policy options

You can copy an example policy to one or more media servers.

See [“Using an example policy”](#) on page 526.

Table 12-3 Copy Policy options

Item	Description
Copy to this media server	Copies the example policy to the media server on which the policy currently resides.
Copy to other media servers	Copies the example policy to other media servers.
Name	Indicates the name of the destination media servers to which you can copy the example policy.
Logon Account	Indicates the logon account for each destination media server.
Add	Lets you add a new media server to the list of destinations.
Edit	Lets you edit information about the selected media server.
Remove	Removes the selected media server from the list of destinations.
Import List	Imports a list of media servers.
Overwrite policies with identical names that already exist on the destination media server	Indicates whether you want to overwrite policies on the destination media server if they have the same name as the example policy you are copying.

Re-creating example policies

You can re-create example policies. If an example policy with the default example policy name already exists when you select the **Re-create Example Policies** option, Backup Exec creates another example policy and adds an incremented number to the example policy name. For example, if the Example: Media Rotation 0002 policy exists, Backup Exec creates another example policy called Example: Media Rotation 0003.

To re-create example policies

- ◆ On the **Tools** menu, select **Re-create Example Policies**.

About using templates in policies

Templates are the building blocks of policies. They contain all the settings for a job, except the resources to be backed up. Each policy must contain at least one template. Backup Exec contains the following types of templates:

Table 12-4 Types of templates

Template type	Description
Backup	Use this template to create backup jobs, such as full, incremental, and differential. See “Adding a backup template to a policy” on page 530.
Duplicate backup set	With this type of template, you can use a staging strategy to create duplicate copies of backup sets. It allows multiple levels of data duplication, either within the backup window or outside of the backup window. See “About duplicate backup set templates” on page 548.
Verify backup sets	Use this template for a verify operation to test the integrity of backup data. You can schedule the verify operation to run at any time after a backup job is completed. See “About the verify backup sets templates” on page 533.
Export media	Use this template to set up an export media utility job that runs automatically after a backup or duplicate backup set job completes. The export media job either moves media from robotic library slots into a portal or displays an alert reminding you to remove the media from a slot. You can also select a vault to move the media to after it is exported from the robotic library slots. See “Adding an export media template to a policy” on page 537.

Table 12-4 Types of templates (*continued*)

Template type	Description
Synthetic backup	<p>This template is available only with the Backup Exec Advanced Disk-based Backup Option.</p> <p>See “About the synthetic backup feature” on page 899.</p>

You can copy backup templates from one policy to another using the **Import Template** option. You can save time by importing templates that contain all or most of the settings you want to use. After you import templates, you can give the template a unique name and change any of the settings.

See [“Exporting expired media from a robotic library”](#) on page 492.

See [“Importing a template into a policy”](#) on page 538.

Adding a backup template to a policy

Each policy you create must include at least one template. The templates include the information that Backup Exec needs to run jobs. Creating a backup template is similar to creating a backup job. You select the device and media that you want to use, the settings for the job, and the schedule for the job. However, in templates you do not select resources to back up. After a policy is complete, create a job by combining the policy with the selection list that includes the resources you want to back up.

See [“About creating jobs using policies and selection lists”](#) on page 544.

To add a backup template to a policy

- 1 Set up a new policy.
 See [“Creating a new policy”](#) on page 522.
- 2 On the **New Policy** dialog box, click **New Template**.
- 3 On the **Template Selection** dialog box, select **Backup Template**, and then click **OK**.
- 4 In the **Properties** pane, under **Destination**, click **Device and Media**. Complete the Device and Media options.
 See [“Device and media options for backup jobs and templates”](#) on page 339.

- 5 In the **Properties** pane, under **Settings**, click **General**. Complete the General options for this template.
See [“General options for backup jobs and templates”](#) on page 342.
If the Advanced Disk-based Backup Option (ADBO) is installed, and you want to create a policy for synthetic backup or true image restore, you must select the option **Collect additional information for synthetic backup and for true image restore**.
See [“Creating a synthetic backup by adding templates to a policy”](#) on page 907.
- 6 In the **Properties** pane, under **Settings**, click **Advanced**. Complete the Advanced options for this template.
See [“Advanced options for backup jobs”](#) on page 348.
- 7 In the **Properties** pane, under **Settings**, click **Pre/Post Commands**. Complete the Pre/Post Command options.
See [“Pre/post commands for backup or restore jobs”](#) on page 352.
- 8 In the **Properties** pane, under **Settings**, click **Network and Security**.
See [“Network and Security backup options”](#) on page 407.
If the Central Admin Server Option (CASO) is installed, an option displays to allow managed media servers to use any network interface to access remote agents.
See [“Enabling managed media servers to use any available network interface card ”](#) on page 1540.
- 9 In the Properties pane, under Settings, select additional options for this job as needed:
 - If you are using the Advanced Open File Option, select **Advanced Open File** and complete the options.
See [“Advanced Open File options ”](#) on page 949.
 - If you want to use offhost backup, select **Advanced Disk-based Backup** and complete the options.
See [“Backup options for the Advanced Disk-based Backup Option”](#) on page 926.
 - If you are backing up other platform types or database agents, such as NetWare, Exchange, SQL, or SharePoint, select the platform type or database agent. Refer to the chapter for that item for instructions on completing the options.
 - If you want Backup Exec to notify someone when the backup job completes, select **Notification**.

See [“Notification options for jobs”](#) on page 682.

- If you want to prevent certain files or file types from being included in the backup, select **Exclusions**.

See [“Exclusions options”](#) on page 532.

10 Set the schedule for the template.

See [“Schedule properties for a template”](#) on page 532.

11 Click **OK**.

Exclusions options

You can exclude a certain file or types of files when you add a backup template to a policy.

See [“Adding a backup template to a policy”](#) on page 530.

Table 12-5 Exclusions options

Item	Description
Edit	Lets you edit the Excludes selection list to add or remove files and folders. You can also edit the selection criteria.
Insert	Lets you create criteria and settings for the Excludes selection list.
Delete	Lets you delete criteria and settings from the Excludes selection list.

Schedule properties for a template

Schedule properties for templates determine when a job should run when it is created using the template.

Table 12-6 Schedule properties for a template

Item	Description
Run now and run according to rules for this template	<p>Runs the job immediately after a selection list is combined with the policy in which this template is contained. If the template is part of a template rule, the job also runs according to the template rule.</p> <p>If you select this option to change the schedule for an existing template in a policy, the existing associated jobs are not run immediately. This prevents you from inadvertently performing a run now operation on all jobs that were created when the policy was combined with selection lists. The jobs run only according to the rules for the template.</p>
Run according to schedule and run according to rules for this template	<p>Configures a schedule for a recurring job. Use the Edit Schedule Details option to set the schedule.</p> <p>See “Scheduling jobs” on page 356.</p> <p>If the template is part of a template rule, the job also runs according to the template rule.</p>
Run only according to rules for this template	<p>Configures the job to run based on a template rule. For example, with the After <Template A> completes, start <Template B> template rule, <Template B> will run whenever <Template A> completes.</p>
Submit job on hold	<p>Submits the jobs that are created using this template with an on-hold status. You should select this option if you want to submit the job, but do not want it to run until you change the job’s hold status.</p>

About the verify backup sets templates

Verify operations test the integrity of data after a backup. Symantec recommends that you verify all backups. By default, Backup Exec includes a verify operation to run immediately after a backup. You can enable or disable the default verify operation using the **Verify after backup** option. You can optionally create a verify backup sets template to schedule and run the verify operation independent of the source backup job.

You can use a verify backup sets template to schedule the verify operation to run outside of your backup window. Running the verify operation outside of your backup window may be helpful if your network resources are scarce. If you have difficulty completing backups within the allotted time window, you can schedule the verify operation to run at a more convenient time.

A verify backup sets template is especially beneficial if you use Backup Exec's Deduplication Option. If you use the default **Verify after backup** option, Backup Exec must perform the verify operation on both the server-side and the source-side. Backup Exec sends the backup sets across the network to verify them. The process can be time consuming depending on the amount of data you are verifying and your network configuration. If you schedule the operation using a verify backup sets template, the operation runs locally and more efficiently.

See [“About using templates in policies”](#) on page 529.

See [“Adding a verify backup sets template to a policy”](#) on page 534.

See [“About the Deduplication Option”](#) on page 1568.

Adding a verify backup sets template to a policy

By default, Backup Exec includes a verify operation to run immediately after a backup. You can enable or disable the default verify operation using the **Verify after backup** option. If you want to run the verify operation independent of the source backup job, you can optionally create a verify backup sets template instead of using the **Verify after backup** option.

See [“About the verify backup sets templates”](#) on page 533.

To add a verify backup sets template to a policy

- 1 Set up a new policy.
See [“Creating a new policy”](#) on page 522.
- 2 Set up a backup template, which is the media-producing template that is the object of the verify job.
See [“Adding a backup template to a policy”](#) on page 530.
- 3 In the **New Policy** dialog box, click **New Template**.
- 4 In the **Template Selection** dialog box, select **Verify Backup Sets Template**, and then click **OK**.

The verify backup sets template displays only if the policy contains a backup template.

- 5 Select the backup template that you want to verify.
For example, if you want to verify data after the monthly full backup completes, select the template for monthly full backups as the source template.
See [“Verify Backup Sets Template properties”](#) on page 535.
- 6 In the **Properties** pane, under **Settings**, click **General**.
- 7 Select the appropriate options.
See [“Verify Backup Sets Template General properties”](#) on page 535.
- 8 If you want to set up notification for this job, in the **Properties** pane, under **Settings**, click **Notification**.
See [“Notification options for jobs”](#) on page 682.
- 9 Set the schedule for the template.
See [“Schedule properties for a template”](#) on page 532.
- 10 Click **OK**.

Verify Backup Sets Template properties

You can perform a verify operation after a backup to test the integrity of the data.
See [“Adding a verify backup sets template to a policy”](#) on page 534.

Table 12-7 Verify Backup Sets Template properties

Item	Description
Template Name	Indicates the name of the job template that is the source of the media that you want to verify.
Job Type	Indicates the template's job type.

Verify Backup Sets Template General properties

You can perform a verify operation after a backup to test the integrity of the data.
See [“Adding a verify backup sets template to a policy”](#) on page 534.

Table 12-8 Verify Backup Sets Template General properties

Item	Description
Template Name	Indicates the name for the verify template you want to create.

Table 12-8 Verify Backup Sets Template General properties (*continued*)

Item	Description
<p>Allow this job to have direct access to the device</p>	<p>Controls which network and computer resources are used to perform the verify operation.</p> <p>If you select this field, the remote agent performs the verify operation. Select this field when the deduplication device exists on the computer on which the remote agent is installed.</p> <p>If you do not select this field, the media server performs the verify operation. The media server should perform the verify operation if the deduplication device exists on the media server.</p>

About export media templates

You can use the export media template to set up an export media utility job that runs automatically after a backup or duplicate backup set job completes. If the targeted device is a library with a portal, the export media job moves the media from its slot in the portal and generates an alert reminding you to remove the media from the portal. You can select a vault to move the media to after it is exported from the robotic library slots. If the device is a library that does not have a portal, the export media job will generate an alert reminding you to remove the media from the indicated slot.

The export media template must be part of a multi-template policy. There must be a source template that produces media and initiates the export job. For example, if you want to export media after your monthly full backup completes, set up a policy with a backup template for the monthly full backup job and set up an export template to run after the monthly full job completes. If a job requires multiple pieces of media, the export media job starts after the source backup job completes, not after each piece of media is filled.

When you create an export media template, Backup Exec automatically adds the After <Template A> completes, start <Template B> to export media template rule to the policy. Backup Exec replaces <Template A> with the name of the template that you select as the source for the export media template, such as a backup template. Backup Exec replaces <Template B> with the export media template.

See [“Adding an export media template to a policy”](#) on page 537.

Adding an export media template to a policy

You can use the export media template to set up an export media utility job that runs automatically after a backup or duplicate backup set job completes.

See [“About export media templates”](#) on page 536.

When you complete this procedure, Backup Exec adds the template rule called After <Template A> completes, start <Template B> to export media. You can add another template or combine the policy with a selection list.

See [“About creating jobs using policies and selection lists”](#) on page 544.

To add an export media template to a policy

- 1 Set up a new policy.
See [“Creating a new policy”](#) on page 522.
- 2 Set up a backup template, which will be the media-producing template that is the source for the export job.
- 3 On the **New Policy** dialog box, click **New Template**.
- 4 On the **Template Selection** dialog box, select **Export Media Template**, and then click **OK**.

The Export Media template displays only if the policy contains a backup template.

- 5 Select the media-producing template to use as the source for the export job.
For example, if you want to export media after the monthly full backup completes, select the template for monthly full backups as the source media set template.
- 6 In the **Properties** pane, under **Settings**, click **General**. Type a name for this export media template.
- 7 To move the media to a media vault after a successful export, in the **Properties** pane, under **Settings**, click **Options**, and select a media vault.
See [“Scanning bar code labels to move media”](#) on page 255.
- 8 If you want to set up notification for this job, in the **Properties** pane, under **Settings**, click **Notification**.

See [“Sending a notification when a job completes”](#) on page 682.

The export media job must run according to the template rule, so you do not need to set any schedule options for this job.

- 9 Click **OK**.

Export Media Template properties

You can use the export media template to set up an export media utility job that runs automatically after a backup or duplicate backup set job.

See [“Adding an export media template to a policy”](#) on page 537.

Table 12-9 Export media template properties

Item	Description
Template Name	Displays the name of the job template that is the source of the media that you want to export.
Job Type	Displays the template's job type.

Importing a template into a policy

Importing templates makes template creation easier. If an existing template contains many of the settings that you want to use again, you can import the existing template into a policy instead of manually creating a new template and duplicating the settings. Backup Exec does not import any template rules that are associated with the imported template.

After you import a template into a policy, you can change the template settings as needed. Backup Exec copies the templates into the policy and provides the default name of <template name> <number>, where <number> indicates that this is the second copy of this template, or the third copy, etc.

See [“Editing a template in a policy”](#) on page 539.

To import a template into a policy

- 1 On the navigation bar, click **Job Setup**.
- 2 If you want to import a policy into an existing policy, in the **Policies** section, double-click the policy.

If you want to create a new policy and then import an existing template into it, in the task pane, under **Policy Tasks**, click **New policy**. Enter a policy name and description.

- 3 Click **Import Template**.

See [“Import Template options”](#) on page 539.

- 4 Select the templates you want to import.
- 5 Click **OK**.

Import Template options

You can import a template into a policy instead of creating a new template.

See [“Importing a template into a policy”](#) on page 538.

Table 12-10 Import Template options

Item	Description
Policy Name	Identifies the existing policy to which the template belongs.
Template Name	Identifies the existing template.
Job Type	Specifies the type of job to which the template applies.

Editing a template in a policy

You can edit a template at any time.

To edit a template in a policy

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Policies** section, double-click the policy that contains the template you want to edit.
- 3 Select the template from the list that displays in the **Job templates** area.
- 4 Click **Edit Template**.
- 5 Change the template properties as needed.

Deleting a template from a policy

When you delete a template from a policy, it is permanently removed from Backup Exec. In addition, Backup Exec deletes any scheduled jobs that were created from the policy that contained the deleted template. Any active jobs that were created from the policy will complete and then will be deleted.

Do not delete a template from a policy unless you are certain that you no longer need the jobs associated with the template. If a policy contains only one template, delete the policy instead of the template.

See [“Deleting a policy”](#) on page 526.

To delete a template from a policy

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Policies** section, double-click the policy.
- 3 Select the template from the list that displays in the **Job templates** area.
- 4 Click **Delete Template**.
- 5 Click **Yes** to confirm that you want to delete this template.

About template rules

Template rules enable you to set up relationships between templates in a policy. For example, you can use template rules to determine which job should be processed first if a start time conflict exists, or to determine which job should start when another job starts, ends, or fails.

When you set a template rule, you must select the rule you want to use and the templates to which the rule will be applied. Backup Exec displays different rules depending on the types of templates that exist in the policy. For example, if a policy contains a duplicate backup template, Backup Exec displays rules for duplicating backup sets. If Backup Exec automatically adds a template rule to a template, you should not delete that rule.

See [“Setting template rules”](#) on page 542.

The following template rules are available:

Table 12-11 Template Rules

Rule	Description
If start times conflict, <Template A> supersedes <Template B>.	Ensures that if two templates in the same policy have the same start time, <Template A> will run first, and <Template B> will run according to the schedule set for it, after the <Template A> job completes. For example, you set a weekly backup to run every Saturday at 5:00 p.m. and a daily backup to run every day at 5:00 p.m. On Saturday, both the weekly backup and the daily backup are scheduled to run at 5:00 p.m. If you set the weekly backup as <Template A> and the daily backup as <Template B>, the weekly backup will run at 5:00 p.m. on Saturday. The daily backup will not run on Saturday, but it will run on Sunday at 5:00 p.m., according to its schedule.

Table 12-11 Template Rules (*continued*)

Rule	Description
If start times conflict, <Template A> will start and upon completion, starts <Template B>.	Ensures that if two templates in the same policy have the same start time, <Template A> will start first. After <Template A> completes, <Template B> will start automatically.
After <Template A> starts, also start <Template B>.	Allows you to run two jobs simultaneously.
After <Template A> completes, start <Template B>.	Starts a job automatically after a selected job completes. The second job will start regardless of whether the first job completes successfully. For example, if the first job fails, the second job will run.
If <Template A> successfully completes, start <Template B>.	Starts a job automatically after a selected job successfully completes. The second job will not start if the first job fails.
If <Template A> fails, start <Template B>.	Starts a new job automatically if a selected job fails.
<Template A> must complete at least once before any other templates will be allowed to start.	Ensures that you complete a baseline backup job before any other synthetic backups begin. This rule is used when baseline backups are needed, such as with synthetic backup templates. If you set the baseline backup template as the <Template A> in this rule, you can ensure that no other synthetic backup jobs will run until the baseline backup completes.
Run <Template A> only once.	Ensures that if you set up a baseline backup, it needs to run only once. This rule applies to synthetic backups.
Duplicate all backup sets that were created by <Template A> using <Template B> as scheduled.	Displays only if the policy contains a duplicate backup template and the template has a schedule associated with it. The rule applies to the duplicate backup template. Backup Exec automatically adds this rule to the policy if you set the schedule for the template to Run now and run according to rules for this template or Run according to schedule and run according to rules for this template. With this rule, you can set up the data duplication job to run outside of the backup window.

Table 12-11 Template Rules (*continued*)

Rule	Description
After <Template A> completes, start <Template B> to duplicate the backup set.	Applies to a duplicate backup template and displays only if the policy contains a duplicate backup template. After you create a duplicate backup template, Backup Exec automatically adds this rule to the policy if you set the schedule for the template to Run only according to rules for this template. Backup Exec replaces <Template A> with the template that you selected as the source for the duplicate backup template, and replaces <Template B> with the duplicate template you created. With this rule, it is likely that the duplication job will run during the backup window. If you do not want to run the duplication during the backup window, use the Duplicate all backup sets that were created by <Template A> using <Template B> as scheduled rule.
After <Template A> completes, start <Template B> to export media.	Applies to export media templates and displays only if the policy contains an export media template. After you create an export media template, Backup Exec automatically adds this rule to the policy. Backup Exec replaces <Template A> with the template that you selected as the source for the export media template, and replaces <Template B> with the export media template you created.

Setting template rules

You can set up template rules to create relationships between templates in a policy.

See [“About template rules”](#) on page 540.

To set a template rule

- 1 On the navigation bar, click **Job Setup**.
- 2 If you want to set template rules for existing templates in an existing policy, in the **Policies** section, double-click the policy.
 If you want to create a new policy, in the task pane, under **Policy Tasks**, click **New policy**. Enter a policy name and description. Then create a new template.
- 3 On the **Policy Properties** screen, click **New Rule**.
- 4 Complete the appropriate fields that display.
 See [“Template Rule properties”](#) on page 543.
- 5 Click **OK**.

Changing template rules

You can change the rules for a template at any time.

To change a template rule

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Policies** section, double-click the policy.
- 3 On the **Policy Properties** screen, select the rule you want to change and click **Edit Rule**.
- 4 Change the template rule as needed.
See “[Template Rule properties](#)” on page 543.
- 5 Click **OK**.

Deleting template rules

You can delete template rules that you added to templates. You should not delete template rules that Backup Exec added to a template automatically. For multi-stage backup templates, at least one template rule must exist. You should not delete template rules for export media templates.

To delete a template rule

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Policies** section, double-click the policy.
- 3 On the **Policy Properties** screen, select the rule you want to delete and then click **Delete Rule**.

Template Rule properties

You can set up template rules to create relationships between templates in a policy. See “[Setting template rules](#)” on page 542.

Table 12-12 **Template Rule** properties

Item	Description
Template rule	Designates the template rule that you want to apply.
Template A	Designates the template that you want to insert into the <Template A> slot in the template rule you selected. Template A is usually the trigger for template B. For example, in the template rule After <Template A> completes, start <Template B>, template A must complete before Backup Exec will start template B.

Table 12-12 **Template Rule** properties (*continued*)

Item	Description
Template B	Designates the template that you want to insert into the <Template B> slot in the template rule you selected. Some rules do not require more than one template. If another template is not required for a template rule, you cannot select a template from this option.

About creating jobs using policies and selection lists

After you create a policy and set up templates in it, you should combine the policy with a selection list to create jobs. When a policy is combined with a selection list, Backup Exec creates jobs based on the settings in the templates. You can combine a policy with several selection lists, and combine a selection list with several policies. You can create new jobs by selecting a policy and then selecting the selection lists to combine with it, or by selecting a selection list and then selecting the policies to combine with it. Backup Exec creates a job for each template and each selection list. For example, if you combine a policy that contains three templates with two selection lists, Backup Exec creates six jobs; three jobs for one selection list and three jobs for the other selection list.

When you create a backup selection list, you can set a time range when the resources in the list will be available for backup. The time range is called the availability window. When you combine a selection list with a policy, Backup Exec compares the schedule of each template in the policy with the availability window for the selection list. If the template schedules do not fall within the availability window, Backup Exec will not create jobs for the policy. When setting the schedule for templates, be sure that the schedule overlaps the availability window for the resources you want to back up with the templates.

See [“Creating new jobs for a policy”](#) on page 544.

See [“Creating new jobs for a selection list”](#) on page 545.

Creating new jobs for a policy

You can create new jobs by combining a policy with a selection list.

See [“About creating jobs using policies and selection lists”](#) on page 544.

To create new jobs for a policy

- 1 On the navigation bar, click **Job Setup**.
- 2 Select the policy for which you want to create jobs.

- 3 Under **Policy Tasks** in the task pane, click **New jobs using policy**.
- 4 Select the selection lists for which you want to create jobs, and then click **OK**.

Creating new jobs for a selection list

You can create new jobs by combining a selection list with a policy.

See “[About creating jobs using policies and selection lists](#)” on page 544.

To create new jobs for a selection list

- 1 On the navigation bar, click **Job Setup**.
- 2 Select the selection list for which you want to create jobs.
- 3 Under **Selection List Tasks**, click **New jobs using policy**.
- 4 Select the policies for which you want to create jobs, and then click **OK**.

New Jobs Using Policy options

When you combine a policy with a selection list, Backup Exec creates jobs based on the settings in the templates.

See “[About creating jobs using policies and selection lists](#)” on page 544.

Table 12-13 New Jobs Using Policy options

Item	Description
Selected policies	Displays the policy or policies you selected with which to work.
Back Up	Indicates that you want to create backup jobs with the selection list and the selected policy.
Name	Displays the name of selection lists that you can combine with the selected policy or policies.

Viewing the policies that are designated to back up selection lists

You can view a list of policies that are designated to back up a selected selection list.

To view a list of policies that are designated to back up a selected selection list

- 1 On the navigation bar, click **Job Setup**.
- 2 Under **Backup Selection Lists**, right-click the selection list for which you want to view policies.
- 3 Click **View Policies That Back Up Selection List**.

Viewing the selection lists that are designated for backup by policies

You can view a list of the selection lists that are designated for backup by a selected policy.

To view a list of selection lists that are designated for backup by a selected policy

- 1 On the navigation bar, click **Job Setup**.
- 2 Under **Policies**, right-click the policy for which you want to view selection lists designated for backup.
- 3 Click **View Selection Lists Backed Up By Policy**.

Editing the next occurrence of a policy-based job

You can edit the next occurrence of a scheduled policy-based job. Only the next occurrence of a scheduled job can be edited. After the next occurrence of the job completes, the job will resume its original settings as created in the policy. Edits made to the associated policy will overwrite any edits made to the job's next occurrence.

To edit the next occurrence of a scheduled policy-based job

- 1 On the navigation bar, click **Job Monitor**.
- 2 Click the **Job List** tab.
- 3 Right-click the job you want to edit, and then click **Edit Next Run**.

Deleting a job created from a policy

In order to delete a job that was created from a policy, you must remove the association between the selection list and the policy. Backup Exec deletes any scheduled jobs that were created from the policy. Any active jobs that were created from the policy will complete and then will be deleted.

To delete a job created from a policy

- 1 On the navigation bar, click **Job Setup**.
- 2 Select the policy or the selection list from which the job was created.

- 3 If you selected the policy, under **Policy Tasks**, click **Delete Jobs Created By Policy**.
 If you selected the selection list, under **Selection List Tasks**, click **Delete Jobs Created By Policy**.
- 4 Check the check box next to the selection list name, and then click **OK**.
- 5 When prompted to continue with the deletion, click **Yes**.

Delete Jobs Created By Policy options

You must remove the association between a selection list and a policy before you can delete a job that you created using a policy.

See [“Deleting a job created from a policy”](#) on page 546.

Table 12-14 Delete Jobs Created By Policy options

Item	Description
Selected policies	Displays the policy you selected.
Delete Jobs	Indicates the selection list or lists from which the jobs you want to delete were created.
Name	Displays the name of the selection lists.

Renaming a job created from a policy

When you create a job from a policy, Backup Exec automatically names the job. The job name is a combination of the selection list name, policy name, and template name. For example, a job created from a policy might be named Backup Selection List 0001-Policy 001-Backup Template 0001. You can rename jobs that were created from policies to make them more meaningful to you.

To rename a job that was created from a policy

- 1 On the navigation bar, click **Job Setup**.
- 2 Under **Jobs**, right-click the policy-created job that you want to rename.
- 3 Click **Rename**.
- 4 Type the new name in the **Name** field, and then click **OK**.

About duplicate backup set templates

The Duplicate Backup Set template enables you to use a multi-stage backup strategy for backing up data to disk and then copying it to tape. The duplicate backup template does not replace the existing Duplicate Backup Sets option. Instead, it provides an automated, alternate method of duplicating backup sets. It allows for multiple levels of data duplication either within the backup window or outside of the backup window.

Duplicate backups are useful in the following situations:

You want to stage your data

For example, you may want to back up data to disk with a 28 day retention (stage 1), then copy the data to another disk for three months for longer term storage (stage 2), and then move the data to tape for offsite storage (stage 3). A policy for this example staging would include a backup template to back up the data to disk for the 28 days, a duplicate backup set template to copy the data from the original disk to the second disk, and another duplicate backup set template to copy the data from the second disk to the tape. Each of these stages may have a different media set to define the data retention period differently for each stage.

You want to reduce your backup window

For example, create a policy that contains a backup job template that uses the Backup-to-Disk option to back up data to disk during the backup window. Then create a duplicate template to copy the backed up data from disk to tape and schedule the duplication job to occur outside the backup window.

You want to create a duplicate set of backup tapes to store offsite

For example, create a backup template to back up data to either disk or tape. Then create a duplicate template and either set the duplication job to run immediately after the first backup job completes or schedule it to run at a specific time.

If you need to restore data from duplicate backups, you can restore from the source backup or from any of the duplicate backups.

You can use the following methods to configure duplicate backups:

The direct link method

This method requires a policy with one backup template and one duplicate backup template. The direct link is established by the After <Template A> completes, start <Template B> to duplicate the backup set template rule, where <Template A> is the backup template and <Template B> is the duplicate template. The template rule provides a direct link between the backup job and the duplication job.

To set up duplicate backups using this method, you must set up a policy and then do the following:

- Add a backup template with a recurring schedule.
- Add a duplicate template and set the Run only according to rules for this template schedule option.

Backup Exec automatically adds the After <Template A> completes, start <Template B> to duplicate the backup set template rule to the policy.

The incremental duplication method

This method requires at least one backup template and at least one duplicate template. If a policy contains several templates, you can use this method to associate one duplicate backup template to several backup and/or duplicate backup templates. With this method, use the Duplicate all backup sets that were created by <Template A> using <Template B> as scheduled template rule.

To set up duplicate backups using this method, you must set up a policy and then do the following:

- Add a backup template with a recurring schedule.
- Add a duplicate backup template with a recurring schedule.
- Set up a template rule using the Duplicate all backup sets that were created by <Template A> using <Template B> as scheduled template rule.

See [“Adding a duplicate backup template to a policy”](#) on page 550.

Adding a duplicate backup template to a policy

To use a multi-stage backup strategy, you must use a duplicate backup template.

See [“About duplicate backup set templates”](#) on page 548.

When you complete this procedure, you can add another template to the policy or combine the policy with a selection list to create jobs.

To add a duplicate backup template

- 1 Set up a new policy.
See [“Creating a new policy”](#) on page 522.
- 2 Set up a backup template.
See [“Adding a backup template to a policy”](#) on page 530.
- 3 On the **New Policy** dialog box, click **New Template**.
- 4 On the **Template Selection** dialog box, select **Duplicate Backup Sets Template**, and then click **OK**.
- 5 Select the source template, which is the template that will provide the backup data to be copied.
- 6 In the **Properties** pane, under **Destination**, select **Device and Media**.
See [“Device and media options for backup jobs and templates”](#) on page 339.
A duplicate backup set template must use a destination device that can be accessed by the same media server as the device specified for the original backup set.
- 7 In the **Properties** pane, under **Settings**, select **General**.
- 8 Complete the appropriate options.
See [“General properties for new duplicate backup set templates”](#) on page 551.
- 9 In the **Properties** pane, under **Settings**, select **Advanced**.
- 10 Complete the appropriate options.
See [“Advanced options for new duplicate backup set templates”](#) on page 552.
- 11 In the **Properties** pane, under **Settings**, click **Network and Security**.
See [“Network and Security backup options”](#) on page 407.

- 12 If you want to set up notification for this job, in the **Properties** pane, under **Settings**, click **Notification**.

See “[Notification options for jobs](#)” on page 682.

- 13 Set the schedule for the backup job.

See “[Schedule properties for a template](#)” on page 532.

You must set a schedule for the template if you want to use the Use scheduled <Template B> to duplicate all backup sets that were created by <Template A> template rule.

- 14 Click **OK**.

Template properties for new duplicate backup sets

To use a multi-stage backup strategy, you must use a duplicate backup template.

See “[Adding a duplicate backup template to a policy](#)” on page 550.

Table 12-15 Template properties for new duplicate backup sets

Item	Description
Template Name	Indicates the template or templates you want to duplicate.
Job Type	Indicates the type of job the template represents.

General properties for new duplicate backup set templates

To use a multi-stage backup strategy, you must use a duplicate backup template.

See “[Adding a duplicate backup template to a policy](#)” on page 550.

Table 12-16 General properties for new duplicate backup sets job templates

Item	Description
Template name	Indicates the name for this template.
Backup set description	Indicates a description of the data in the backup set.
Preferred source device	Designates the device that was used as the destination device for the original backup job.

Advanced options for new duplicate backup set templates

To use a multi-stage backup strategy, you must use a duplicate backup template.

See [“Adding a duplicate backup template to a policy”](#) on page 550.

Table 12-17 Advanced options for new duplicate backup set job templates

Item	Description
Verify after job completes	Enables Backup Exec to automatically verify that the media can be read after the backup completes. This option is selected by default. Symantec recommends that you verify all backups.
Compression type	Lets you apply any of the following types of compression: <ul style="list-style-type: none"><li data-bbox="559 631 1192 805">■ None. Select this option to copy the data to the media in its original form. If the data was backed up using software compression, then it is copied in its software compression form. Using some form of data compression can help expedite backups and preserve storage media space. Hardware data compression should not be used in environments where the devices that support hardware compression are used interchangeably with devices that do not have that functionality. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive<li data-bbox="559 991 1192 1104">■ Hardware [if available, otherwise none]. Select this option to use hardware data compression (if the storage device supports it). If the drive does not feature data compression the data is backed up uncompressed.

Administering Backup Exec

This chapter includes the following topics:

- [About administrating Backup Exec](#)
- [Copying jobs, selection lists, or policies](#)
- [Editing job properties](#)
- [Job Monitor options](#)
- [Filtering jobs](#)
- [About managing custom filters](#)
- [Viewing the job workload for a media server from the Calendar tab](#)
- [Viewing jobs for specific days on the calendar](#)
- [Managing jobs from the Calendar tab](#)
- [Viewing the Symantec Endpoint Protection Security Summary](#)
- [About error-handling rules](#)
- [How thresholds are used to stall, fail, and recover jobs](#)
- [Setting thresholds to recover jobs](#)

About administrating Backup Exec

Backup Exec includes features that enable you to manage Backup Exec and jobs created in Backup Exec.

You can perform the following operations:

- Copy jobs, selection lists, and policies to local or remote servers.

See [“Copying jobs, selection lists, or policies”](#) on page 554.

- Monitor jobs.
See [“Job Monitor options”](#) on page 557.
- Filter jobs.
See [“Filtering jobs”](#) on page 582.
- Monitor the Symantec Endpoint Protection Security Summary.
See [“Viewing the Symantec Endpoint Protection Security Summary”](#) on page 590.
- Configure error-handling rules.
See [“About error-handling rules”](#) on page 590.
- Configure thresholds to recover jobs.
See [“Setting thresholds to recover jobs”](#) on page 596.

Copying jobs, selection lists, or policies

Backup Exec enables you to copy all jobs (including backup, report, and utility jobs), selection lists, and policies that were created on your media server to the same media server, or to another media server.

To copy jobs, selection lists, or policies to other media servers, the Copy Server Configurations feature must be installed.

See [“About Backup Exec’s standard features”](#) on page 112.

After you select the items to copy and the media server to which you want to copy the items, the operation is queued. The default time-out is five minutes; if the transfer cannot be completed within five minutes, the transfer is terminated and an alert is issued. The queue checks for copy jobs every 60 seconds, and then sends all the copy jobs that are queued.

Backup Exec sends an alert with the job success or failure status along with a log file that allows you to view results. The job log for Copy to Media Server jobs does not display with the other job logs in the Job History.

To copy jobs, selection lists, or policies

- 1 On the media server’s navigation bar, click **Job Setup**.
- 2 Select the job, backup selection list, or policy you want to copy.
- 3 In the task list, under **General Tasks**, click **Copy**.

- 4 Select the media server to which you want to copy the job, selection list, or policy, and then select any applicable overwrite options
 See [“Copy to Media Server options”](#) on page 555.
- 5 Click **OK**.

Copy to Media Server options

Backup Exec enables you to copy all jobs (including backup, report, and utility jobs), selection lists, and policies that were created on your media server to the same media server, or to another media server.

See [“Copying jobs, selection lists, or policies”](#) on page 554.

Table 13-1 Copy to Media Server options

Item	Description
Copy to this media server	Indicates that you want to copy data to this media server.
Copy to other media servers	Indicates that you want to copy data to a different media server. In the Destination media servers field, you must select the media server you want to copy to.
Destination media servers	Indicates the media server you want to copy to, if you are copying to a different media server. If the media server doesn't appear on the list, you can add it.
Overwrite jobs with identical names that already exist on the destination media server	Overwrites an existing job, selection list, or policy that has the same name.
Overwrite logon accounts used by this job that already exist on the destination server	Overwrites the logon accounts for an existing job that has the same name. This option only appears if you are copying a job to another media server.
Add	Lets you add a media server to the Destination media servers list.
Edit	Lets you edit information for a selected media server, such as logon account information.
Remove	Removes a selected media server from the Destination media servers list.

Table 13-1 Copy to Media Server options (*continued*)

Item	Description
Import List	Lets you import a list of media servers to the Destination media servers list. The list should include only the media server name, with one per line.

Viewing the job log for a copy to media server job

A copy to media server job copies jobs, selection lists, or policies from one media server to another media server. Backup Exec sends an alert with the job success or failure status along with a log file that allows you to view results. The job log for Copy to Media Server jobs does not display with the other job logs in the Job History.

To view the job log for a copy to media server job

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts** or **Alert History**.
- 3 Click the **Source** column heading.
- 4 Locate an alert with "Job" as a Source and "Copy to Media Server job" as a Job Name.
- 5 Right-click the alert, and then select **View Job Log**.
- 6 If a Copy to Media Server alert does not exist, do one of the following:

To enable the alert from the task pane

- In the task pane, under **Alert Tasks**, click **Configure alert categories**.
- Enable the Job Failed and Job Success alert categories.

To enable the alert from the Tools menu

- On the **Tools** menu, click **Options**.
- In the **Properties** pane, under **Settings**, click **Preferences**.
- Check **Automatically display new alerts**.

Editing job properties

You can edit existing job properties.

To edit job properties

- 1 On the navigation bar, click **Job Setup**.
- 2 On the **Jobs** pane, click the job you want to edit.
- 3 In the task pane, under **General Tasks**, click **Properties**.

Job Monitor options

Backup Exec’s **Job Monitor** lets you monitor and perform tasks on the active, scheduled, or completed jobs that have been submitted for processing.

The **Job Monitor** provides the following tabs:

Table 13-2 Job Monitor options

Tab	Description
Job List	<p>Displays the active jobs and the scheduled jobs in the Current Jobs pane. The Job History pane displays the jobs that are successful, completed with exceptions, failed, and canceled.</p> <p>See “Viewing properties for active jobs” on page 557.</p> <p>See “Viewing the properties for completed jobs” on page 572.</p>
Calendar	<p>Displays the scheduled, active, and completed jobs in a daily, weekly, or monthly view.</p> <p>See “Viewing the job workload for a media server from the Calendar tab” on page 588.</p>
Security Summary	<p>Displays a summary from Symantec Endpoint Protection of the viruses that were found, and potential threats and risks to the media server.</p> <p>Note: This tab appears only if the Symantec Endpoint Protection Manager component is installed.</p> <p>See “Viewing the Symantec Endpoint Protection Security Summary” on page 590.</p>

See [“About managing custom filters”](#) on page 582.

Viewing properties for active jobs

View properties of active jobs in the Job Monitor.

To view active job properties

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, click the active job that you want to view.
- 3 In the task pane, under **General Tasks** , click **Properties**.
 See “[Job activity options](#)” on page 558.

Job activity options

View properties of active jobs in the **Job Monitor**.

See “[Viewing properties for active jobs](#)” on page 557.

Table 13-3 Job activity options

Item	Description
Job name	Shows the job name that was entered during job configuration.
Job type	Shows the type of job that was submitted for processing.
Job log	Shows the file name of the job log. The job log cannot be viewed until the job has completed. The job log is located in Program Files\Symantec\Backup Exec\Data.
Status	Shows the status of the operation. See “ Active job statuses ” on page 563.
Current Operation	Shows the type of operation that is currently in progress (Backup, Catalog, Restore, Verify, etc.).
Created On	Shows the type of server on which the job was created, either a central administration server or a managed media server. This only displays if you have a central administration server or managed media server.
Server name	Shows the name of the media server that is processing the job.
Device name	Shows the name of the storage device that is processing the job. Only data from the first stream displays for multistream jobs.

Table 13-3 Job activity options (*continued*)

Item	Description
Source	<p>Shows the name of the media or the share that is being processed.</p> <p>The icon field to the left of the field name displays either of the following:</p> <ul style="list-style-type: none"> ■ A disk drive icon when a backup operation is running. ■ A tape drive icon when a restore operation or a verify operation is running. <p>Only data from the first stream displays for multistream jobs.</p>
Destination	<p>Lists the location where the data is being written.</p> <p>The icon field to the left of the field name displays either of the following:</p> <ul style="list-style-type: none"> ■ A tape device icon when a backup operation is running. ■ A disk drive icon when a restore operation is running. <p>Only data from the first stream displays for multistream jobs.</p>
Current directory	<p>Lists the name of the current directory that is being processed.</p> <p>The icon field to the left of the field displays either of the following:</p> <ul style="list-style-type: none"> ■ A folder if the active job is a backup or restore operation. ■ No icon, if the active job is not a backup or restore operation, but a job such as an Erase or Format operation. <p>Only data from the first stream displays for multistream jobs.</p>
Current file	<p>Lists the name of the current file that is being processed.</p> <p>The icon field to the left of the field name displays either of the following:</p> <ul style="list-style-type: none"> ■ A page, if the active job is a backup or restore operation. ■ No icon, if the active job is not a backup or restore operation, but a job such as an Erase or Format operation. <p>Only data from the first stream displays for multistream jobs.</p>
Media server	<p>Lists the name of the media server on which this job is running.</p> <p>If the Central Admin Server Option is installed, this media server is the managed media server that the central administration server has delegated this job to.</p> <p>See “About managing jobs in CASO” on page 1559.</p>

Table 13-3 Job activity options (*continued*)

Item	Description
Delegation status	<p>Indicates the current status of a job that is being delegated from the central administration server to the managed media server. This option appears only if the Central Admin Server Option is installed.</p> <p>The following statuses may appear, where <x> is replaced with the name of the managed media server:</p> <ul style="list-style-type: none"> ■ Preparing to delegate job to <x> ■ Delegating job to <x> ■ Job has been delegated to <x> ■ Job has been received by <x> ■ Job is actively running on <x> ■ Job has completed on <x> ■ Error in delegating job ... re-submitting job to <x> <p>See “About managing jobs in CASO” on page 1559.</p>
Directories	Indicates the number of directories that have processed.
Files	Indicates the number of files that have processed.
Skipped files	Indicates the number of files that were skipped during the operation.
Corrupt files	Indicates the number of corrupt files that were encountered during the operation.
Files in use	Indicates the number of files that were in use during the operation.
Job rate	Indicates the number of megabytes that were processed per minute.
Bytes	Indicates the number of bytes that were processed.
Start time	Indicates the time when the operation started.
Elapsed time	Indicates the length of time that has elapsed since the operation started.
Percent complete	<p>Indicates the percentage of the job that has completed. This option appears only if Display progress indicators for backup jobs is selected as a preference.</p> <p>See “Application settings for Backup Exec preferences” on page 197.</p>
Estimated total bytes	<p>Indicates the total number of bytes that is estimated for the backup job during a prescan. This option appears only if Display progress indicators for backup jobs is selected as a preference.</p> <p>See “Application settings for Backup Exec preferences” on page 197.</p>

Table 13-3 Job activity options (*continued*)

Item	Description
Estimated time remaining	Indicates the estimated time it will take for the job to complete. This option appears only if Display progress indicators for backup jobs is selected as a preference. See “ Application settings for Backup Exec preferences ” on page 197.
Note	Indicates that the option to show job estimates is not selected. This option appears only if Display progress indicators for backup jobs is selected as a preference. See “ Application settings for Backup Exec preferences ” on page 197.

Searching for text in the job history or job properties

You can search for specific text in the job history or the job properties log.

To search for text in the job history or job properties

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, click the active job that you want to view.
- 3 In the task pane, under **General Tasks**, click **Properties**.
See “[Job activity options](#)” on page 558.
- 4 Click **Find**.
- 5 Enter the text that you want to find.
See “[Find options](#)” on page 561.
- 6 Click **Next** to find the next occurrence of the text.

Find options

You can search for specific text in the job history or the job properties log.

See “[Searching for text in the job history or job properties](#)” on page 561.

Table 13-4 Find options

Item	Description
Find	Indicates the text that you want to find.

Table 13-4 Find options (*continued*)

Item	Description
Match whole word only	Indicates that you want to search for the whole word you typed. If you do not select this option, Backup Exec finds the text that includes part of the word. For example, if you search for the word "file" and do not select this option, Backup Exec finds all occurrences of "file", "files", "filed", and any other words that contain "file". If you do select this option, Backup Exec finds only the occurrences of "file".
Match case	Indicates that you want to use the exact capitalization for the word you typed. For example, if you search for the word "File" and select this option, Backup Exec finds all occurrences of "File", but does not find any occurrences of "file".
Highlight all matches	Highlights the text that matches the search criteria.

Canceling an active job

You can cancel a job that is in progress. If the job is scheduled, it runs again at the next scheduled time.

It may take several minutes for a job to cancel. While Backup Exec processes the cancellation of a job, the Cancel Pending status appears in the status column.

To cancel an active job

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, click the active job that you want to cancel.
- 3 In the task pane, under **Active Job Tasks**, click **Cancel**.

To select multiple jobs in the **Job List** view in the **Job Monitor**, select a job, and then press the <Ctrl> or <Shift> keys while you click other jobs that you want to select. This allows you to perform tasks such as Cancel on more than one job at a time, as long as the jobs are of similar type.

- 4 Confirm the cancellation of the job.

Placing all scheduled occurrences of an active job on hold

If an active job is scheduled to run again, you can place the scheduled jobs on hold. The active job continues to run when you place the scheduled occurrences of the active job on hold.

See “[Scheduled job statuses](#)” on page 565.

To place all scheduled occurrences of an active job on hold

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, click the active job for which you want to hold all scheduled occurrences.

To select multiple jobs in the **Current Jobs** pane, select a job, and then press the <Ctrl> or <Shift> keys while you click other jobs that you want to select. This allows you to perform Hold Schedule on more than one job at a time, as long as the state of the jobs are the same.

- 3 In the task pane, under **General Tasks**, click **Hold Schedule**.

Removing the hold on a scheduled job

You can remove the hold on a scheduled job at any time.

To remove the hold on a scheduled job

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, click the occurrence of the job that you want to remove from a hold.
- 3 In the task pane, under **General Tasks**, click **Hold Schedule** to clear the check box.

Active job statuses

Possible statuses for an active job include the following:

Table 13-5 Active job statuses

Item	Description
Running	The operation is underway.
Queued	The job has been initiated, but Backup Exec is actively looking for a suitable drive or media.

Table 13-5 Active job statuses (*continued*)

Item	Description
Cancel Pending	Backup Exec cannot process the Cancel request immediately. This status is displayed until the job is actually canceled. The job is then displayed in Job History with a status of Canceled.
Loading Media	The media is being loaded and positioned on the target device.
Pre-processing	<p>This status can indicate any or all of the following:</p> <ul style="list-style-type: none"> ■ Backup Exec is calculating the amount of data that will be backed up, if the Display progress indicators for backup jobs option is enabled in Preferences. See “Changing application settings for Backup Exec preferences” on page 196. ■ Backup Exec is waiting for a pre- or post-command to complete. ■ Backup Exec is retrieving the set maps and is positioning the tape to the append point location for an append job.
Snapshot processing	Backup Exec is processing a snapshot operation.
Device Paused	<p>The device that the job was sent to is paused.</p> <p>See “Pausing storage devices” on page 446.</p>
Server Paused	<p>The media server is paused.</p> <p>See “Pausing a media server” on page 446.</p>
Stalled	<p>The Backup Exec services have become unresponsive.</p> <p>See “Setting thresholds to recover jobs” on page 596.</p>
Media Request	You must insert media for the job to continue.
Communication Stalled	<p>Communications between the managed media server and the central administration server have not occurred within the configured time threshold.</p> <p>See “Setting communication thresholds and active job status updates for CASO ” on page 1533.</p>

Table 13-5 Active job statuses (*continued*)

Item	Description
No Communication	<p>No communication about jobs is being received at the central administration server from the managed media server. The configured time threshold has been reached.</p> <p>See “Setting communication thresholds and active job status updates for CASO” on page 1533.</p>
Consistency Check	<p>Backup Exec is running a consistency check of the databases before backup.</p>
Updating Catalogs	<p>Backup Exec is updating the catalog information.</p>
Schedule, CPS backup job running	<p>The Exchange logs are being continuously protected with Continuous Protection Server (CPS). The status appears if you check the option Continuously back up transaction logs with Backup Exec Continuous Protection Server on the backup job properties for Exchange backups.</p> <p>See “About continuous protection for Exchange data” on page 1112.</p>

See [“Scheduled job statuses”](#) on page 565.

See [“Completed job statuses”](#) on page 577.

Scheduled job statuses

Possible statuses for scheduled jobs are listed in the following table:

Table 13-6 Scheduled job statuses

Scheduled job status	Description
Blocked by template rule	<p>The scheduled job cannot run because it was created by a policy that contains a job template with the following template rule:</p> <p>< Template A> must complete at least once before any other templates will be allowed to start</p> <p>The job designated as <Template A> in the policy must run before this scheduled job can run.</p> <p>See “Setting template rules” on page 542.</p>

Table 13-6 Scheduled job statuses (*continued*)

Scheduled job status	Description
Invalid Schedule	<p>The scheduled job will not run because of one of the following:</p> <ul style="list-style-type: none"> ■ An associated record in the database is missing. ■ The availability window and the schedule for the selection list being backed up by this job do not have a time in common. <p>See “Setting priority and availability windows for selection lists” on page 307.</p>
Not in Time Window	<p>The job was ready to be sent for processing, but the time window for the job closed. This probably occurred because appropriate destination devices were not available during the common time between the job’s time window and the selection list’s availability window.</p> <p>See “Setting priority and availability windows for selection lists” on page 307.</p>
On Hold	<p>The job has been placed on hold.</p>
Queued	<p>A temporary state that displays when Backup Exec is applying an error-handling rule that is enabled to retry the job.</p> <p>See “Custom error-handling rule for recovered jobs” on page 594.</p>

Table 13-6 Scheduled job statuses (*continued*)

Scheduled job status	Description
Ready	<p>The job is ready to run, but cannot for one of the following reasons:</p> <ul style="list-style-type: none"> ■ Internal error. No devices are available, but the cause is unknown. ■ Invalid job. The job type is unknown; there may be an internal error or the database is corrupted. ■ Invalid target. This is a device type that no longer exists. ■ Media server not available. ■ No license for option name. A license must be purchased and installed on the targeted media server. ■ No media servers are available. ■ No media servers available in media server pool. ■ Specified destination device pool is empty. ■ Specified destination device is not in media server pool. ■ Specified destination device not on local media server. ■ Specified destination device pool on local media server is empty. ■ The destination device cannot be a device pool. ■ The destination device cannot be a media server. ■ There is another job running in the system that is blocking execution of this job. This job will run after the other job completes. ■ Invalid input. ■ Incompatible Resumes. ■ No server license available. ■ No multi-server license available. ■ No Windows license. ■ No Windows server. ■ No NetWare server. ■ Need local media server. ■ Local server is not a media server. ■ No idle devices are available. ■ No eligible devices within the device pool are available. ■ Blocked by an active, linked Duplicate Backup Sets job.
Scheduled	<p>The job is scheduled to run in the future. Scheduled jobs that are linked to another job, such as a job to duplicate backup sets, will not display a scheduled job status.</p>

Table 13-6 Scheduled job statuses (*continued*)

Scheduled job status	Description
Server Paused	<p>The job is ready, but the Backup Exec media server has been paused. No jobs are dispatched while the media server is paused.</p> <p>See “Pausing a media server” on page 446.</p>
Superseded by job x	<p>The scheduled job cannot run because it was created by a policy that contains a job template with the following template rule:</p> <p>If start times conflict, <Template A> supersedes <Template B>.</p> <p>The <Template B> job will run according to the schedule set for it, after the <Template A> job completes.</p> <p>See “Setting template rules” on page 542.</p>
To Be Scheduled	<p>A state that the scheduled job transitions through as it is being sent for processing.</p>

See [“Active job statuses”](#) on page 563.

See [“Completed job statuses”](#) on page 577.

Running a scheduled job immediately

You can run a scheduled job immediately. The job will also run on the next scheduled occurrence.

To run a scheduled job immediately

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, select the scheduled job.
- 3 In the task pane, under **Scheduled Job Tasks**, click **Run now**.

Placing a scheduled job on hold

You can place a scheduled job on hold to prevent the job from running. You can also place the entire job queue on hold to make changes to your environment. Jobs do not run until you change the hold status.

To place a scheduled job on hold

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, select the scheduled job.

To select multiple jobs, select a job, and then press the <Ctrl> or <Shift> keys while you click other jobs that you want to select. This allows you to apply **Hold Schedule** to more than one job at a time, as long as the state of the jobs are the same.

- 3 In the task pane, under **General Tasks**, check **Hold Schedule**.

Removing the hold on a scheduled job

When you remove the hold on a scheduled job, the job then runs according to the schedule.

To remove the hold on a scheduled job

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, select the scheduled job.

To select multiple jobs, select a job, and then press the <Ctrl> or <Shift> keys while you click other jobs that you want to select. This allows you to remove the hold from more than one job at a time, as long as the state of the jobs are the same.

- 3 In the task pane, under **General Tasks**, uncheck **Hold Schedule**.

Placing the job queue on hold

You can place the entire job queue on hold to make changes to your environment. The server is paused to place the job queue on hold. When the job queue is on hold, only active jobs continue to run. No other jobs run until the job queue is taken off hold.

See [“Removing the hold on the job queue”](#) on page 570.

To place the job queue on hold

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the task pane, under **General Tasks**, check **Hold job queue**.
- 3 Click **Yes**.

Removing the hold on the job queue

When you remove the hold on the job queue, the server is unpaused and jobs then run according to the schedule.

To remove the hold on the job queue

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the task pane, under **General Tasks**, uncheck **Hold job queue**.

Changing the priority for a scheduled job

The priority determines the order that jobs run. If two jobs are scheduled to run at the same time, the priority you set determines which job runs first. The priority is changed for all occurrences of the scheduled job.

To change the priority for a scheduled job

- 1 On the navigation bar, select **Job Monitor** or **Job Setup**.
- 2 Do one of the following:

To increase or decrease the priority by one level, for example to increase the priority from Lowest to Low

Do the following in the order listed:

- Select the job.
- In the task pane, click **Increase Priority** or **Decrease Priority**.

To increase or decrease the priority by more than one level, for example to increase the priority from Lowest to Highest

Do the following in the order listed:

- Right-click the job.
- Click **Change Priority**.
- Select the new priority.

Running a test job for a scheduled job

The test run option determines if a scheduled backup will complete successfully. During the test run, the tape capacity, credentials, and media are checked. If the test job determines there is a problem, the job will continue to run and the problem will appear in the job log.

To run a test job for a scheduled job

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, select the active job.

- 3 In the task pane, under **Scheduled Job Tasks**, click **Test run**.
To select multiple jobs in the **Job List** view in the **Job Monitor**, select a job, and then press the <Ctrl> or <Shift> keys while you click other jobs that you want to select. This allows you to perform tasks such as Test Run on more than one job at a time, as long as the jobs are of similar type.
- 4 Enter test run properties for the job.
See [“Setting test run default options”](#) on page 387.
- 5 In the **Properties** pane, under **Settings**, click **Notification** and enter notification information for the items.
See [“Sending a notification when a job completes”](#) on page 682.
- 6 In the **Properties** pane, under **Frequency**, click **Schedule** and then click **Submit job on hold** if you want to submit the job with an on-hold status.
Select this option if you want to submit the job, but do not want the job to run until you change the job’s hold status.
- 7 Click **Run Now** to submit the test run job.

Deleting scheduled jobs

Deleting a scheduled job from the **Job List** tab in the **Job Monitor** removes all scheduled occurrences of the job. To delete only the occurrence of a scheduled job on a specific date, you can edit the schedule to remove that date.

If the job was created by a policy, you must first remove the association between the policy and the selection list.

See [“Deleting a job created from a policy”](#) on page 546.

To delete a scheduled job

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, right-click the scheduled job.
- 3 Click **Delete**.
- 4 If you want to delete the backup selection lists that are associated with the job, click **If the selection list used by this job is no longer in use, delete it**.
- 5 Click **Yes**.

Viewing the properties for completed jobs

You can view detailed job-related properties for each job that has been processed. For some jobs, you can right-click the job and choose to retry the job, or choose to configure a custom error-handling rule for the error that the job failed with.

Errors that are reported in the job log contain hyperlinks that you can click to go to the Symantec Technical Support Web site.

The **Job History** dialog box contains two tabs: **Job History** and **Job Log**. The **Job History** tab provides summary information for the job. The **Job Log** tab provides job and file statistics. Most job logs display in HTML format. However, some logs may display in text.

See [“Completed job statuses”](#) on page 577.

See [“Configuring default job log options”](#) on page 580.

To view the properties for completed jobs

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job list** tab, in the **Job History** pane, select the completed job that you want to view.
- 3 In the task pane, under **General Tasks**, click **Properties**.
- 4 View the information on the **Job History** tab or on the **Job Log** tab.

See [“Job History properties for completed jobs”](#) on page 572.

See [“Job Log properties for completed jobs”](#) on page 575.

Job History properties for completed jobs

The **Job History** tab provides summary information for a job.

See [“Viewing the properties for completed jobs”](#) on page 572.

Table 13-7 Job History properties for completed jobs

Item	Description
Previous	Displays the job history of the previous job that was run as part of this recurring job.
Next	Displays the job history of the next job that was run as part of this recurring job.
Job name	Shows the job name that was entered during job configuration.

Table 13-7 Job History properties for completed jobs (*continued*)

Item	Description
Job type	Shows the type of operation that was performed, such as Back up, Back Up and Delete, Catalog, Restore, or Verify.
Job status	Shows the status of the operation.
Job log	Shows the file name and location of the job log.
Server name	Shows the name of the media server that processed the job.
Selection list name	Shows the name of the selection list that was processed in the job.
Device Name	Shows the name of the device that processed the job.
Target name	Shows the name of the device that was selected during job configuration.
Media set name	Shows the name of the media set that processed the job.
All Media Used	Lists all the media that was used to process the job.
Byte count	Shows the number of bytes that were processed. (This item does not appear in the Job History for catalog jobs.)
Job rate	Shows the amount of data that was backed up per minute for the entire job. (This item does not appear in the Job History for catalog jobs.)
Files	Shows the total number of files that were processed.
Directories	Shows the total number of directories that were processed.
Skipped files	Shows the number of files that were skipped during the operation.
Corrupt files	Lists the number of corrupt files that were encountered during the operation.
Files in use	Lists the number of open files that were encountered during the operation.
Original start time	Shows the time the job was submitted for processing.

Table 13-7 **Job History** properties for completed jobs (*continued*)

Item	Description
Job started	Shows the time the operation started.
Job ended	Shows the time the operation ended.
Elapsed time	Shows the length of time the operation took.
Set type	Lists the type of operation that was performed on the media set, such as Back up, Back Up and Delete, Catalog, Restore, or Verify.
Set status	Shows the status of the operation.
Set description	Shows the job name that was entered during job configuration.
Resource name	Lists the name of the resource for the job.
Logon account	Lists the name of the logon account that was used for the job.
Encryption key	Indicates if an encryption key was used for the job.
Error	<p>Shows the error code, if an error occurred.</p> <p>You can use the job log to locate where the error occurred and to get additional information about the error from the Unique Message Identifier.</p> <p>See “Linking from the job log to the Symantec Technical Support Web site” on page 577.</p> <p>In addition, you can user error-handling rules to enable retry options and final job disposition for jobs when this error occurs.</p> <p>See “About error-handling rules” on page 590.</p>
Agent used	Indicates if a Backup Exec agent was used during the operation.
Advanced Open File Option used	Indicates if the Advanced Open File Option was used during the operation.
Start time	Shows the time the operation started.
End time	Shows the length of time the operation took.

Job Log properties for completed jobs

The **Job Log** tab provides job and file statistics. Most job logs display in HTML format. However, some logs may display in text.

See [“Viewing the properties for completed jobs”](#) on page 572.

Table 13-8 Job Log properties for completed jobs

Item	Description
Job Information	Displays the job server, job name, the date and time the job started, the type of job, and the job log name.
Device and Media Information	Displays the drive name, media label, the overwrite protection and append periods, and the media set that this job was targeted to.
Utility Job Information	Displays information about the slot, bar code, media label, status, and device that the utility job was processed on. See “About creating utility jobs to help manage devices and media” on page 480.
Job Completion Status	Displays the job end time, completion status, error codes, error description, and error category. The job completion section is green, orange, or red, depending on the job status. See “Completed job statuses” on page 577.
Errors	Displays a detailed description of the errors that were encountered during job processing. The errors are grouped by set and labeled. The label includes the operation and the destination resource name for that set. The error section is red in the job log. To locate where the error occurred in the Backup Set Detail Information, click the error text. Then, if additional information on an error is available, click the underlined error code number to go to the Symantec Technical Support Web site.
Exceptions	Displays a detailed description of the minor errors that were encountered during job processing. The exceptions section is orange in the job log.
NDMP Log	Provides details about the NDMP environment variables that were selected for an operation, and about duplicate sets for NDMP.

Viewing the history of a job, policy, or selection list

You can view the history of each active job, scheduled job, policy, and selection list. For job histories, you can view details of each recurring job instance.

For example, if a job has run 20 times, there are 20 job histories for that job. The **View History** dialog box for that job lists all 20 job histories for that job.

To view the history of a job, policy, or selection list

- 1 Do one of the following:

To view the history of a policy or selection list On the navigation bar, click **Job Setup**.

To view the history of an active job, scheduled job, or job history item On the navigation bar, click **Job Monitor**.

- 2 Right-click the item for which you want to view the history.

- 3 Do one of the following:

- For current jobs, on the shortcut menu, click **View History**.
- For job histories, on the shortcut menu, click **View Recurring Job Instances**.

- 4 View the history, and then click **OK**.

Deleting completed jobs

You can delete a job from the **Job Monitor**, or have Backup Exec automatically delete the job history using database maintenance. You can also set a default when you schedule jobs to automatically delete jobs that are set to run once, and that are not created from templates.

If you delete a job, it is removed from the computer and cannot be recovered.

To delete a completed job

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Job History** pane, select the job that you want to delete.

- 3 In the task pane, under **General Tasks**, click **Delete**.

You can select multiple jobs in the **Job list** view in the **Job Monitor** by selecting a job, and then pressing the <Ctrl> or <Shift> keys while you click other jobs that you want to select. This allows you to perform tasks such as Delete on more than one job at a time, as long as the jobs are of similar type.

You can delete up to 2500 jobs from the Job History. If you attempt to delete more than 2500 jobs, you are prompted to continue with the deletion.

- 4 Click **Yes**.

Linking from the job log to the Symantec Technical Support Web site

Errors that are reported in the job log each have a unique code, called a Unique Message Identifier (UMI). These codes contain hyperlinks that you can click to go to the Symantec Technical Support Web site. From the Web site, you can access technical notes and troubleshooting tips that are related to a specific message. Unique Message Identifier (UMI) codes establish unique message codes across all Symantec products.

Some alerts also contain a UMI. For example, if a Warning alert appears when a job fails, the alert includes the UMI code.

See [“Responding to active alerts”](#) on page 654.

You can create or enable an error-handling rule for errors. These rules let you set options to retry or stop a job when the error occurs.

See [“About error-handling rules”](#) on page 590.

To link from the job log to the Symantec Technical Support Web site

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Job History** pane, select the completed job that you want to view.
- 3 In the task pane, under General Tasks , click **Properties**.
- 4 Click **Expand All** to view all of the information that is contained in the topics. Click **Collapse All** to hide the information in the topics.
- 5 Scroll to the Job Completion Status section.
- 6 Click the UMI code, which appears as a blue hyperlink.

Completed job statuses

Possible job completion statuses for jobs that were processed include the following:

Table 13-9 Job completion status

Status	Description
Successful	The job completed without errors.
Completed with exceptions	The job completed, but one of the following types of files was encountered during the operation: <ul style="list-style-type: none">■ In use■ Skipped■ Corrupted
Failed over	The job ran in a cluster environment and was active on one computer, and then the cluster performed a failover and the job was restarted on another computer in the cluster. There are two separate sets of job history when a job is failed over. The first job history will have the Failed over status and the second job history will have the status that is appropriate for the completed job.
Resumed	The status is the same as the failed over status, however the Apply CheckPoint Restart option was selected.
Canceled	The administrator terminated the operation as it was running.
Canceled, timed out	The Enable automatic cancellation feature in the Frequency - Schedule property was enabled and the job was not completed within the specified timeframe. See “Scheduling jobs” on page 356.

Table 13-9 Job completion status (*continued*)

Status	Description
Failed	<p>The operation took place, but one or more significant errors occurred. The job log should indicate what caused the errors so that you can decide if you want to run the job again. For example, if a job failure occurred due to a lost connection during job processing, you could choose to resubmit the job when the connection is restored.</p> <p>If a drive loses power during a backup operation, you should restart the backup job using a different tape. You can restore the data written to the tape up to the point of the power loss, but you should not reuse the tape for subsequent backup operations.</p> <p>A failed job will have an error message in the Errors section of the job log with a navigable link to the Symantec Technical Support Web site.</p> <p>See “Linking from the job log to the Symantec Technical Support Web site” on page 577.</p> <p>A job may fail for the following reasons:</p> <ul style="list-style-type: none"> ■ Devices specified by the job were not available when the job was run. ■ The logon account information used in the backup job is incorrect. Verify the logon account information is valid for the resource being backed up. ■ There was a problem with the storage device when the job was run. ■ The computer being backed up was shut down before or during the backup job.
Recovered	<p>The job was active when the status of the managed media server changed from Communication Stalled to No Communication. The custom error-handling rule for Recovered Jobs was applied to the job.</p> <p>See “Setting communication thresholds and active job status updates for CASO ” on page 1533.</p>
Missed	<p>The job did not run during the scheduled time window. The job is rescheduled to run based on the time window that you configured.</p> <p>See “Setting the time window for a scheduled job” on page 364.</p>

Configuring default job log options

You can configure default options for job logs that specify the amount of detail you want to include in the completed job log. For jobs that produce large job logs, for example, a backup of a considerable number of separate files, you may want to reduce the amount of detail in the job log. The size of the job log increases proportionally to the level of detail configured for the job log.

To configure default job log options

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Job Logs**.
- 3 Select the appropriate options.

See “[Default job log options](#)” on page 580.

Default job log options

You can configure default options for job logs that specify the amount of detail you want to include in the completed job log.

See “[Configuring default job log options](#)” on page 580.

Table 13-10 Default job log options

Item	Description
Summary information only	<p>Includes the following information in the job log:</p> <ul style="list-style-type: none"> ■ Job name ■ Job type ■ Job log name ■ Media server name ■ Storage device ■ Starting date and time ■ Errors encountered ■ Ending date and time ■ Completion statistics <p>This option also includes the name of files that were skipped, the name of the media set, the backup type and results of the verify operation if one was performed.</p>
Summary information and directories processed	Includes summary information and a list of all processed subdirectories in the job log.

Table 13-10 Default job log options (*continued*)

Item	Description
Summary information, directories, and files processed	Includes summary information, processed subdirectories, and a list of all the file names that were processed in the job log.
Summary information, directories, files and file details	Includes Summary information, processed subdirectories, a list of all the file names and their attributes in the job log. This option increases the job log sizes significantly.
Prefix for the job log file name	Indicates a prefix to add to the job logs that are processed. The default prefix is BEX. The job log file name consists of Prefix_ServerName_Count, where Prefix is the label that you enter in this field, ServerName is the name of the media server that ran the job, and Count is the number of job logs that this job has produced.
Attach job logs as html	Attaches the job logs in an HTML format when an email notification is sent.
Attach job logs as text	Attaches the job logs in a text format when an email notification is sent.
Job log path	Shows the current location of the job log. To change the path you can use BE Utility.

About using job logs with vertical applications

The Backup Exec Administration Console provides a view of the job logs in HTML format. If necessary, you can convert the job logs to a text format for use with vertical applications.

To convert a job log file to a text format, type the following at a command prompt from the default directory C:\Program Files\Symantec\Backup Exec\Data, or wherever Backup Exec was installed to:

```
bemcmd -o31 -f"<pathname>\job log filename">
```

For example, to display the job log C:\program files\Symantec\Backup Exec\Data\bex00001.xml in text format to the command prompt, you would type:

```
bemcmd -o31 -f"C:\program files\Symantec\Backup Exec\Data\bex00001.xml"
```

To redirect the job log to a file, you would type one of the following:

```
bemcmd -o31 -f"C:\program files\Symantec\Backup Exec\Data\bex00001.xml" >  
bex00001.txt
```

or

```
bemcmd -o31 -l"bex00001.txt" -f"C:\program files\Symantec\Backup  
Exec\Data\bex00001.xml"
```

Filtering jobs

You can select predefined filters to limit the jobs that appear in the **Job Setup** view, and in the **Job Monitor** view.

To filter jobs

- 1 Do one of the following:

To filter jobs on the Job Setup View Click **Job Setup**.

To filter jobs on the Job List tab Click **Job Monitor**, and then click **Job List**.

To filter jobs on the Calendar tab Click **Job Monitor**, and then click **Calendar**.

- 2 In the **Filter** list, click the filter that you want to use.

About managing custom filters

You can create and edit custom filters for jobs. Backup Exec has an XML file for each pane in which you can create custom filters. Backup Exec stores the custom filters that you create in the following location:

```
\Documents and Settings\\Local Settings  
\Application Data\Symantec Corporation  
\BkupExec.exe_StrongName_qlwvfcithy432w2rcmdl1dn0kfn1fr5rb  
\<BE_Version_Number><BE_Build#>\user.config
```

The logon account name that you use to log on to the computer is appended to each custom filters file. Each Backup Exec user has custom filter files. You can only view the custom filters that you create.

See [“Creating a custom filter for jobs”](#) on page 583.

See [“Creating a custom filter for current jobs”](#) on page 583.

See [“Creating a custom filter for jobs in the job history”](#) on page 585.

See [“Deleting custom filters”](#) on page 587.

See [“Editing custom filters”](#) on page 588.

Creating a custom filter for jobs

You can create custom filters to limit the jobs that appear in job lists.

To create a custom filter for jobs

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Jobs** pane, in the **Filter** list, click **<new custom filter>**.
- 3 Type a unique name and a description for this filter.
- 4 Under **Criteria**, select the type of information on which you want to filter jobs.
- 5 Check **Enable this filter**.
- 6 Check the check boxes for the types of data on which you want to filter. Uncheck the check boxes for the types of data on which you do not want to filter.
- 7 Click **OK**.

New Jobs Custom Filter options

You can create custom filters to limit the jobs that appear in job lists.

Table 13-11 **New Jobs Custom Filter options**

Item	Description
Name	Indicates the unique name of the custom filter.
Description	Describes the filter.
Enable this filter	Indicates that the selected criteria are included in the filter. After you enable a filter, you can select the specific criteria on which to filter.

Creating a custom filter for current jobs

You can create custom filters to limit the current jobs that appear in job lists.

To create a custom filter for current jobs

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Current Jobs** pane, in the **Filter** list, click **<new custom filter>**.
- 3 Type a unique name and a description for this filter.
- 4 Under **Criteria**, select the type of information on which you want to filter jobs.
- 5 Check **Enable this filter**.
- 6 Check the check boxes for the types of data on which you want to filter. Uncheck the check boxes for the types of data on which you do not want to filter.
- 7 Click **OK**.

New Current Jobs Custom Filter options

You can create custom filters to limit the current jobs that appear in job lists. See [“Creating a custom filter for current jobs”](#) on page 583.

Table 13-12 **New Current Jobs Custom Filter options**

Item	Description
Name	Indicates the unique name of the custom filter.
Description	Describes the filter. This description appears on the Custom Filter Management dialog box. However, it does not appear in the Job Monitor view.
Enable this filter	Indicates that the selected criteria are included in the filter. After you enable a filter, you can select the specific criteria on which to filter.
Do not use a date range	Indicates that you do not want to filter the current jobs list based on when jobs are scheduled to run. For example, you enable the Job Type filter and select Backup as the job type. If you select this option, all backup jobs that are scheduled to run on any date appear on the job history list.

Table 13-12 **New Current Jobs Custom Filter options** (*continued*)

Item	Description
Show next	<p>Indicates that you want to filter the current jobs list based on when the jobs are scheduled to run. You can select a specific number of hours or days on which to filter.</p> <p>For example, you enable the Job Type filter and select Backup as the job type. If you select this option and set 24 hours as the range, only backup jobs that are scheduled to run during the next 24 hours appear in the job history list.</p>

Creating a custom filter for jobs in the job history

You can create custom filters to limit the jobs that appear in the job history.

To create a custom filter for jobs in the Job History

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Job List** tab, in the **Job History** pane, in the **Filter** list, click **<new custom filter>**.
- 3 Type a unique name and a description for this filter.
- 4 Under **Criteria**, select the type of information on which you want to filter jobs.
- 5 Check **Enable this filter**.
- 6 Check the check boxes for the types of data on which you want to filter. Uncheck the check boxes for the types of data on which you do not want to filter.
- 7 Click **OK**.

New Job History Custom Filter options

You can create custom filters to limit the jobs that appear in the job history.

See [“Creating a custom filter for jobs in the job history”](#) on page 585.

Table 13-13 **New Job History Custom Filter options**

Item	Description
Name	Indicates the unique name of the custom filter.
Description	Describes the filter. This description appears on the Custom Filter Management dialog box. However, it does not appear in the Job Monitor view.
Enable this filter	Indicates that the selected criteria are included in the filter. After you enable a filter, you can select the specific criteria on which to filter.
Show job histories that backed up or restored data on this server	Indicates that you want to view histories for the jobs that backed up or restored a particular server.
Show job histories with the following job name	Indicates that you want to view job histories for jobs with a particular name.
Do not use a date range	<p>Indicates that you do not want to filter the job history list based on when jobs ran.</p> <p>For example, you enable the Job Type filter and select Backup as the job type. If you select this option, all backup jobs that have been run appear on the job history list.</p>
Show last	<p>Indicates that you want to filter the job history list based on when the jobs ran. You can select a specific number of hours or days on which to filter.</p> <p>For example, you enable the Job Type filter and select Backup as the job type. If you select this option and set 24 hours as the range, only backup jobs that ran in the last 24 hours appear in the job history list.</p>

Table 13-13 **New Job History Custom Filter options** (*continued*)

Item	Description
Show using the following date range	<p>Indicates that you want to filter the job history list based on a particular range of dates. You can filter by date and time.</p> <p>For example, you enable the Job Type filter and select Backup as the job type. If you select this option and set the data range from December 1 to December 8, only the backup jobs that ran from December 1 to December 8 appear in the job history list.</p>
Show only the last occurrence of a job	<p>Indicates that you want to filter the job history list based on the last occurrence of the specified job types.</p> <p>For example, you enable the Job Type filter and select Backup as the job type. If you select this option, only the last occurrence of each backup job appears in the job history list.</p>
Filter job histories that have a byte count	<p>Lets you filter the job history to show only the jobs that have a byte count that is larger or smaller than the size you indicate.</p>
Filter job histories that have an elapsed time	<p>Lets you filter the job history to show only the jobs that have an elapsed time that is longer or shorter than the time you indicate.</p>

Deleting custom filters

You can delete custom filters that you no longer need.

To delete a custom filter

- 1 Do one of the following:

To delete a custom filter from the **Job Setup** view Click **Job Setup**.

To delete a custom filter from the **Job Monitor** view Click **Job Monitor > Job List**.

- 2 In the task pane, under **Custom Filter** tasks, click **Manage custom filters**.

- 3 Select the filter that you want to delete.
- 4 Click **Delete**.
- 5 When prompted to delete the custom filter, click **Yes**.
- 6 Click **Close**.

Editing custom filters

You can change options in your custom filters.

To edit a custom filter

- 1 Do one of the following:

To edit a custom filter from the **Job Setup** view Click **Job Setup**.

To edit a custom filter from the **Job Monitor** view Click **Job Monitor > Job List**.

- 2 In the task pane, under **Custom Filter** tasks, click **Manage custom filters**.
- 3 Select the filter that you want to edit.
- 4 Click **Edit**.
- 5 Edit the custom filter options.
- 6 Click **OK**.
- 7 Click **Close**.

Viewing the job workload for a media server from the Calendar tab

You can view a media server's job workload for a month, for a week, or for a day.

The month and the week views list the number of jobs and the number of job instances. You can see at a glance what the scheduled work load is for any day. Job details do not appear on these views.

The day view provides a graphical view of the scheduled work load, and lists jobs for that day in chronological order. Available tasks for the jobs appear in the task pane. A preview pane provides detailed information for each job, such as final status, and the rate and byte count for job histories.

Note: Custom filters are not available in the Calendar views.

To view the job workload for a media server from the Calendar tab

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Calendar** tab, do one of the following:

Click the day icon.



Click the week icon.



Click the month icon.



Viewing jobs for specific days on the calendar

Use the calendar on the task pane to view the number of jobs for a specific day. The day can be any number of days, weeks, or months in the future or in the past.

To view jobs for specific days on the calendar

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Calendar** tab, in the task pane, click an arrow key to set the calendar forward or backward in monthly increments.
- 3 To go to the current date, at the bottom of the calendar, click **Today**.

Managing jobs from the Calendar tab

Some common tasks are available on the **Calendar** tab. You can also right-click a job to access a shortcut menu that contains additional tasks.

To manage jobs from the Calendar tab

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Calendar** tab, in the task pane, on the calendar, click the day for which you want to view details.

- 3 Click the day icon.
- 4 Do one of the following:
 - To view available tasks in the task pane Select the name of the job.
 - To view available tasks on the right-click menu Right-click the job.
- 5 Click the task that you want to perform.
See “[Job Monitor options](#)” on page 557.

Viewing the Symantec Endpoint Protection Security Summary

The Security Summary provides details from the Symantec Endpoint Protection application about viruses, threats, and risks to the media server.

See “[About using Backup Exec with Symantec Endpoint Protection](#)” on page 408.

Note: To enable the Security Summary, you must install the Symantec Endpoint Protection Manager component on the media server.

See the *Administrator’s Guide for Symantec Endpoint Protection* for more information about the Security Summary.

You can configure a backup job to start automatically when the Symantec ThreatCon level reaches a specified level.

To view the Security Summary

- 1 On the navigation bar, click **Job Monitor**.
- 2 On the **Security Summary** tab, view Symantec Endpoint Protection summary information.

About error-handling rules

You can enable default rules or create custom rules to set retry options and final job disposition for failed or canceled jobs. Retry options let you specify how often to retry a job if it fails and the time to wait between retry attempts. The final job disposition lets you either place the job on hold until you can fix the error, or reschedule the job for its next scheduled service.

To apply an error-handling rule for a group of similar errors, or error categories, you can enable a default error-handling rule. Each default error-handling rule applies to one category of errors, such as Network Errors or Security Errors. Default error-handling rules are disabled by default, so you must edit a rule and enable it before the retry options and job disposition settings will apply to jobs that fail with errors in the selected category. You cannot delete default error-handling rules, or add specific error codes to a category, or add new error categories. Before the error-handling rules will apply, the final error code must be in an error category that is associated with a rule, and the rule must be enabled.

To apply an error-handling rule for a specific error code that is in an error category, you can create a custom error-handling rule. You can select up to 28 error codes in an error category that a custom error-handling rule can apply to. You can also add an error code to an existing custom rule.

One custom error-handling rule, named Recovered Jobs, is created when Backup Exec is installed and is enabled by default. This rule applies retry options and a final job disposition to jobs that fail and that are not scheduled to run again.

See [“Custom error-handling rule for recovered jobs”](#) on page 594.

If both a custom error-handling rule and a default error-handling rule apply to a failed job, the settings in the custom rule are applied to the job.

Note: If the server on which Backup Exec is installed is in a cluster environment, the Cluster Failover error-handling rule is displayed on the list of error-handling rules. This rule is enabled by default.

See [“Cluster failover error-handling rule”](#) on page 595.

Creating a custom error-handling rule

You can create custom rules to set retry options and final job disposition for failed or canceled jobs.

See [“About error-handling rules”](#) on page 590.

To create a custom error-handling rule

- 1 On the **Tools** menu, click **Error-Handling Rules**.
- 2 Click **New**.
- 3 Complete the items in the **Error-Handling Rule Settings** dialog box, and then click **OK**.

See “[Error-Handling Rule Settings options](#)” on page 592.

See “[Custom error-handling rule for recovered jobs](#)” on page 594.

See “[Cluster failover error-handling rule](#)” on page 595.

Error-Handling Rule Settings options

You can create custom rules to set retry options and final job disposition for failed or canceled jobs. You can also edit existing rules.

See “[Creating a custom error-handling rule](#)” on page 591.

Table 13-14 Error-Handling Rule Settings options

Item	Description
Name	Indicates the name for the error-handling rule. To add or update a custom error-handling rule, you must enter a rule name.
Final Job Status	Indicates the status for the job that will activate the rule. The job status can be viewed, but not modified. The following statuses are available: <ul style="list-style-type: none">■ Error■ Canceled■ Failed

Table 13-14 **Error-Handling Rule Settings options** (*continued*)

Item	Description
Error Category	<p>Indicates the category of error for which the rule will be applied.</p> <p>If you are editing a default or custom error-handling rule, the error category can be viewed, but not modified.</p> <p>If you are creating a custom error-handling rule, you must select an error category that contains the errors to apply this rule to.</p> <p>Available error categories include the following:</p> <ul style="list-style-type: none"> ■ Other ■ Network ■ Server ■ Resource ■ Security ■ Backup Device ■ Backup Media ■ Job ■ System ■ Dispatch
Enabled	<p>Enables or disables the error-handling rule. This check box must be selected before you can set the retry options and the final job disposition options.</p>
Available errors	<p>Lists the error codes that are not associated with a custom error-handling rule. This field will not be displayed if you are editing a default error-handling rule.</p> <p>If you are creating or editing a custom error-handling rule, you must select the check box of the error code that you want this rule to apply to. You can select up to 28 error codes.</p> <p>To change the list of available errors, select a different error category.</p>
Retry job	<p>Allows Backup Exec to retry the job.</p>
Maximum retries	<p>Indicates the number of times you want the job retried. The maximum number of times the job can be retried is 99.</p>
Retry interval	<p>Indicates the number of minutes to wait before the job is retried. The maximum number of minutes is 1440.</p>

Table 13-14 Error-Handling Rule Settings options (*continued*)

Item	Description
Place job on hold until error condition has been manually cleared	Places the job on hold until you can manually clear the error. After you clear the error, you must remove the hold for the job.
Reschedule for its next scheduled service	Runs the job at the next scheduled occurrence.
Notes	Shows any miscellaneous information for the error-handling rule.

Custom error-handling rule for recovered jobs

Recovered Jobs is a custom error-handling rule that is used by Backup Exec to recover jobs that were failed with specific errors. This rule is created when Backup Exec is installed, and is enabled by default.

The retry options for this rule are to retry the job twice, with an interval of five minutes between the retry attempts. The final job disposition is to place the job on hold until you have manually cleared the error condition.

The following table describes the error codes that are selected by default for the Recovered Jobs custom error-handling rule.

Table 13-15 Error codes for recovered jobs custom error-handling rule

Error code	Description
0xE00081D9 E_JOB_ENGINE_DEAD	The displayed error message is: The Backup Exec job engine system service is not responding. See “Setting thresholds to recover jobs” on page 596.
0xE0008820 E_JOB_LOCAL RECOVERNORMAL	The displayed error message is: The local job has been recovered. No user action is required.
0xE000881F E_JOB_REMOTE RECOVERNORMAL	The displayed error message is: The remote job has been recovered. No user action is required.

Table 13-15 Error codes for recovered jobs custom error-handling rule
(continued)

Error code	Description
0xE0008821 E_JOB_STARTUP RECOVERY	The displayed error message is: Job was recovered as a result of Backup Exec RPC service starting. No user action is required.

Note: If the Central Admin Server Option is installed, additional error codes are selected.

See [“About error-handling rules”](#) on page 590.

See [“Cluster failover error-handling rule”](#) on page 595.

Cluster failover error-handling rule

If the server on which Backup Exec is installed is in a cluster environment, the cluster failover error-handling rule is displayed on the list of error-handling rules. This rule is enabled by default.

You cannot configure any options for this rule. You can only enable or disable the cluster failover error-handling rule.

The cluster failover error-handling rule and the **Apply CheckPoint Restart** option in **Cluster Backup Job Properties** work together to enable you to resume jobs from the point of failover. The **Apply CheckPoint Restart** option is dependent on the cluster failover error-handling rule; if you disable the rule, the option will automatically be disabled to match the rule’s setting.

See [“Enabling or disabling checkpoint restart ”](#) on page 824.

How thresholds are used to stall, fail, and recover jobs

If the Backup Exec services become unresponsive or jobs no longer run, you can set the threshold at which Backup Exec changes the status of active jobs to stalled. You can also set the threshold at which Backup Exec fails the jobs that were stalled, and then recovers them.

See [“Setting thresholds to recover jobs”](#) on page 596.

By setting a fewer number of seconds before Backup Exec reaches the threshold for changing a job's status to stalled, you can receive an earlier notification that jobs have stalled. A shorter time between the stalled and recovered thresholds also allows Backup Exec to fail and then recover the stalled jobs earlier. However, setting the thresholds too low may force a job to be recovered when it is not necessary.

Backup Exec recovers the jobs by using the custom error-handling rule named Recovered Jobs. This custom error-handling rule is created and enabled when Backup Exec is installed, and specifies that stalled/failed/recovered jobs are retried two times, with an interval of five minutes between the retries.

See [“Custom error-handling rule for recovered jobs”](#) on page 594.

See [“About error-handling rules”](#) on page 590.

Jobs that are stalled and then failed and recovered by Backup Exec because of unresponsive Backup Exec services are displayed differently in Backup Exec than jobs that fail because of errors in normal daily activities. The stalled/failed/recovered jobs are not indicated in red text in the job history as other failed jobs are. Instead, these jobs are displayed in gray text with a job status of **Recovered**.

In the job history, the error category is listed as Job Errors. The job history indicates the type of internal communication error that occurred and that the job was recovered. Based on the type of error that occurred, there may or may not be a log file associated with the recovered job.

Setting thresholds to recover jobs

If the Backup Exec services become unresponsive or jobs no longer run, you can set the threshold at which Backup Exec changes the status of active jobs to stalled. You can also set the threshold at which Backup Exec fails the jobs that were stalled, and then recovers them.

See [“How thresholds are used to stall, fail, and recover jobs”](#) on page 595.

To set thresholds to recover jobs

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Settings**, click **Job Status and Recovery**.
- 3 Change the appropriate fields, and then click **OK**.

See [“Application settings for job status and recovery options”](#) on page 597.

Application settings for job status and recovery options

If the Backup Exec services become unresponsive or jobs no longer run, you can set the threshold at which Backup Exec changes the status of active jobs to **Stalled**. You can also set the threshold at which Backup Exec fails the jobs that were stalled, and then recovers them.

See “[Setting thresholds to recover jobs](#)” on page 596.

Table 13-16 Application settings for job status and recovery options

Item	Description
Stalled	Indicates the number of seconds before the statuses for active jobs are changed to Stalled when the Backup Exec job engine service is not responding.
Recovered	Indicates the number of seconds before jobs are failed and then recovered by Backup Exec. A custom error-handling rule named Recovered Jobs is applied to recovered jobs. If this rule is disabled, then any other error-handling rules that have been enabled will apply to the recovered jobs. If there are no error-handling rules that apply to the job, then the job fails.

Restoring data

This chapter includes the following topics:

- [About restoring data](#)
- [Restore jobs and the catalog](#)
- [Restoring data by using the Restore Wizard](#)
- [Preventing the Restore Wizard from launching from the Restore button](#)
- [Configuring the Restore Wizard to launch from the Restore button](#)
- [Restoring data by setting job properties](#)
- [About selecting data to restore](#)
- [About redirecting restore jobs](#)
- [About redirecting restore jobs to native Microsoft Virtual Hard Disk \(VHD\) files](#)
- [Using redirected restore for Active Directory, Active Directory Application Mode for Windows Server 2003/2008](#)
- [Setting defaults for restore jobs](#)
- [Canceling a restore job](#)

About restoring data

With Backup Exec, you can retrieve information from storage media, including media created with backup software other than Backup Exec, and restore it to any server or remote workstation.

In most cases, you will need to restore only one file, but there may be times when directories, groups of files, or an entire system will need to be restored.

Backup Exec offers the following methods for finding the files you need to restore:

Table 14-1 Methods for finding files to restore

Method	Description
Resource view	Lists backed up data by the resource from which it was backed up. This feature is useful for finding files that were located on a certain server or workstation.
Media view	Lists the data that is contained on a piece of media. This feature is useful for viewing the contents of a tape that was backed up from another media server.
Selections detail view	Enables you to specify file and date attributes for the data you want to restore.
Search catalogs	Enables you to find files or other items that you want to restore, or to make sure that you have backups of certain files. This feature also enables you to see all cataloged, backed up versions of a file, so you can restore earlier versions if needed.

You can select options that you want to use for most restore jobs. Backup Exec will use the default options unless you override them when setting up a specific restore job.

When creating your restore jobs, you can do the following:

- Restore data to the system from which it was originally backed up or redirect the restore to another system.
- Specify if the restore job should begin processing immediately or schedule it to run at a future time.
- Specify which local network is to be used for restoring data, ensuring that other connected critical networks are not affected by this Backup Exec job.

See [“About selecting data to restore”](#) on page 626.

Restore jobs and the catalog

While backing up data from a resource, Backup Exec creates a set of catalog files that reside on the media server and on the media. These catalog files contain information about the contents of all media and are used when selections are made for restore jobs.

Media backed up at other Backup Exec installations must be cataloged by the local media server before data can be viewed in the **Restore Job Properties** dialog box because the catalog for the media does not exist on the media server. The media must have a Catalog job performed on it before files can be selected to restore.

See [“Creating a new catalog”](#) on page 248.

See [“Changing application settings for the catalogs”](#) on page 601.

See [“Catalog levels”](#) on page 603.

Changing application settings for the catalogs

Backup Exec uses the catalogs based on the application settings.

See [“Restore jobs and the catalog”](#) on page 600.

See [“Catalog levels”](#) on page 603.

To change application settings for the catalogs

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Settings**, click **Catalog**.
- 3 Select the appropriate options.
See [“Application settings for the catalogs”](#) on page 601.
- 4 Click **OK**.

Application settings for the catalogs

Defaults are enabled for application settings for the catalogs when you install Backup Exec.

See [“Changing application settings for the catalogs”](#) on page 601.

Table 14-2 Application settings for the catalogs

Option	Description
<p>Request all media in the sequence for catalog operations</p>	<p>Catalogs media starting with the lowest known tape number in the tape family. For example, if you don't have tape one, the catalog job will begin with tape two. If you uncheck this option, the catalog job begins on the tape that you specify.</p> <p>If you uncheck Request all media in the sequence for catalog operations, then you cannot select the Use storage media-based catalogs check box.</p>
<p>Use storage media-based catalogs</p>	<p>Allows Backup Exec to read the catalog information from the media.</p> <p>Media-based catalogs allow quick cataloging of media that are not included in the disk-based catalog (for example, media that was written by another installation of Backup Exec). This feature enables media to be cataloged in minutes, rather than the hours required with traditional file-by-file cataloging methods.</p> <p>If you want to create a new catalog by having Backup Exec read each file block, clear this option. Clearing this option is only recommended if normal catalog methods are unsuccessful.</p> <p>If you uncheck Request all media in the sequence for catalog operations, then the Use storage media-based catalogs check box cannot be selected.</p> <p>You cannot use Granular Recovery Technology if you select this option.</p> <p>See “How to restore individual items by using Granular Recovery Technology” on page 321.</p>

Table 14-2 Application settings for the catalogs (*continued*)

Option	Description
Truncate catalogs after	<p>Retains only the header information and removes all file and directory details after the specified amount of time. This option reduces the size of the catalogs considerably. After the catalogs have been truncated, the files and directories cannot be restored unless the media is recataloged.</p> <p>The last access date is not reset when catalogs are truncated.</p> <p>You can perform a full restore of backup sets from truncated catalogs.</p> <p>This option does not apply to synthetic backup jobs or true image restore jobs.</p>
Current path	Designates the path where you want catalogs to be located. This path defaults to \Program Files\Symantec\Backup Exec\Catalogs.
Catalog drive	Designates a volume where you want the catalog files to be located. This is useful if you have limited disk space on your media server.
Catalog path	Designates a path on the volume for the catalog files. If the path you provide does not exist, you will be prompted to create the path.

Catalog levels

The amount of information that can be viewed through the catalog for media is determined by the media's catalog level. Backup Exec fully catalogs each backup; however, there may be instances where media does not appear as fully cataloged in the **Restore Job Properties** dialog box.

The following are possible catalog levels:

Table 14-3 Media catalog levels

Item	Description
Fully cataloged media	With fully-cataloged media, you can do the following: <ul style="list-style-type: none">■ View information on all the directories and files contained in each backup set.■ Search for files to restore.
Truncated cataloged media	Truncated cataloged media lists only backup set information. No files or file attributes can be viewed. This version of Backup Exec writes only full catalogs.
Uncataloged media	There is no catalog information for the media. You must catalog the media to view and select files to restore.

See [“Changing application settings for the catalogs”](#) on page 601.

See [“Restore jobs and the catalog”](#) on page 600.

Restoring data by using the Restore Wizard

The Restore Wizard guides you through the creation of a restore job. The Restore Wizard is helpful for users who are new to Backup Exec.

To restore data using the Restore Wizard

- 1 From the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job using Wizard**.
- 3 Follow the on-screen prompts.

Preventing the Restore Wizard from launching from the Restore button

By default, the Restore Wizard displays when you select **Restore** on the navigation bar. If you prefer to set up restore jobs manually, you can prevent the Restore Wizard from displaying. If you disable the Restore Wizard, you can re-enable it at any time.

To prevent the Restore Wizard from launching from the Restore button

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job using Wizard**.

- 3 Uncheck **Always launch the Restore Wizard from the Restore button**.
- 4 Click **Next**.

Configuring the Restore Wizard to launch from the Restore button

By default, the Restore Wizard displays when you select **Restore** on the navigation bar. If you disable the Restore Wizard, you can re-enable it at any time.

To configure the Restore Wizard to launch from the Restore button

- 1 On the **Tools** menu, click **Wizards > Restore Wizard**.
- 2 Check **Always launch the Restore Wizard from the Restore button**.
- 3 Click **Next**.

Restoring data by setting job properties

If you are familiar with Backup Exec, you can restore data by selecting the options you want to use in the restore job.

See [“About selecting data to restore”](#) on page 626.

See [“About redirecting restore jobs”](#) on page 634.

See [“Filtering jobs”](#) on page 582.

See [“About restoring file permissions”](#) on page 619.

To protect remote resources, you must install the Backup Exec Remote Agent for Windows Systems on the remote computer.

See [“About the Remote Agent for Windows Systems”](#) on page 1945.

Depending on your file system environment, byte counts for restored data may not match the byte count recorded when the data was backed up. This is normal and does not mean that files were excluded in the restore job.

See [“Troubleshooting restore issues”](#) on page 797.

If you are restoring System State, restart your system before you restore more data.

To restore data by setting job properties

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.

- 3 Select the data that you want to restore.
See [“Selections options for restore jobs”](#) on page 609.
- 4 In the **Properties** pane, under **Source**, click **Device**.
- 5 Select the device from which you want to restore data.
See [“Device options for restore jobs”](#) on page 611.
- 6 Complete the following as necessary:

To change or test the logon credentials for the selected resources

In the **Properties** pane, under **Source**, click **Resource Credentials**.

See [“Resource Credentials options”](#) on page 337.

To redirect this job to another system other than the one from which the data was backed up

In the **Properties** pane, under **Destination**, do any of the following:

- Click **File Redirection** to redirect file sets.
See [“File Redirection restore options”](#) on page 634.
- Select the name of an agent or option to redirect specific agent data.
See [“Microsoft SQL Redirection options”](#) on page 1306.
See [“Microsoft Exchange redirection options”](#) on page 1163.
See [“Microsoft SharePoint redirection options”](#) on page 1250.
See [“Oracle redirection options”](#) on page 1347.
See [“DB2 redirection options”](#) on page 970.
See [“Redirection options for Enterprise Vault”](#) on page 1037.
See [“Archive redirection options for Archiving Option components”](#) on page 1490.
See [“VMware Redirection options”](#) on page 1404.
See [“Microsoft Hyper-V Redirection options”](#) on page 1184.

See [“About redirecting restore jobs”](#) on page 634.

- | | |
|--|---|
| To set general options for the restore job | In the Properties pane, under Settings , click General .
See “General options for restore jobs” on page 612. |
| To set advanced options for the restore job | In the Properties pane, under Settings , click Advanced .
See “Advanced options for restore jobs” on page 613. |
| To specify a local network to be used for this restore job | In the Properties pane, click Network and Security , and then enter the network information.
See “Network and security restore options” on page 617. |
| To set commands to run before or after the job | In the Properties pane, under Settings , click Pre/Post Commands .
See “Running pre and post commands for restore jobs” on page 618. |

- To configure restore options for an agent or option
- In the **Properties** pane, under **Settings**, select the name of the agent.
- See [“SQL restore options”](#) on page 1289.
- See [“Microsoft Exchange restore options”](#) on page 1155.
- See [“Microsoft SharePoint restore options”](#) on page 1247.
- See [“Lotus Domino restore options”](#) on page 1082.
- See [“Oracle restore options”](#) on page 1344.
- See [“DB2 restore options”](#) on page 968.
- See [“Enterprise Vault restore options”](#) on page 1034.
- See [“Restore job properties for Archiving Option databases”](#) on page 1484.
- See [“Restore job options for Linux, UNIX, and Macintosh computers”](#) on page 1898.
- See [“NDMP restore options for NetApp/IBM/Fujitsu”](#) on page 1864.
- See [“VMware restore options”](#) on page 1401.
- See [“Microsoft Hyper-V restore options”](#) on page 1183.
- To configure Backup Exec to notify someone when the restore job completes
- In the **Properties** pane, under **Settings**, click **Notification**.
- See [“Notification options for jobs”](#) on page 682.
- 7** Do one of the following:
- To run the backup job now
- Click **Run Now**.
- To schedule the backup job for later
- In the **Properties** pane, under **Frequency**, click **Schedule**.
- See [“Schedule options”](#) on page 356.

Selections options for restore jobs

When the **Restore Job Properties** dialog box appears, **Selections** is chosen by default in the **Properties** pane. Through the **Selections** options, you choose the data you want to include in the restore job. You can also choose how the data will appear in this dialog box.

Options on this dialog box include:

Table 14-4 Selections options for restore jobs

Item	Description
Selection list	Designates the selection list or lists you want to use. You can also use the default Selection list, which creates a new selection list using this name.
Load selections from existing list	Loads a previously created selection list or merges existing selection lists.
Search Catalogs	Lets you search for files or other items that you want to restore.
Include/Exclude	Allows you to select files to include in or exclude from the restore job. See “Restore Include/Exclude Selections options” on page 610.
Include subdirectories	Selects the contents of all the subfolders when a directory is selected.
Show file details	Displays details, such as the media label, the last backup date, and the backup set count, about the files available for selection.
Preview pane	Displays the preview pane at the bottom of the dialog box.
Beginning backup date	Indicates the date of the earliest backup set that you want to appear in the selection list. By default, the selection list includes backup sets that were created in the last 30 days. This option appears only when the View by Media and View by Resource tabs are selected.

Table 14-4 Selections options for restore jobs (*continued*)

Item	Description
Ending backup date	Indicates the date of the latest backup set that you want to appear in the selection list. By default, the selection list includes backup sets that were created in the last 30 days. This option appears only when the View by Media and View by Resource tabs are selected.
View by Resource	Displays backed up data by the resource from which it was backed up. This feature is useful for finding files that were located on a certain server or workstation.
View by Media	Displays the data that is contained on a piece of media. This feature is useful for viewing the contents of a tape that was backed up from another media server.
View Selection Details	Displays details about the media selected on either the View by Resource tab or the View by Media tab. The details that display include the date and time when the media was created, the media label, and the backup set to which the media belongs.

Restore Include/Exclude Selections options

The following include and exclude options are available when you restore jobs:
See [“Restoring data by setting job properties”](#) on page 605.

Table 14-5 Restore Include/Exclude Selections options

Item	Description
Media	Indicates the media that contains the files you want to restore.
Backup set	Indicates the backup set for which you want to specify attributes.
Path	Indicates any available directory or subdirectory. Enter the full path to the subdirectory.

Table 14-5 Restore Include/Exclude Selections options (*continued*)

Item	Description
File	Specifies a filename to be included or excluded. The default for this field is *.* , which means every file name with every extension is selected. Wildcard characters are permitted. The asterisk (*) in a file name or extension is a wildcard character that represents all characters occupying any remaining position in the file name or extension. For example, to specify all files with the .exe extension, type *.exe The question mark symbol (?) wildcard for a single character is also supported, as well as a double asterisk (**) to represent any number of characters, irrespective of any backslashes.
Include subdirectories	Indicates that all subdirectories and their contents in the path you have entered are included in (or excluded from) the job. If you want to process only the directory listed in the Path field, leave this option cleared.
Include	Includes the files in the operation. This is the default option.
Exclude	Excludes the files from the job.
Files dated	Includes or excludes files that were created or modified during the specified time period.

Device options for restore jobs

The following device options are available:

Table 14-6 Device options for restore jobs

Item	Description
Device	Specifies the device that contains the media for the data you want to restore. If the media is in another device, Backup Exec ignores this option.
Maximum number of devices to use for resources that support multiple data streams	Specifies the maximum number of devices that this restore job can use. Only one device per stream can be used. This option is applicable only for restores of Oracle and DB2 data.

General options for restore jobs

General options for restore jobs, including the name of the job, can be set through the **Restore Job Properties** dialog box.

See [“Restoring data by setting job properties”](#) on page 605.

Options for this dialog box include the following:

Table 14-7 General settings options for restore job

Item	Description
Job name	Specifies a name that describes the data that you are restoring. This is the name that is used to identify this job in the job schedule.
Job priority	Displays the priority of the access to the devices for this job. See “About job priority” on page 195.
Restore over existing files	Overwrites files on the target resource that have the same name as files that are being restored. Use this option only when you are sure that you want to restore an older version of a file.
Skip if file exists	Prevents Backup Exec from overwriting files on the target disk with files that have the same names that are included in the restore job.
Overwrite the file on disk only if it is older	Prevents Backup Exec from restoring over files that exist on the disk if they are more recent than the files included in the restore job. This option is useful if you are rebuilding a system. For example, after installing the operating system on a crashed computer, you could restore a previous full backup of the system without worrying about overwriting later versions of operating system files.
Restore all information for files and directories	Restores all information, including security information, for files and directories.
Restore only security information for files and directories	Restores only security information for files and directories. No other information is restored. This option is valid only for NTFS volumes. If you select this option along with the Skip if file exists option, no information will be restored. You can use this option with the Restore over existing files option and the Overwrite the file on disk only if it is older option.

Table 14-7 General settings options for restore job (*continued*)

Item	Description
Restore all information except security for files and directories	Restores all information except security information for files and directories. This option is valid only for NTFS volumes.
Restore corrupt files	<p>Allows you to restore corrupt files. Select this option only if you do not want to have Backup Exec automatically exclude corrupt files from the restore process.</p> <p>This option is only recommended if a job has failed because a catalog query could not determine the corrupt files on the tape. Normally, when a restore job is run, Backup Exec queries the catalog to determine if any corrupt files are on the tape and excludes them from the restore job. If, during the query process, Backup Exec cannot determine if a file is corrupt, the Restore job will not continue and will be marked as Failed. If a corrupt file cannot be excluded automatically, you can manually exclude corrupt files in the Restore selections window and run the job with the Restore Corrupt File option enabled.</p>
Preserve tree	<p>Restores the data with its original directory structure intact. This option is enabled by default. If you clear this option, all data (including the data in subdirectories) is restored to the path you specify in the Redirection dialog box.</p> <p>Clearing the Preserve Tree option is useful when restoring several subdirectories or individual files from media, but it should not be cleared when restoring an entire drive.</p>

Advanced options for restore jobs

You can set the following advanced options for restore jobs:

See [“Restoring data by setting job properties”](#) on page 605.

Table 14-8 Advanced options for restore

Item	Description
Restore Removable Storage data	<p>Restores the Removable Storage data. The Removable Storage database is stored in the <i>Systemroot\System32\Ntmsdata</i> directory and is automatically backed up when the system directory is selected for backup.</p> <p>Removable Storage is a service used to manage removable media and storage devices; it enables applications to access and share the same media resources.</p>
Restore disk quota data	<p>Restores disk quota data. Disk quota data is automatically backed up when the root directory of a volume is selected for backup.</p> <p>Disk quotas track and control disk usage on a per user, per volume basis; the values can be restored to the limits that were set before the backup.</p>
Restore Terminal Services database	<p>Restores the Terminal Services database. The default location for the Terminal Services database, which contains licensing data for client licenses, is the <i>Systemroot\System32\LServer</i> directory and is automatically backed up when the system directory is selected for backup.</p> <p>Terminal Services allow client applications to be run on a server so that client computers can function as terminals rather than independent systems.</p>
Restore Windows Management Instrumentation repository	<p>Restores the Windows Management Instrumentation (WMI) repository. The WMI repository is stored in the <i>Systemroot\System32\wbem\Repository</i> directory and is automatically backed up when the system directory is selected for backup.</p> <p>The Windows Management Instrumentation repository provides support for monitoring and controlling system resources and provides a consistent view of your managed environment.</p>
Restore Cluster Quorum	<p>Restores the cluster configuration.</p> <p>See “About restoring data to a Microsoft cluster” on page 840.</p>

Table 14-8 Advanced options for restore (*continued*)

Item	Description
<p>Force the recovery of the cluster quorum even if other nodes are online and/or disk signatures do not match.</p>	<p>Restores the cluster configuration if you are not able to take the other nodes in the cluster offline or if the disk that the cluster quorum previously resided on has been changed. This option is only available for computers that run Windows Server 2000/2003/2008 and if Restore Cluster Quorum is also selected.</p> <p>If this option is selected, the cluster service for any nodes that are online is stopped. This option also enables the drive letter of the disk that the cluster quorum was on to remain the same, even if the configuration has changed and the disk signatures contained in the restore media do not match the disk signatures contained in the cluster quorum.</p> <p>Any changes made to the cluster quorum after the last backup will be lost.</p>
<p>Mark this server as the primary arbitrator for replication when restoring folders managed by the File Replication Service, or when restoring SYSVOL in System State.</p>	<p>Designates this server as the primary replicator for all members in the set when restoring FRS-managed folders or SYSVOL as part of System State.</p> <p>If all members of a replication set are to be restored, then stop replication, restore all the member servers, and then when restoring the last member server, select this option to designate the server as the primary replicator. If this option is not selected, replication may not function.</p> <p>Note: In this version of Backup Exec, all restores of SYSVOL and FRS-managed folders are non-authoritative. An authoritative restore can only be performed by redirecting the restore and then copying the files to the server. Refer to your Microsoft documentation for details on performing an authoritative restore.</p>
<p>Allow managed media server to use any network interface to access remote agents</p>	<p>Enables a job that is delegated or copied to a managed media server to use any network interface to access remote agents. This option is applicable only for the Central Admin Server Option. By default, jobs that are delegated or copied to a managed media server from the central administration server use the network and security settings that are set on the managed media server. If the network specified on the managed media server is unavailable, selecting this check box enables the managed media server to use an alternate network to run important jobs.</p>

Table 14-8 Advanced options for restore (*continued*)

Item	Description
Merge the existing hardware configuration and registry services with the data to be restored	Merges the existing hardware and registry services with the data you selected to be restored. This option should only be used for restoring System State.
Overwrite the existing hardware configuration and registry services with the data to be restored	Overwrites hardware configuration and registry services with the data you selected to be restored. This option should only be used for restoring System State and there have been no hardware changes.
Restore junction points, symbolic links, files and directories from backup media	<p>Restores the information for the junction points, symbolic links, and the files and directories to which they are linked. If you select this option, existing junction points or symbolic links are overwritten.</p> <p>A junction point or symbolic link must have been backed up with one of the following options selected:</p> <ul style="list-style-type: none"> ■ Back up files and directories by following junction points ■ Back up files and directories by following symbolic links <p>Otherwise, these files and directories are not restored unless the junction point was linked to a mounted drive that did not have an assigned drive letter.</p> <p>See “Advanced options for backup jobs” on page 348.</p>
Preserve existing junction points and symbolic links and restore files and directories from backup media	<p>Restores files and directories that were backed up from junction points and symbolic links while retaining the destination computer's current junction points and symbolic links. This option prevents current junction points and symbolic links from being overwritten with the junction point and symbolic link information on the backup media.</p> <p>When you select this option and identical junction points, symbolic links, or directory names exist on both the destination computer and the media, the files and directories are restored to the destination computer's junction point, symbolic link, or directory.</p> <p>See “Advanced options for backup jobs” on page 348.</p>

Table 14-8 Advanced options for restore (*continued*)

Item	Description
<p>Path on an NTFS volume that is local to the media server for temporary storage of restore data</p>	<p>Creates a temporary staging area for restore data.</p> <p>This option is applicable only when you restore individual items in the following conditions:</p> <ul style="list-style-type: none"> ■ The backup of Microsoft Hyper-V, Microsoft Exchange, SharePoint, Active Directory, and VMware Virtual Infrastructure was enabled for Backup Exec Granular Recover Technology (GRT). ■ The backup is on a tape. ■ The backup is on a backup-to-disk folder that is not on an NTFS volume. <p>Type the path to a folder on an NTFS volume on this media server. Restore data and metadata for this job are stored here temporarily before the individual items are restored. The staged data is automatically deleted when the restore is complete.</p> <p>Symantec recommends that you avoid using system volumes for temporary staging locations.</p> <p>You can also specify a location that all applicable restore jobs can use.</p> <p>See “Setting defaults for restore jobs” on page 638.</p> <p>See “How to restore individual items by using Granular Recovery Technology” on page 321.</p>

Network and security restore options

You can override the default settings for a restore job by changing network and security options.

See [“Restoring data by setting job properties”](#) on page 605.

Table 14-9 Network and security restore options

Item	Description
Network interface	<p>Designates the name of the network interface card that connects the media server to the network you want to use for the restore network. The list includes all available network interface cards on the media server.</p> <p>If you are using the Central Admin Server Option (CASO), select the Use the default network interface for the managed media server option if you want CASO delegated restore jobs to be processed using the network interface card configured as the default in the managed media server.</p>
Protocol	<p>Designates the network protocol.</p> <p>You have the following options:</p> <ul style="list-style-type: none"> ■ Use any available protocol ■ Use IPv4 ■ Use IPv6
Subnet	<p>Displays the 32-bit number that determines the subnet to which the network interface card belongs.</p>
Allow use of any available network interface, subnet, or protocol for remote agents not bound to the above network interface, subnet, or protocol	<p>Ensures that the data from a remote system is backed up or restored over any available network if the remote system that you selected for backup or restore is not part of the specified restore network.</p> <p>If you do not select this check box and you selected data from a remote system that is not part of the specified restore network, the job fails because Backup Exec cannot back up or restore the data from the remote system.</p>
Interface details	<p>Designates the Media Access Control (MAC) address, Adapter type, Description, IP addresses, and subnet prefixes of the network interface you selected for the restore network.</p>

Running pre and post commands for restore jobs

You can run commands before or after a restore job, and set the following conditions for these commands:

- Run the job only if the pre-job command is successful
- Run the post-job command only if the pre-job command is successful
- Run the post-job command even if the job fails

- Allow Backup Exec to check the return codes (or exit codes) of the pre- and post-job commands to determine if the commands completed successfully. An exit code of zero returned to the operating system by the pre- or post-job command is interpreted by Backup Exec to mean the command completed successfully. A non-zero exit code is interpreted by Backup Exec to mean the command ended with an error.

See [“About pre/post commands”](#) on page 399.

See [“Setting default pre/post commands”](#) on page 400.

See [“Restoring data by setting job properties”](#) on page 605.

To set up commands to run before or after a restore job

1 In the **Properties** pane, under **Settings**, click **Pre/Post Commands**.

2 Complete the options as needed.

See [“Pre/post commands for backup or restore jobs”](#) on page 352.

About restoring file permissions

This section contains details on restoring data using the Restore Security option, which affects file security. This security feature applies only to NTFS partitions. To enable the Restore Security option, in the **Properties** pane, under **Settings**, you select **General**.

When restoring data with the Restore Security option, Backup Exec overwrites all directory security information presently on the disk with the security levels associated with the data being restored. This overwrite begins at the root of the restored directory structure and updates each directory in the tree until it reaches the data contained in the last directory.

For example:

With the following data on the storage media (a backup made prior to making security changes on disk):

\(root) Security applied: Users - Full

\Users Security applied: Users - Full

\User1 Security applied: User1 - Full

DATA.TXT Security applied: User1 - Full

With the following data on the disk (recently changed directory and file security):

\(root) Security applied: Users - Read

\Users Security applied: Users - Change

\User1 Security applied: User1 - Full

DATA.TXT Security applied: User1 - Full

After a restore with the Restore Security option selected, the security level of the data on the disk looks like this:

(root) Security applied: Users - Full

\Users Security applied: Users - Full

\User1 Security applied: User1 - Full

DATA.TXT Security applied: User1 - Full

If the data is restored without the Restore Security option selected, data.txt would inherit the permissions of the directory in which it was restored. In this case, it would inherit User1 directory's security level of Full.

See [“General options for restore jobs”](#) on page 612.

About System State

The system-specific data that comprises System State includes the registry, the COM+ Class Registration database, and boot and system files. The Certificate Services database will also be included if the server is operating as a certificate server. If the server is a domain controller, the data also includes Active Directory services database and SYSVOL directory. The System State data is backed up only as a collection. However, you can use the Active Directory Recovery Agent to restore individual objects.

If you are restoring Active Directory to a computer that is a domain controller, you must start the computer in safe mode and use the Directory Services Restore Mode to perform the restore. System State cannot be restored unless the target computer is in Directory Services Restore Mode. To restore System State data to a server that is not a domain controller, you can perform a basic restore.

If you have more than one domain controller in the network and you want Active Directory replicated to the other domain controllers, you must perform an authoritative restore of the Active Directory.

To perform an authoritative restore of the Active Directory, you must run Microsoft's Ntdsutil utility after the Backup Exec restore job completes and you have restored the System State data, but before you restart the server. An authoritative restore ensures that the restored data is replicated to all of the servers. For more information about authoritative restore and the Ntdsutil utility, see your Microsoft documentation.

Note: A System State backup is always a full backup. Therefore, when restoring, only the most recent backup of the System State must be restored. You should not cancel a System State restore job. Canceling this job could leave the system unusable.

See [“Restoring System State”](#) on page 621.

Restoring System State

The system-specific data that comprises System State includes the registry, the COM+ Class Registration database, and boot and system files. The Certificate Services database will also be included if the server is operating as a certificate server. If the server is a domain controller, the data also includes Active Directory services database and SYSVOL directory. The System State data is backed up only as a collection. However, you can use the Active Directory Recovery Agent to restore individual objects.

See [“About System State”](#) on page 620.

See [“Restoring data by setting job properties”](#) on page 605.

See [Table 14-8](#) on page 614.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 888.

To start the Backup Exec services and perform a local restore of System State data on a domain controller

- 1 Start the target server, press <F8> when prompted for Startup Options, and then select the **Directory Services Restore Mode** option.
- 2 Do one of the following:

To open Services on Windows 2000/2003 Do the following in the order listed:

- Right-click **My Computer**.
- Click **Manage**.
- Expand **Services and Applications**.

To open Services on Windows 2008 Do the following in the order listed:

- Right-click **My Computer**.
- Click **Manage**.
- Expand **Configuration**.

- 3 Click **Services**.

- 4 For each Backup Exec service listed, do the following in the order listed:
 - Click **Properties** on the shortcut menu.
 - Click the **Log On** tab, click **This account**, enter a user account with local administrator's rights, and then click **OK**.
 - Right-click the service, and then click **Start**.
- 5 After the Backup Exec services have started, run Backup Exec and perform a restore of the System State. Set the following option on the Advanced screen: **Mark this server as the primary arbitrator for replication when restoring folders managed by the File Replication Service, or when restoring SYSVOL in System State.**
- 6 If you are restoring System State, restart your system before you restore more data.

About restoring Shadow Copy Components

The Backup Exec Shadow Copy Components file system uses Microsoft's Volume Shadow Copy Service to protect critical operating system and application service data, and third-party application and user data on Windows resources.

A Writer is specific code within an application that participates in the Volume Shadow Copy Service framework to provide point-in-time, recovery-consistent operating system and application data. Writers appear as Shadow Copy Components, which are listed as resources in backup and restore selections. When expanded, the Backup Exec Shadow Copy Components file system includes the following selections:

Table 14-10 Backup Exec Shadow Copy Components

Item	Description
System State Writers	Lets you select System State Writers to restore. See "Restoring System State" on page 621.
Service State Writers	Lets you select Service State Writers to restore. See "Restoring data by setting job properties" on page 605.
User Data Writers	Lets you restore user data and the Microsoft Hyper-V. See "Restoring data to the Hyper-V host" on page 1182.

The User Data Writer in Backup Exec is the Active Directory Application Mode Writer (ADAM Writer). When restoring data with the ADAM Writer, Backup Exec stops the service for the ADAM instance you want to restore before the restore

job starts. However, Backup Exec does not restart the ADAM service when the restore job completes because post-processing jobs, such as authoritative restores using Adamutil.exe, may be needed. You must restart the ADAM service. If Backup Exec cannot stop the ADAM service or if Backup Exec cannot restore all of the ADAM files, the ADAM restore fails.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 888.

See [“Using redirected restore for Active Directory, Active Directory Application Mode for Windows Server 2003/2008”](#) on page 636.

About restoring utility partitions

Utility partitions, which are usually small partitions installed on the disk by OEM vendors like Dell, Hewlett-Packard, and IBM, can be selected for restore. These utility partitions contain system diagnostic and configuration utilities and are usually restored during disaster recovery.

However, utility partitions can be selected during a normal restore job provided the following requirements are met:

- Utility partitions, but not the data belonging to the partitions, must be present on the system.
- You must have Administrator rights to restore utility partitions.
- The system on which the utility partition data is being restored should be the same system from which the data was originally backed up, unless you are required to do a redirected restore.
See [“About performing redirected restores of utility partitions”](#) on page 623.
- Utility partitions being restored must belong to the same vendor. For example, Dell utility partitions cannot be restored to a Compaq system.
- The size of the utility partition on which the data is being restored must be equal or greater in size than the utility partition that was backed up.

See [“Restoring data by using the Restore Wizard”](#) on page 604.

See [“About selecting data to restore”](#) on page 626.

About performing redirected restores of utility partitions

You may need to perform a redirected restore of a utility partition if, during a disaster recovery, the system being recovered has been renamed. A redirected restore could also be required if a new system is replacing a crashed system. In the latter case, the system being restored must be the same model as the system originally backed up.

When doing a redirected restore of utility partitions, the following conditions must be met:

- Utility partitions, but not the data belonging to the partitions, must be present on the system.
- You must have Administrator rights to restore utility partitions.
- Utility partitions being restored must belong to the same vendor. For example, Dell utility partitions cannot be restored to a Compaq system.
- The size of the utility partition on which the data is being restored must be equal or greater in size than the utility partition that was backed up.
- The system on which the redirected restore is targeted must be the same make and model and have the same size utility partitions as the system from which the utility partition was backed up.

See “[Restoring data by setting job properties](#)” on page 605.

See “[File Redirection restore options](#) ” on page 634.

See “[About manual disaster recovery of Windows computers](#)” on page 780.

About restoring media created with other backup software

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Backup Exec supports restoring NetWare SMS volume backups to non-SMS volumes. For example, the data that is backed up with Backup Exec for NetWare Servers or Novell’s SBackup can be restored to the Windows media server or another network share.

About restoring data from ARCserve media

You can restore data from ARCserve media.

See “[Restoring data from ARCserve media](#)” on page 625.

The following types of data cannot be restored from ARCserve tapes:

- Databases, such as Microsoft SQL and Exchange Server and NetWare Directory Services
- Windows registry
- Interleaved files
- Compressed files

- Encrypted files
- Long filenames and Extended Attributes for OS/2 files
- Long filenames and resource forks for Macintosh files

Media containing ARCserve backups can be overwritten; however, backup append jobs are not supported. All Backup Exec media utility functions can be performed on ARCserve media.

See [“About inventorying media”](#) on page 448.

See [“Creating a new catalog”](#) on page 248.

See [“Restoring data by setting job properties”](#) on page 605.

See [“About restoring media created with other backup software”](#) on page 624.

Note: If the ARCserve backup spans multiple tapes, you must have all the tapes that were included in the ARCserve backup available. Make sure you start both the catalog and restore jobs with the first tape used in the ARCserve backup.

Restoring data from ARCserve media

You can restore data from ARCserve media.

See [“About restoring data from ARCserve media”](#) on page 624.

Table 14-11 Restoring data from ARCserve media

Step	Action
1	Inventory all the tapes included in the ARCserve backup. See “Inventorying media in a device” on page 448.

Table 14-11 Restoring data from ARCserve media (*continued*)

Step	Action
2	<p>Catalog all the tapes included in the ARCserve backup.</p> <p>See “Creating a new catalog” on page 248.</p> <p>During cataloging, Backup Exec reports file formats that it can read. Files that cannot be read do not appear in the catalogs. The media description that appears in the Backup Exec catalog comes from the session description used by ARCserve.</p> <p>Media-based catalogs are not supported on tapes created by other vendors’ backup products. Because of this, cataloging ARCserve tapes takes considerably longer than cataloging a tape made with Backup Exec.</p>
3	<p>Restore selected data to a server or workstation.</p> <p>See “Restoring data by setting job properties” on page 605.</p> <p>Due to the naming conventions ARCserve uses for some systems, it may be necessary to select a different location for the data using Backup Exec’s File Redirection.</p>

About selecting data to restore

When you are setting up a restore job, the first thing you do is select the data you want to restore. You can select data from the **View by Resource** tab or the **View by Media** tab.

You can find a list of icons that appear in the backup selections pane at the following URL:

<http://entsupport.symantec.com/umi/V-269-12>

On the **View by Resource** tab, restore selections are listed by the resource from which they were backed up.

Figure 14-1 View by Resource



The **View by Media** tab displays nodes that represent the media that contain backup sets. Each node displays the media label for the media on which the backup set is contained.

Figure 14-2 View by Media



If a backup set spans multiple pieces of media, the node for that backup set displays the media labels for all of those pieces of media. Beneath that node, the backup sets displays.

Note: True image restore selections do not appear on the **View by Media** tab. You can view true image restore selections on the **View by Resource** tab.

To expand the view for a resource or piece of media, click the adjacent box that contains the plus sign (+). To collapse the view, click the minus sign (-).

When the view is expanded, backup sets contained on the resource or media are displayed. You can expand the backup set to view the data included in the backup. The data that has been backed up from the resource appears in the right pane of the **Restore Job Properties - Selections** dialog box. Remember that only media cataloged or backed up at this server are displayed in the views. If you want to restore data backed up at another installation of Backup Exec, you must catalog the media first.

You can traverse file levels from either side of the window by clicking folders and subfolders as they appear.

To select data, select the check box next to the drives, directory, or file you want to restore. If the **Include subdirectories** option is selected on the **Restore Job**

Properties dialog box, all files and directories at or below the selected directory level are included in the restore job.

The check box and check mark displayed vary depending on the item's status.

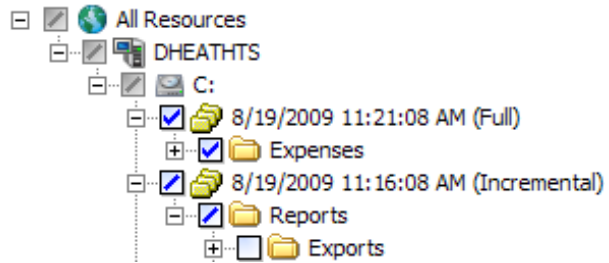
Figure 14-3 Data Selections

A slash in a shaded check box means that some items below the check box are selected, but the item itself cannot be selected.

A check mark in a check box means all items at or below the directory or drive level are selected.

A slash in a check box means some items below the directory or drive level are selected.

A clear check box means the item can be selected.



See [“Restore jobs and the catalog”](#) on page 600.

See [“Creating a new catalog”](#) on page 248.

See [“Restoring data by setting job properties”](#) on page 605.

Creating a restore selection list

A restore selection list includes all of the resources that you want to restore.

To create a restore selection list

- 1 On the navigation bar, click **Job Setup**.
- 2 In the Task pane, under **Selection List Tasks**, click **New restore selection list**.
- 3 Select the resources that you want to include in the selection list.
See [“New Restore Selection List options”](#) on page 628.
- 4 Select the appropriate options.
See [“New Restore Selection List options”](#) on page 628.
- 5 (Optional) To change or test a logon account for the resources, in the **Properties** pane, under **Source**, click **Resource Credentials**.
See [“Edit Logon Credentials options”](#) on page 190.
- 6 Click **OK**.

New Restore Selection List options

You can create a restore selection list that includes all of the resources that you want to restore.

See “[Creating a restore selection list](#)” on page 628.

Table 14-12 New Restore Selection List options

Item	Description
Selection list	Specifies the name of this selection list. You can use the name that Backup Exec provides.
Load selections from existing list	Loads an existing selection list. You can use the Load Selections from Existing List option to merge multiple selection lists. See “ Merging selection lists ” on page 300.
Search Catalogs	Searches the catalog to find files or other items that you want to restore, or to make sure that you have backups of certain files. This feature also enables you to see all cataloged, backed up versions of a file, so you can restore earlier versions if you need to. You can also use this feature to make sure that you have multiple copies of a file. Then, you can remove the file by running a full backup job that uses the method to back up and delete the files.
Include/Exclude	Selects files that you want to include in or exclude from this selection list.
Include subdirectories	Selects the contents of all the subfolders when a directory is selected.
Show file details	Displays details about the files available for selecting.
Preview pane	Displays the preview pane at the bottom of the dialog box. Clear this check box to remove the preview pane.
Beginning backup date	Displays only when the View by Media and View by Resource tabs are selected. To enable date ranges, check the check box next to the date. To display only the catalogs for data that was backed up during a specific date range, enter the beginning date in this field and enter the ending date in the Ending backup date field.
Ending backup date	Displays only when the View by Media and View by Resource tabs are selected. To enable date ranges, check the check box next to the date. To display only the catalogs for data that was backed up during a specific date range, enter the ending date in this field and enter the beginning date in the Beginning backup date field.
View by Resource	Displays backed up data by the resource from which it was backed up. This feature is useful for finding files that were located on a certain server or workstation.

Table 14-12 New Restore Selection List options (*continued*)

Item	Description
View by Media	Displays the data that is contained on a piece of media. This feature is useful for viewing the contents of a tape that was backed up from another media server.
View Selection Details	Displays details about the media selected on either the View by Resource tab or the View by Media tab. The details that display include the date and time when the media was created, the media label, and the backup set to which the media belongs.

Changing and testing resource credentials for restore jobs

If the logon account needed to restore data is different from the default logon account, you can change the account through the **Resource Credentials** dialog box. You can also use this dialog box to overwrite logon accounts for redirected restores. You can also verify that a logon account can access a resource.

You can change or test the default resource credentials when you create a new restore job.

See [“Restoring data by setting job properties”](#) on page 605.

To change and test resource credentials for restore jobs

- 1 On the navigation bar, click **Job Setup**.
- 2 Do one of the following:

To work with a job that is associated with a policy	In the Backup Selection Lists pane, click the job with which you want to work.
To work with a job that is not associated with a policy	In the Jobs pane, click the job with which you want to work.
- 3 In the Task pane, under **General Tasks**, click **Properties**.
- 4 In the **Properties** pane, under **Source**, click **Resource Credentials**.
- 5 Select the resource whose logon account you want to edit.
- 6 Click **Change**.
- 7 Select the logon account you want to use for this selection, or click **New** and create a new logon account.

See [“Logon Account Selection options”](#) on page 631.

- 8 To verify that the logon account you are using can access the resources selected for restore, click **Test All**.

While Backup Exec attempts to connect to the resources, "Testing" displays in the **Test Results** column. After a test has completed, the **Test Results** column will display one of the following results: Successful, Not tested, or an error message. The Not Tested result indicates that either the logon accounts have not been tested or that the tests have been performed but the server that contains the selection could not be accessed

Some tests may take a long time to complete. To cancel a logon account test, click **Cancel Test**.

- 9 Click **OK**.

Logon Account Selection options

The **Logon Account Selection** dialog box may appear for the following reasons:

- The Backup Exec logon account does not have sufficient rights to access the selected resource.
- You selected the option to change a logon account for a backup job.

From this dialog box, you can do the following:

- Select one of the existing logon accounts that is listed.
- Create a new logon account.
- Edit an existing logon account.

Searching for files to restore

You can search the catalog to easily find files that you want to restore, or to make sure that you have backups of certain files. This feature also enables you to see all cataloged, backed up versions of a file, so you can restore earlier versions if you need to. You can also use this feature to make sure that you have multiple copies of a file. Then, you can remove the file by running a full backup job that uses the method to back up and delete the files.

To search for files to restore

- 1 On the **Edit** menu, click **Search Catalogs**.
- 2 Complete the appropriate options.
See "[Search Catalogs options](#)" on page 632.

3 Click **Find Now**.

Click **Stop** to halt the search, or **New Search** to search for another file.

The **Search Catalogs** results window appears. All of the backed up versions of the file appear in the **Search Catalogs** window. Double-click the file to view the file's properties.

To sort the listings by filename, size, type or date modified, click the appropriate column heading.

4 Check the version of the file you want to restore and click **Apply**.

5 Submit the job using the same procedures required for other restore jobs.

Backup Exec will prompt you to insert the correct media if it is not already located in a drive.

Search Catalogs options

The **Search Catalogs** dialog box contains two tabs. The **Name & Resource** tab lets you search for data using file and media information. The **Date Modified** tab lets you search for data using dates.

The **Name & Resource** tab contains the following options:

Table 14-13 Name & Resource options

Item	Description
File/item name	Indicates the name of the file or item that you want to find. If you do not want to limit the search to a particular file, leave this field blank to search all files. You can use wildcard characters. Use a question mark (?) to represent any single character. Use an asterisk (*) to represent any number of characters. For example, to include all files with the .exe extension, type *.exe.
Path	Indicates the directory in which to search. If you do not want to limit the search to a particular directory, leave this field blank to search all directories. To search NetWare catalogs, use a forward slash (/).
Resource	Indicates the server and share in which to search. If you do not want to limit the search to a particular resource, leave this field blank to search all resources.

Table 14-13 Name & Resource options (*continued*)

Item	Description
Media	Indicates the cataloged media to search. You can select All Cataloged Storage Media to search the entire catalog, or you can select individual media to narrow the search.
Find directories	Searches for directories listed in the Path or File/item name fields.
Include subdirectories	Searches all subdirectories below the directory listed in the Path field.

The **Date Modified** tab contains the following options:

Table 14-14 Date Modified options

Item	Description
All files/items	Searches for all files or items.
Find all files/items created or modified	Searches only for files or items that have been created or modified in a specified time period.
Between x/x/x and x/x/x	Specifies specific dates by month, day, and year for the search.
During the previous x month(s)	Restricts the search to the previous month or number of months specified.
During the previous x day(s)	Restricts the search to the previous day or number of days specified.

About restore jobs and media libraries

For restore jobs, Backup Exec accesses the source media (if it is contained in the magazine) regardless of its sequential placement in the magazine. For example, if the data specified for a restore job resides on two media in the magazine, the media do not have to be placed in adjacent slots for Backup Exec to restore the data. Backup Exec's ability to randomly access media in this manner minimizes the amount of administrator attention required at the media server.

If Backup Exec does not find the media required for the restore job in the robotic library or other accessible storage devices, an alert is issued requesting the media necessary to complete the job.

About redirecting restore jobs

Backup Exec defaults to restoring data to the resource from which the data originated. By using the **Restore Job Properties** dialog box, you can restore data to any protected server or share.

See [“File Redirection restore options”](#) on page 634.

To redirect database files that are protected by licensed Backup Exec agents such as SQL or Exchange, in the Task pane, under **Destination**, select the redirection option for the agent.

See [“Using redirected restore for Active Directory, Active Directory Application Mode for Windows Server 2003/2008”](#) on page 636.

See [“About redirecting restore jobs to native Microsoft Virtual Hard Disk \(VHD\) files”](#) on page 636.

See [“Redirecting restores for SQL”](#) on page 1305.

See [“Redirecting Exchange restore data”](#) on page 1162.

See [“Redirecting a restore job for SharePoint 2003”](#) on page 1241.

See [“Redirecting a restore job for SharePoint 2007”](#) on page 1227.

See [“Redirecting a restore of Oracle data”](#) on page 1347.

See [“Redirecting a restore of DB2 data”](#) on page 970.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

See [“Redirecting the restore of a VMware virtual machine”](#) on page 1404.

See [“Restoring a virtual machine to a different host”](#) on page 1184.

File Redirection restore options

Backup Exec defaults to restoring data to the resource from which the data originated. You can redirect a restore job to any protected server or share.

See [“Restoring data by setting job properties”](#) on page 605.

Table 14-15 File Redirection restore options

Item	Description
Redirect file sets	Specifies target paths or resources other than where the data was originally backed up.
Restore to drive	Designates the destination for the restored data. Click the Browse button (...) to view local and network drives.

Table 14-15 File Redirection restore options (*continued*)

Item	Description
Server logon account	<p>Displays the current logon account being used by the server. If you need to use another logon account, click Change, and then select or create another account.</p> <p>See “Creating a new Backup Exec System Logon Account” on page 193.</p> <p>Click Clear to clear this field.</p>
Restore to path	<p>Specifies the target path on the device listed in the Restore to Drive field. If you want to retain the original directory structure, make sure that the Preserve Tree option is selected in the Restore Job Properties - Settings - General dialog box. If the Preserve tree option is not selected, all of the data will be restored to the path designated in this field.</p>
Path logon account	<p>Displays the logon account required for the target path. If you need to use another logon account, click Change, and then select or create another account. Click Clear to clear this field.</p>
Create a Microsoft Virtual Hard Disk for redirected files (Windows Server 2008 R2 or later)	<p>Creates one or more Microsoft Virtual Hard Disk files from the redirected data. This option is only available on computers with Microsoft Windows Server 2008 R2 or later.</p> <p>See “About redirecting restore jobs to native Microsoft Virtual Hard Disk (VHD) files” on page 636.</p>
Create a different Microsoft Virtual Hard Disk for each backup set that is restored	<p>Creates one Microsoft Virtual Hard Disk file for each backup set you want to restore.</p>
Create a single Microsoft Virtual Hard Disk that contains the merged files and folders from all redirected backup sets	<p>Creates one Microsoft Virtual Hard Disk file into which Backup Exec merges all of the files and folders that are contained in the backup sets.</p>
File name	<p>Designates the name of the Microsoft Virtual Hard Disk file you selected to create.</p>

About redirecting restore jobs to native Microsoft Virtual Hard Disk (VHD) files

You can redirect a restore job to a native Virtual Hard Disk (VHD) by selecting a supported computer as the restore job's destination. Supported computers include computers running Microsoft Windows 2008 R2, or Windows 7 clients with RAWFS installed. When you redirect a restore job to a native VHD, Backup Exec creates a VHD file that expands dynamically as you save data to it. The file can expand until it reaches 2040 GB, which is the maximum size for a native VHD file. You can create one VHD file that merges the data from all redirected backup sets. Or you can create a VHD file for each backup set.

See [“File Redirection restore options”](#) on page 634.

See [“About managing Microsoft Virtual Hard Disk \(VHD\) files in Backup Exec”](#) on page 293.

Using redirected restore for Active Directory, Active Directory Application Mode for Windows Server 2003/2008

When you want to install a new Windows Server Domain Controller into an existing domain, the Active Directory and SYSVOL data are replicated from the existing Domain Controller that is in the domain to the new Domain Controller. If there is a large amount of data to be replicated or if the connection between the Domain Controllers is slow or intermittent, the replication time can be lengthy. The Active Directory Application Mode replication time is also affected by the amount of data to be replicated and the connection speed. To decrease the replication time for Active Directory and Active Directory Application Mode, you can use the Install from Media feature.

For Active Directory, you can use the Install from Media feature to perform a System State backup of an existing Domain Controller in the domain in which you want to add a new Domain Controller. Then, you can perform a redirected restore of the data from the System State backup to the target Domain Controller.

For Active Directory Application Mode, you can back up data using the ADAM Writer. Then, you can perform a redirected restore of the data from the ADAM backup to the target system.

See [“About inventorying media”](#) on page 448.

See [“Creating a new catalog”](#) on page 248.

See “[Restoring data by setting job properties](#)” on page 605.

To install Active Directory using the Install from Media feature

- 1 Perform a standard System State backup of an active Windows Server Domain Controller that is in the target domain.
- 2 Transport the tape to the location of the system that will be installed into the target Domain.

The tape is not encrypted or protected. Symantec recommends that you encrypt the tape. Use caution when transporting it to the location of the target domain.

- 3 Inventory the drive where the tape is loaded.
- 4 Catalog the tape.
- 5 Perform a redirected restore of the System State backup to a temporary location on a volume or directory on the target system. In the **Properties** pane, under **Destination**, click **File Redirection**. And then select redirection options.

See “[File Redirection restore options](#)” on page 634.

When you redirect restored data, Backup Exec creates a sub-directory for each type of System State data being restored. Backup Exec creates the following sub-directories: Active Directory, SYSVOL, Registry, Boot Files, COM+ Class Registration Database, Certificate Server (if installed), and Cluster Quorum (if installed). Backup Exec also creates Automated System Recovery for Windows Server 2008.

- 6 To begin the Domain Controller installation, click **Start** on the target system, and then click **Run**.
- 7 Type `dcpromo /adv`
- 8 Click **OK**.
- 9 Click **Next** when the Active Directory Installation Wizard appears.
- 10 Select **Additional domain controller for an existing domain**.
- 11 Click **Next**.
- 12 Select **From these restored backup files**, and then enter the temporary location to which you redirected the System State data in step 5.
- 13 Click **Next**.
- 14 Complete the Active Directory Installation Wizard by following the prompts on the screen.
- 15 Complete the Domain Controller installation.

- 16 Reboot the system that has the new Domain Controller.
- 17 Delete any remaining temporary redirected System State files.
For more information, refer to your Microsoft documentation.

Setting defaults for restore jobs

The default options for all restore jobs are set through the **Options - Set Application Defaults** dialog box. Configure these items to match the settings that you want to use for most restore jobs. You can override these defaults while setting up a restore job, if necessary.

See [“Restoring data by setting job properties”](#) on page 605.

To set defaults for restore jobs

- 1 On the **Tools** menu, select **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Restore**.
- 3 Select the appropriate options.

See [“Default restore options ”](#) on page 638.

Default restore options

Configure the default restore options to match the settings that you want to use for most restore jobs.

See [“Setting defaults for restore jobs”](#) on page 638.

Table 14-16 Default restore options

Option	Description
Restore over existing files	Overwrites files on the target resource that have the same name as files that are being restored. Use this option only when you are sure that you want to restore an older version of a file.
Skip if file exists	Prevents Backup Exec from overwriting files on the target disk with files that have the same names that are included in the restore job.

Table 14-16 Default restore options (*continued*)

Option	Description
Overwrite the file on disk only if it is older	<p>Prevents Backup Exec from restoring over files that exist on the disk if they are more recent than the files included in the restore job.</p> <p>This option is useful if you are rebuilding a system. For example, after installing the operating system on a crashed computer, you could restore a previous full backup of the system without worrying about overwriting later versions of operating system files.</p>
Restore corrupt files	<p>Allows you to restore corrupt files. Select this option only if you do not want to have Backup Exec automatically exclude corrupt files from the restore process.</p> <p>Warning: Corrupt files, which appear in the restore selections window with a red X, could be incomplete files. Restoring corrupt files could result in corrupt data. Symantec recommends performing redirected restore of corrupt files rather than restoring to the original location.</p>
Restore junction points, symbolic links, files and directories from backup media	<p>Restores the information for the junction points and symbolic links and the files and directories to which they are linked. If you select this option, existing junction points are overwritten.</p> <p>If a junction point was originally backed up without the Backup files and directories by following junction points check box selected, then the files and directories to which the junction point is linked will not be restored, unless the junction point was linked to a mounted drive that did not have an assigned drive letter.</p> <p>See “Advanced options for backup jobs” on page 348.</p>

Table 14-16 Default restore options (*continued*)

Option	Description
<p>Preserve existing junction points and symbolic links and restore files and directories from backup media</p>	<p>Restores files and directories backed up from junction point links and symbolic links while retaining the system's current junction points. This option prevents current junction points from being overwritten with the junction point information restored from the backup media.</p> <p>When this option is selected and identical junction points or directory names exist on both the target system and the media, the files and directories are restored to the target system's junction point or directory.</p> <p>If a junction point or directory does not already exist in the same location and with the same name as the junction point to be restored, then the information for the junction point and the files and directories to which they point will be restored.</p> <p>If a junction point was originally backed up without the Back up files and directories by following junction points check box selected, then the files and directories to which the junction point is linked will not be restored, unless the junction point was linked to a mounted drive that did not have an assigned drive letter.</p> <p>See “Advanced options for backup jobs” on page 348.</p>

Table 14-16 Default restore options (*continued*)

Option	Description
Path on an NTFS volume that is local to the media server for temporary storage of restore data	<p>Creates a temporary staging area for restore data.</p> <p>This option is applicable only when you restore individual items in the following conditions:</p> <ul style="list-style-type: none"> ■ The backup of Microsoft Exchange, SharePoint, and Active Directory was enabled for Backup Exec Granular Recovery Technology (GRT). ■ The backup is on tape. ■ The backup is on a backup-to-disk folder that is not on an NTFS volume. <p>Note: This option also applies to archive jobs for the Exchange Mailbox Archiving Option.</p> <p>See “Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option” on page 1415.</p> <p>Type the path to a folder on an NTFS volume on this media server. Restore data and metadata are stored here temporarily before the individual items are restored.</p> <p>Symantec recommends that you avoid using system volumes for temporary staging locations.</p> <p>You can also specify a location for an individual job.</p> <p>See “Advanced options for restore jobs” on page 613.</p> <p>See “How to restore individual items by using Granular Recovery Technology” on page 321.</p>

Canceling a restore job

Canceling a restore job while it is in progress will result in unusable data, and may leave the drive in an unusable state. You may want to redirect the restore to a noncritical target, and then copy the data to a final destination when the job completes successfully.

You should not cancel a System State restore job. Canceling this job could leave the system unusable.

To cancel a restore job

- 1 On the navigation bar, click **Job Monitor**.
- 2 Click the **Job List** tab.

- 3** In the **Current Jobs Filter** box, click **Active jobs**.
- 4** Select the restore job you want to cancel.
- 5** In the task pane, under **Active Job Tasks**, click **Cancel**.
- 6** Click **Yes**.

Alerts and notifications

This chapter includes the following topics:

- [About alerts and notifications](#)
- [About alert views](#)
- [Viewing alerts](#)
- [Responding to active alerts](#)
- [Configuring alert category properties](#)
- [Enabling or disabling alerts from the Active Alerts pane](#)
- [Deleting alerts from the Alert History](#)
- [Setting up notification for alerts](#)
- [Configure Recipients options](#)
- [Assigning recipients to alert categories for notification](#)
- [Stopping alert notification for a recipient](#)
- [Sending a notification when a job completes](#)
- [Sending a notification when a selection list is used in a job](#)
- [About SNMP notification](#)
- [Installing and configuring the SNMP system service](#)
- [Installing the Windows Management Instrumentation performance counter provider](#)
- [Installing the Windows Management Instrumentation provider for SNMP](#)

- [Uninstalling the Windows Management Instrumentation performance counter provider](#)
- [Uninstalling the Windows Management Instrumentation provider for SNMP](#)

About alerts and notifications

An alert is any event in Backup Exec that is important enough to display a message or require a response from you.

Alert categories are conditions that cause alerts. Alert categories encompass many circumstances or problems that affect the system, jobs, media, or device sources. Each alert category can include one or more events that generate an alert. For example, a Job Failed error can be caused for many reasons.

Each alert category has one of the following alert types, which helps you distinguish the severity of the alert or whether Backup Exec needs a response from you.

Table 15-1 Alert types





Item	Description
Attention required 	Indicates issues that require a response before the job or operation can continue.
Error 	Indicates issues that affect job processing or the integrity of your backup.
Warning 	Indicates conditions that may or may not cause jobs to fail. You should monitor the conditions and take actions to resolve them.

Table 15-1 Alert types (*continued*)

Item	Description
<p>Information</p> 	<p>Provides status messages for the conditions that you might want to know about.</p>

Most alerts are enabled. However, you choose which alerts to display by editing alert category properties.

See [“Configuring alert category properties”](#) on page 658.

Alerts remain in the Active Alerts pane until they receive a response. You can respond to an alert manually or you can configure Backup Exec to respond to them automatically after a specified length of time. Depending on the alert type, a response might not be required, such as with informational alerts. After you respond to an alert, Backup Exec moves it to alert history, where it remains for the length of time you specify or until you delete it.

See [“Responding to active alerts”](#) on page 654.

You can configure Notifications to inform recipients when alerts occur. For example, you can notify a backup administrator via email or cell phone text message when a critical alert occurs.

See [“Setting up notification for alerts”](#) on page 661.

To assist with hardware troubleshooting, Backup Exec displays alerts for SCSI event ids 9 (device timeout), 11 (controller error), and 15 (device not ready).

About alert views

Backup Exec has two views for alerts: **Active Alerts** and **Alert History**.

The **Active Alerts** view displays the alerts that are active in the system. The **Alert History** view displays the alerts that have been responded to, or the alerts that have been automatically cleared from the system.

By default, Backup Exec displays all enabled alerts. However, when you select an alert view, you can choose filters to limit the type of alerts that appear in the pane.

See [“Filtering alerts”](#) on page 648.

You can double-click alerts in either view to see more detailed information.

See [“Viewing alert properties”](#) on page 651.

It may be necessary to view the job log to troubleshoot an alert. You can view the job log from an active alert or from a previous alert.

See [“Viewing the job log from an alert”](#) on page 653.

The status bar at the bottom of the screen displays an alert icon. The icon that displays in the status bar is for the most severe type of active alert, which may not be the most recent alert.

Active Alerts view and Alert History view

Backup Exec has two views for alerts: **Active Alerts** and **Alert History**.

See [“About alert views”](#) on page 645.

The **Active Alerts** view shows the following properties.

Table 15-2 Active Alerts view options

Item	Description
Type	Indicates the severity of the alert. The type helps you determine how quickly you want to respond. The following alert types may appear: <ul style="list-style-type: none">■ Errors■ Warnings■ Information■ Attention Required
Category	Indicates the condition that caused the alert. Categories include Database Maintenance, General Information, Device Error, or Job Failed.
Message	Indicates the text of the error message.
Time Alert Received	Shows the date and time when the alert was received.
Job Name	Indicates the name of the job that triggered the alert. This column is blank if the alert was not triggered by a job, such as for general information alerts.
Device Name	Shows the name of the device on which the alert occurred.
Server Name	Shows the name of the server on which the alert occurred.

Table 15-2 Active Alerts view options (*continued*)

Item	Description
Source	<p>Indicates the cause of the alert.</p> <p>Alerts can originate from one of the following sources:</p> <ul style="list-style-type: none"> ■ System ■ Job ■ Media ■ Device

The **Alert History** view shows the following properties.

Table 15-3 Alert History view options

Item	Description
Type	<p>Indicates the severity of the alert. The type helps you determine how quickly you want to respond.</p> <p>The following alert types may appear:</p> <ul style="list-style-type: none"> ■ Errors ■ Warnings ■ Information ■ Attention Required
Category	Indicates the condition that caused the alert. Categories include Database Maintenance, General Information, Device Error, or Job Failed.
Message	Indicates the text of the error message.
Time Alert Received	Shows the date and time when the alert was received.
Time User Responded	Shows the date and time when the user responded to the alert.
User Who Responded	Shows the user ID of the user who responded to the alert.
Response Machine	Shows the name of the computer on which the user responded to the alert.
Job Name	Indicates the name of the job that triggered the alert. This column is blank if the alert was not triggered by a job, such as for general information alerts.
Device Name	Shows the name of the device on which the alert occurred.
Server Name	Shows the name of the device on which the alert occurred.

Table 15-3 Alert History view options (*continued*)

Item	Description
Source	Indicates the cause of the alert. Alerts can originate from one of the following sources: <ul style="list-style-type: none">■ System■ Job■ Media■ Device

Viewing alerts

Active Alerts display the alerts that are active in the system. **Alert History** displays alerts that have been responded to, or the alerts that have been automatically cleared from the system.

To view alerts

- 1 On the navigation bar, click **Alerts**.
- 2 Select the **Active Alerts** tab or the **Alert History** tab.

Filtering alerts

You can filter the alerts that appear in the **Active Alerts** view or the **Alert History** view. Filters are useful when you have many alerts and you want to only view specific alert types. You can also filter the **Alert History** by alert types to expedite finding the alerts that were generated in the past.

See [“Creating custom filters for alerts”](#) on page 649.

To filter alerts

- 1 On the navigation bar, click **Alerts**.
- 2 Click the **Active Alerts** tab or the **Alert History** tab.
- 3 In the **Filter** box, select the type of alerts that you want to view.
- 4 If you have the Central Admin Server Option installed, you can select the **Media server alerts** filter. Then, select the media server for which you want to view alerts.

Select **All media servers** to view alerts for all media servers.

Creating custom filters for alerts

You can view various types of alerts and the sources that cause them by creating custom filters. For example, you can create a custom filter that displays only the Attention Required alerts and the Error alerts that are generated from devices and from media sources.

To create custom filters for alerts

- 1 On the navigation bar, click **Alerts**.
- 2 Select one of the following tabs:
 - **Active Alerts**
 - **Alert History**
- 3 In the task pane, under **Custom Filter Tasks**, click **Manage custom filters**.
- 4 Click **New**.
- 5 Type a unique name and a description for the filter.
- 6 On the **Properties** pane, under **Criteria**, click **Alert Type**.
- 7 Check **Enable this filter**.
- 8 Uncheck the check boxes for the alert types that you do not want to display.
- 9 On the **Properties** pane, under **Criteria**, click **Source**.
- 10 Check **Enable this filter**.
- 11 Uncheck the check boxes for the sources that you do not want to filter on.
- 12 Do one of the following:

If you have the Central Admin Server Option installed

Do the following:

- On the **Properties** pane, under **Criteria**, click **Media Server**.
- Select the media servers on which you want to filter.
If a media server is not listed, you can add it to the list.
- On the **Properties** pane, under **Criteria**, click **Media Server Pool**.
- Select the media server pools on which you want to filter.

If you do not have the Central Admin Server Option installed Go to step 13.

13 Click **OK**.

Editing custom filters for alerts

You can change custom filters at any time.

To edit a custom filter for alerts

- 1 On the navigation bar, click **Alerts**.
- 2 Select one of the following tabs:
 - **Active Alerts**
 - **Alert History**
- 3 In the task pane, under **Custom Filter Tasks**, click **Manage custom filters**.
- 4 Select the filter that you want to edit.
- 5 Click **Edit**.
- 6 Edit the custom filter options.
- 7 Click **OK**.
- 8 Click **Close**.

Deleting custom filters for alerts

You can delete custom filters when you no longer need them.

To delete a custom filter for alerts

- 1 On the navigation bar, click **Alerts**.
- 2 Select one of the following tabs:
 - **Active Alerts**
 - **Alert History**
- 3 In the task pane, under **Custom Filter Tasks**, click **Manage custom filters**.
- 4 Select the filter that you want to delete.
- 5 Click **Delete**.
- 6 When you are prompted to delete the custom filter, click **Yes**.
- 7 Click **Close**.

Viewing alert properties

Alert properties provide detailed information about each alert. In addition to the alert properties information, you can view category properties. If the alert is in the alert history, you can view response information.

To view alert properties

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts** or **Alert History**.
- 3 Select the alert from either the **Active Alerts** pane or the **Alert History** pane.
- 4 In the task pane, under **General Tasks**, click **Properties**.
See “[Alert properties](#)” on page 651.
- 5 After you view the alert properties, click **OK**.

Alert properties

Alert properties provide detailed information about each alert.

See “[Viewing alert properties](#)” on page 651.

The **Alert Properties** dialog box provides the following items:

Table 15-4 Alert properties

Item	Description
Category	Shows the title of the alert.
Type	Indicates the severity of the alert. The type helps you determine how quickly you want to respond. Following are the alert types: <ul style="list-style-type: none">■ Errors■ Warnings■ Information■ Attention Required
Server	Shows the name of the media server on which the alert occurred.
Device	Shows the name of the device on which the alert occurred.
Job Name	Shows the name of the job that is associated with the alert.
Time alert received	Shows the date and time the alert occurred.

Table 15-4 Alert properties (*continued*)

Item	Description
Source	<p>Indicates the cause of the alert.</p> <p>Alerts can originate from one of the following sources:</p> <ul style="list-style-type: none"> ■ System ■ Job ■ Media ■ Device
SNMP trap Identification	<p>Lists the SNMP message from Backup Exec regarding status and error conditions. SNMP must be installed to view this message.</p>
Enabled	<p>Indicates whether the alert is activated or disabled.</p>
Send notifications	<p>Indicates whether notifications are enabled or cleared for the alert. Recipients must be configured to use this option.</p>
Send SNMP notifications	<p>Indicates whether SNMP notifications are enabled or cleared for the alert. SNMP must be installed to use this option.</p>
Record in event log	<p>Indicates whether the alert is entered into the Windows Event Viewer. The Windows Event log displays all the property information for the alert.</p> <p>If a link appears in the Windows Event log you can search the Symantec Technical Support Web site for information about the Event ID.</p>
Event ID	<p>Shows the alert's ID in the Windows Event Viewer.</p>
Automatically clear after <i>x</i> days/hours/minutes	<p>Shows the length of time the alert remains active before it is moved to the Alert history.</p> <p>For the Attention Required alerts, you can set a default response. For more information, see the documentation for Backup Exec Utility.</p>
Respond with	<p>Lists the response Backup Exec automatically sends. This option is available only for the Media Overwrite and Media Insert alert categories.</p>
Include job log	<p>Sends the job log to the recipient that is configured for notification. This option can only be used for the recipients that are configured for email or printer notification.</p>
User who responded	<p>Shows the user ID that responded to the alert.</p>

Table 15-4 Alert properties (*continued*)

Item	Description
Response machine	Shows the name of the computer from which the user responded.
Time user responded	Shows the date and time when the user responded to the alert.
User response	Shows the response that the user entered for the alert.

Viewing the job log from an alert

The job log provides detailed job information, device and media information, job options, file statistics, and job completion status for completed jobs. You can view the job log for the jobs that have generated alerts from either **Active Alerts** or **Alert History**, depending on where the alert is located.

To view the job log from an alert

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts** or **Alert History**.
- 3 Select the alert for which you want to view the job log from either the **Active Alerts** or **Alert History** pane.
- 4 Perform one of the following:
 - If the alert is in **Active Alerts**, under **Alert Tasks** in the task pane, click **View job log**.
 - If the alert is in the **Alert History**, under **Alert History Tasks**, click **View job log**.
- 5 Do any of the following:
 - To search for a specific word or phrase, click **Find**. Type the text you want to find, and then click **Find Next**.
Be sure to expand all sections of the job log. The Find feature searches only the expanded sections of the job log.
 - To print the job log, click **Print**. To print the log, you must have a printer attached to your system and configured.
 - To save the job log as an .html or .txt file, click **Save As** and then select the file name, file location, and file type.
- 6 After you have finished viewing the job log, click **OK**.

Responding to active alerts

You can respond to active alerts and depending on the alert condition, continue or cancel the operation. By default, Backup Exec displays all enabled alerts, and all alerts that require a response. If you have set filters, only those alerts that are selected appear, in addition to any alerts that require a response. After the alert condition is resolved, the alert is moved to the **Alert History**.

If you click **Close** on the alert response dialog box, you will close the dialog box, but the alert remains active. To clear the alert and move it to the alert history, you must select a response such as **OK**, **Yes**, **No**, or **Cancel**.

You can configure automatic responses for alert categories.

Some alerts provide a Unique Message Identifier (UMI) code. This code is a hyperlink to the Symantec Technical Support Web site. You can access the technical notes that are related to the alert.

See [“Configuring automatic responses for alert categories”](#) on page 655.

To respond to an active alert

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Select the alert to which you want to respond, and then under **Alert Tasks** in the task pane, click **Respond**.

If you have more than one alert selected, click **Respond OK to all** to clear the selected alerts to the alert history. The alerts must have an OK response available to be automatically cleared.

- 4 Click a response for the alert.

See [“Alert response options”](#) on page 657.

About automatic responses for alert categories

Using the Alert Autoresponse Wizard, you can do the following:

- Configure the length of time to keep alerts active
- Specify the response that you want to send for alerts

You can use the Alert Autoresponse Wizard for the following alert categories:

Table 15-5 Alert categories for the **Alerts Autoresponse Wizard**

Alert category	Description
Library Insert	The Library Insert alert is a request to insert overwritable media into the robotic library by using the import command.
Media Insert	The Media Insert alert is a request to insert overwritable media into a tape drive. Most tape drives provide status to Backup Exec when media is inserted, which clears the alert. Other tape drives do not generate this status and the alert must be cleared with a response.
Media Overwrite	The Media Overwrite alert is displayed when Media Overwrite Protection is configured to prompt before overwriting media.
Media Remove	The Media remove alert is a request to acknowledge that media has been removed from a tape drive. Most tape drives provide status to Backup Exec when media is removed, which clears the alert. Other tape drives do not generate this status and the alert must be cleared with a response.

See [“Configuring automatic responses for alert categories”](#) on page 655.

Configuring automatic responses for alert categories

Automatic responses are useful if you routinely get alerts in the following categories and your responses to the alerts are always the same:

- Library Insert
- Media Insert
- Media Overwrite
- Media Remove

See [“About automatic responses for alert categories”](#) on page 654.

Settings that you configure by using this wizard are overridden if the alert category properties are reconfigured.

To configure automatic responses for alert categories

- 1 On the **Tools** menu, click **Wizards>Alert Autoresponse Wizard**.
- 2 On the **Welcome to the Alert Autoresponse Wizard** panel, click **Next**.

- 3 On the **Configure Library Insert Category** panel, select the appropriate options as follows:

Automatically clear after	Select the amount of time to display the alert before Backup Exec clears it and moves it to the alert history. When this job queues again, the alert reappears in the Active Alerts view and clears automatically after the time interval elapses.
Respond with	This option is not available for the Library Insert alert category.
Do not autorespond	Select this to keep this alert in the Active Alerts view until media is added to the library.

- 4 Click **Next**.

- 5 On the **Configure Media Insert Category** panel, select the appropriate options as follows:

Automatically clear after	Select the amount of time to display the alert before Backup Exec clears it.
Respond with	Do one of the following: <ul style="list-style-type: none">■ Select Yes to acknowledge that the media was inserted.■ Select No to retry the media insert operation on another tape drive (if multiple tape drives were selected for the job).■ Select Cancel to cancel this occurrence of the job.
Do not autorespond	Select this to keep this alert in the Active Alerts view until media is added to the tape drive.

- 6 Click **Next**.

- 7 On the **Configure Media Overwrite Category** panel, select the appropriate options as follows:

Automatically clear after	Select the amount of time to display the alert.
Respond with	Do one of the following: <ul style="list-style-type: none">■ Select Yes to overwrite the media automatically.■ Select No to try other media.■ Select Cancel to cancel the occurrence of this job.
Do not autorespond	Select this to keep this alert in the Active Alerts view until you acknowledge the alert by clicking OK .

- 8 Click **Next**.
- 9 On the **Configure Media Remove Category** panel, select the appropriate options as follows:

Automatically clear after	Select the amount of time to display the alert in the Active Alerts view before Backup Exec clears it and moves it to Alert History.
Respond with	This option is not available for the Media Remove alert category.
Do not autorespond	Select this to keep this alert in the Active Alerts view until you acknowledge the alert by clicking OK.

- 10 Click **Next**.
- 11 Read the **Alert Autoresponse Summary** panel, and then click **Next**.
- 12 On the **Completing the Alert Autoresponse Wizard** panel, click **Finish**.

Clearing informational alerts from the Active Alerts pane

Informational alerts can originate from the system, jobs, media, or devices. The alerts are set by default to move to the **Alert History** after 24 hours; however, some informational alerts appear frequently and fill the **Active Alerts** pane. You may want to clear these informational alerts to the **Alert History** pane before they are automatically moved by the system.

To clear informational alerts from the active alerts pane

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Alert Tasks** in the task pane, click **Clear all informational alerts**.

Alert response options

You can respond to active alerts and depending on the alert condition, you can either continue or cancel the operation.

See [“Responding to active alerts”](#) on page 654.

Table 15-6 Alert response options

Item	Description
Category Name	Shows the title of the alert.

Table 15-6 Alert response options (*continued*)

Item	Description
Message	Describes the event that caused the alert and provides suggestions for responding to the alert.
Click here for more information: V-XXX-XXXXX	<p>Appears if a TechNote is associated with an error. Click the Unique Message Identifier (UMI), which starts with the letter V, and appears as a blue hyperlink. A new browser window opens to the Symantec Technical Support Web site.</p> <p>If the computer does not have access to the Internet, you can type the following URL in a browser window on another computer: <a href="http://entsupport.symantec.com/umi/<UMI Code>">http://entsupport.symantec.com/umi/<UMI Code></p>
Server name	Shows the name of the computer on which the alert occurred.
Device name	Shows the name of the device on which the alert occurred.
Job name	Shows the name of the job that is associated with the alert.
Time	Shows the date and time the alert occurred.
Automatically display new alerts	<p>Enables alerts to appear automatically on the Backup Exec console when they are sent. If you do not select this option, you must respond to alerts through the Active Alerts pane.</p> <p>Note: Alerts that require a response always appear on the Backup Exec console.</p> <p>See “Changing application settings for Backup Exec preferences” on page 196.</p>
View job log	Lets you view the job log for the job that triggered the alert.
Automatically respond to and clear all alerts in the category	<p>Lets you provide automatic responses for this alert. You must select the amount of time to wait before a response is provided and select the response.</p> <p>See “Configuring automatic responses for alert categories” on page 655.</p>

Configuring alert category properties

You can set up alert categories to enable or disable alerts and to determine what actions should take place when an alert occurs.

Alternately, you can quickly enable and disable alerts from the **Active Alerts** pane without configuring other options.

See “[Enabling or disabling alerts from the Active Alerts pane](#)” on page 661.

Most alerts are enabled by default, however the following alert categories are initially disabled:

- Backup job contains no data
- Job Start
- Job Success

Each time you change the alert configuration, it is recorded in the audit log. You can view the audit log at any time to view the changes made to the alert category.

To configure alert category properties

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Alert Tasks** in the task pane, click **Configure alert categories**.
- 4 Under **Alert Categories**, click the alert for which you want to view or change properties.

You can change the options for more than one alert category simultaneously. To select consecutive alert categories, click the first alert, press and hold down <Shift>, and then click the last item. To select alert categories that are not consecutive, press and hold down <Ctrl>, and then click each item.

- 5 Under **Category Properties**, select the appropriate options.
See “[Configure Alert Categories options](#)” on page 659.
- 6 Click **Apply** to apply the properties to the alert and continue configuring additional alerts.
- 7 Click **OK** to exit the **Configure Alert Categories** dialog box.

Configure Alert Categories options

You can set up alert categories to enable or disable alerts and to determine what actions should take place when an alert occurs.

See “[Configuring alert category properties](#)” on page 658.

Table 15-7 Configure Alert Categories options

Item	Description
Alert categories	Lists the categories that are available.
Category name	Shows the title of the alert. This property can be viewed, but not edited.

Table 15-7 **Configure Alert Categories options** (*continued*)

Item	Description
Enable alerts for this category	<p>Activates or disables the alert. You cannot disable alert types such as error and attention required.</p> <p>You can also enable an alert category from the task pane.</p> <p>See “Enabling or disabling alerts from the Active Alerts pane” on page 661.</p>
Send notifications to selected recipients	<p>Sends a notification when an alert occurs. You must have recipients configured to use this option.</p> <p>To configure recipients to receive the notification, you must click Recipients.</p> <p>See “Configure Recipients options” on page 666.</p>
Include job log with a notification to an email or printer recipient	<p>Sends the job log to the recipient that is configured for notification. The recipient must be configured to receive email or printer notifications.</p>
Send SNMP Notifications	<p>Enables SNMP notification. SNMP must be installed to use this option.</p> <p>See “About SNMP notification” on page 683.</p>
Record event in the Windows Event Log	<p>Enters the alert into the Windows Event Viewer. The Windows Event log displays all the property information for the alert.</p> <p>If a link appears in the Windows Event log you can search the Symantec Technical Support Web site for information about the Event ID.</p>
Automatically clear after x days/hours/minutes	<p>Lets you enter the number of minutes, hours, or days you want the alert to remain active before it is moved to the Alert History.</p> <p>For Attention Required alerts, you can set up automatic responses.</p> <p>See “Configuring automatic responses for alert categories” on page 655.</p>
Respond with	<p>This option is available only for the Media Overwrite and Media Insert alert categories. Indicates the response that you want Backup Exec to send automatically. The choices are Cancel, No, Yes, or OK.</p>

Enabling or disabling alerts from the Active Alerts pane

You can quickly enable or disable alerts from the task pane instead of configuring them in the **Configure Alert Categories** dialog box. The error and attention required alert types cannot be disabled.

See [“Configuring alert category properties”](#) on page 658.

To enable or disable an alert from the Active Alerts pane

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 On the **Active Alerts** pane, select the alert that you want to enable or disable.
- 4 Under **Alert Tasks** in the task pane, click **Alert category enabled**.

Deleting alerts from the Alert History

Alerts that have been responded to or automatically cleared from the system are kept in the **Alert History**. All alerts are displayed, except for entries that have been filtered and selected for exclusion. The alerts remain in the **Alert History** for the length of time you set in the database maintenance option or until you delete the alert.

To delete an alert from the Alert History

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Alert History**.
- 3 Select the alert you want to delete, and then under **Alert History Tasks** in the task pane, click **Delete**.
- 4 Click **Yes** to confirm that you want to delete the alert.

Setting up notification for alerts

You can configure Backup Exec to notify recipients when alerts occur.

The following notification methods are available:

- SMTP email or phone text messaging
- MAPI email
- VIM email

■ **Pager**

You can use printer and Net Send notification methods, but they do not require configuration before creating and assigning recipients. You can use one or more methods for each recipient.

Table 15-8 How to set up notification for alerts

Step	Action
Step 1	Configure the method you want to use to notify the recipient. See “Configuring SMTP for email or mobile phone text message notification” on page 662. See “Configuring MAPI email for notification” on page 664. See “Configuring VIM email for notification” on page 665. See “Configuring a pager for alert notification” on page 666.
Step 2	Configure recipients. Recipients are individuals, computer consoles, printers, or groups. See “Configuring SMTP email or mobile phone text messaging for a person recipient” on page 667. See “Configuring MAPI mail for a person recipient” on page 668. See “Configuring VIM mail for a person recipient” on page 670. See “Configuring a pager for a person recipient” on page 671. See “Configuring a Net Send recipient” on page 674. See “Configuring a printer recipient” on page 675. See “Configuring a group recipient” on page 677.
Step 3	Assign the recipients to alerts or jobs for notification. See “Assigning recipients to alert categories for notification” on page 680.

Configuring SMTP for email or mobile phone text message notification

You must have an SMTP-compliant email system such as a POP3 mail server to receive alert notification messages using the SMTP notification method.

To configure SMTP for email or mobile phone text message notification

- 1** On the navigation bar, click **Alerts**.
- 2** Click **Active Alerts**.
- 3** Under **Notification Tasks** in the task pane, click **Configure email and pagers**.

- 4 On the **SMTP Configuration** tab, click **Enable**.
- 5 Select the appropriate options for this notification method.
See [“SMTP Configuration options”](#) on page 663.
- 6 Click **OK**.

SMTP Configuration options

You must have an SMTP-compliant email system such as a POP3 mail server to receive alert notification messages using the SMTP notification method.

See [“Configuring SMTP for email or mobile phone text message notification”](#) on page 662.

Table 15-9 SMTP Configuration options

Item	Description
Enable	Activates the notification method.
SMTP mail server	Shows the name of an SMTP mail server on which you have a valid user account. Backup Exec does not check the server name or the email address for validity.
SMTP port	Defaults to a standard SMTP port. In most cases, the default should not have to be changed.
Sender name	Indicates the sender's name. Spaces and special characters are allowed here.
Sender email address	Indicates the email address of the user from whom the notification message is sent. The email address should contain a name that identifies the user to the mail server, followed by an at sign (@) and the host name and domain name of the mail server. For example, john.smith@company.com. For a mobile phone: type the number of the mobile phone in email address format. For example: 1231231234@mymobile.com. Check with the mobile service provider for the correct email address for text messages.
Enable SMTP Authentication	Enables SMTP authentication, which logs the sender in to the mail server the SMTP notification is sent. For SMTP authentication to work properly, anonymous access and TLS encryption must be disabled on the Exchange server.
SMTP server login	Indicates the login name of the sender on the SMTP mail server.

Table 15-9 SMTP Configuration options (*continued*)

Item	Description
Sender password	Indicates the sender's password on the SMTP mail server. Make sure that you provide a confirmation password. Backup Exec does not check the server name or the email address for validity.

Configuring MAPI email for notification

You must have a MAPI-compliant email system such as Microsoft Exchange to receive alert notification messages using the MAPI notification method.

If you install Microsoft Outlook after installing Backup Exec, you must stop and restart the Backup Exec services for MAPI email notification to work and in order to save the MAPI configuration settings.

To configure MAPI email for notification

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure email and pagers**.
- 4 Click the **MAPI Configuration** tab and select the appropriate options.
 See [“MAPI Configuration options”](#) on page 664.
- 5 Click **OK**.

MAPI Configuration options

You must have a MAPI-compliant email system such as Microsoft Exchange to receive alert notification messages using the MAPI notification method.

See [“Configuring MAPI email for notification”](#) on page 664.

Table 15-10 MAPI Configuration options

Item	Description
Enable	Activates the notification method.
Mail server name	Indicates the name of the Exchange server. You must use an Exchange server to which the Backup Exec service account has access. See “About changing Windows security” on page 108.

Table 15-10 MAPI Configuration options (*continued*)

Item	Description
Mailbox name to send email from	<p>Indicates the mailbox from whom the notification message is sent, such as John Smith. The name appears in the From field in the message and does not require a full address.</p> <p>The Backup Exec services must be running under a domain account that has rights to the Exchange mailbox that is used for MAPI notification. Otherwise, the MAPI configuration settings are not saved.</p>

Configuring VIM email for notification

You must have a VIM (Lotus Notes) compliant email system to receive alert notification messages using the VIM notification method.

To configure VIM email for notification

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure email and pagers**.
- 4 Click the **VIM Configuration** tab and select the appropriate options.
See “[VIM Configuration options](#)” on page 665.
- 5 Click **OK**.

VIM Configuration options

You must have a VIM (Lotus Notes) compliant email system to receive alert notification messages using the VIM notification method.

See “[Configuring VIM email for notification](#)” on page 665.

Table 15-11 VIM Configuration options

Item	Description
Enable	Activates the notification method.
Notes client directory	Indicates the path of the directory in which the Notes client is located.
Mail password	Indicates the password that enables you to connect to the Notes client.
Confirm mail password	Indicates the password that enables you to connect to the Notes client. You must retype the password to confirm it.

Configuring a pager for alert notification

You can configure Backup Exec to page you with alert notification messages. A modem is required for pager notification. Make sure that the modem can communicate properly with your paging service. Before setting up pager notification, contact your paging service for information about modems compatible with the service.

To configure a pager for alert notification

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure email and pagers**.
- 4 Click the **Pager Configuration** tab and select **Enable** to activate this alert notification method.
- 5 Select a modem from the **Select a modem for sending pages** option.
 Only modems recognized in Windows appear in the list box.
- 6 Click **OK**.

Configure Recipients options

Recipients are individuals with a predefined notification method, computer consoles, printers, or groups. Recipient configuration consists of selecting a notification method and defining notification limits. After you create entries for the recipients, you can assign them to alerts, jobs, or selection lists.

The following types of recipients can be configured for notifications:

Table 15-12 **Configure Recipients options**

Item	Description
Person	Lets you set up an individual as a recipient for alerts. The individual must have a predefined method of notification such as SMTP, MAPI, or VIM email, or a pager. You must configure the notification method before you can enable it for the recipient.
Net Send	Lets you set up a computer as a notification recipient.
Printer	Lets you set up a specific printer to which notifications can be sent.
Group	Lets you set up a group of one or more recipients, including person recipients, Net Send recipients, and other groups.

See [“Configuring SMTP for email or mobile phone text message notification”](#) on page 662.

See [“Assigning recipients to alert categories for notification”](#) on page 680.

See [“Sending a notification when a job completes”](#) on page 682.

See [“About selection lists”](#) on page 295.

Configuring SMTP email or mobile phone text messaging for a person recipient

You can configure a person recipient to receive SMTP email or mobile phone text message notification messages if you have configured the SMTP notification method.

To configure SMTP email or mobile phone text messaging for a person recipient

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure recipients**.
- 4 Click **New**, click **Person** in the **Recipient Type** dialog box, and then click **OK**.
- 5 In the **Name** field, type the name of the recipient that you want to configure.
- 6 Click the **SMTP Mail** tab and select the appropriate options.
See [“SMTP Mail options”](#) on page 667.
- 7 Click **OK**.

SMTP Mail options

You can configure a person recipient to receive SMTP email or mobile phone text message notification messages if you have configured the SMTP notification method.

See [“Configuring SMTP email or mobile phone text messaging for a person recipient”](#) on page 667.

Table 15-13 SMTP Mail options

Item	Description
Enable	Activates this notification method for the recipient.

Table 15-13 SMTP Mail options (continued)

Item	Description
Address	<p>For email, indicates the email address of the recipient to whom the notification message is sent. For example, john.smith@company.com.</p> <p>For a mobile phone, indicates the number of the mobile phone in email address format. For example: 1231231234@mymobile.com. Check with the mobile service provider for the correct email address for text messages.</p>
Test	Enables you to test the notification configuration for the recipient.
Enable	Activates the option.
Notify me a maximum of x times within x minutes	Indicates the total number of notifications you want sent to the recipient for all alerts that are generated within the specified number of minutes. After the specified number of notifications have been sent, additional notifications are not sent until the specified minutes have been reached. The maximum number of minutes that can be set is 1440, which is the number of minutes in a day.
Reset the notification limits after x minutes	Indicates the number of minutes that must be reached before the notification limits are reset. When the time limit has been reached, the number of sent notifications is reset to zero.
Enable	Activates the option and lets you configure the length of time the recipient is available for notification.
Schedule	<p>Enables you to select the days and times when notifications can be sent to the recipient.</p> <p>See “Scheduling notification for recipients” on page 677.</p>

Configuring MAPI mail for a person recipient

You can configure a person recipient to receive MAPI email notification messages if you have configured the MAPI notification method.

To configure MAPI mail for a person recipient

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure recipients**.

- 4 Click **New**, click **Person** in the **Recipient Type** dialog box, and then click **OK**.
- 5 In the **Name** field, type the name of the recipient that you want to configure.
- 6 Click the **MAPI Mail** tab and select the appropriate options.
See [“MAPI mail options”](#) on page 669.
- 7 Click **OK**.

MAPI mail options

You can configure a person recipient to receive MAPI email notification messages if you have configured the MAPI notification method.

See [“Configuring MAPI mail for a person recipient”](#) on page 668.

Table 15-14 MAPI mail options

Item	Description
Enable	Activates this notification method for the recipient.
Mailbox	Indicates the email address or mailbox name of the recipient to whom the notification message is sent. For example, john.smith@company.com or John Smith.
Test	Enables you to test the notification configuration for the recipient.
Enable (Limit the number of notifications sent)	Activates the option that lets you specify the number of notifications to send to the recipient.
Notify me a maximum of x times within x minutes	Indicates the total number of notifications you want sent to the recipient for all alerts that are generated within the specified number of minutes. After the specified number of notifications have been sent, additional notifications are not sent until the specified minutes have been reached. The maximum number of minutes that can be set is 1440, which is the number of minutes in a day.
Reset the notification limits after x minutes	Enables you to enter the number of minutes that must be reached before the notification limits are reset. When the time limit has been reached, the number of sent notifications is reset to zero.
Enable (Limit when notifications can be sent)	Lets you configure the length of time the recipient is available for notification.

Table 15-14 MAPI mail options (*continued*)

Item	Description
Schedule	Enables you to select the days and times when notifications can be sent to the recipient. See “Scheduling notification for recipients” on page 677.

Configuring VIM mail for a person recipient

You can configure a person recipient to receive VIM email notification messages if you have configured the VIM notification method.

To configure VIM mail for a person recipient

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure recipients**.
- 4 Click **New**, click **Person** in the **Recipient Type** dialog box, and then click **OK**.
- 5 In the **Name** field, type the name of the recipient that you want to configure.
- 6 Click the **VIM Mail** tab and select the appropriate options.

See [“VIM Mail options”](#) on page 670.

VIM Mail options

You can configure a person recipient to receive VIM email notification messages if you have configured the VIM notification method.

See [“Configuring VIM mail for a person recipient”](#) on page 670.

Table 15-15 VIM Mail options

Item	Description
Enable	Activates this notification method for the recipient.
Address	Indicates the email address of the recipient to whom the notification message is sent. For example, JohnSmith@company.com.
Test	Enables you to test the notification configuration for the recipient.
Enable (Limit the number of notifications sent)	Activates the option.

Table 15-15 VIM Mail options (*continued*)

Item	Description
Notify me a maximum of x times within x minutes	Indicates the total number of notifications you want sent to the recipient for all alerts that are generated within the specified number of minutes. After the specified number of notifications have been sent, additional notifications are not sent until the specified minutes have been reached. The maximum number of minutes that can be set is 1440, which is the number of minutes in a day.
Reset the notification limits after x minutes	Indicates the number of minutes that must be reached before the notification limits are reset. When the time limit has been reached, the number of sent notifications is reset to zero.
Enable (Limit when notifications can be sent)	Activates the option and lets you configure the length of time the recipient is available for notification.
Schedule	Enables you to select the days and times when notifications can be sent to the recipient. See “Scheduling notification for recipients” on page 677.

Configuring a pager for a person recipient

You can configure a person recipient to receive notification messages by pager if you have configured the pager notification method.

To configure a pager for a person recipient

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure recipients**.
- 4 Click **New**, click **Person** in the **Recipient Type** dialog box, and then click **OK**.
- 5 In the **Name** field, type the name of the recipient that you want to configure.
- 6 Click the **Pager** tab and select the appropriate options.
See [“Pager options”](#) on page 672.
- 7 Click **Advanced Setup** to configure the advanced pager setup options in the **Advanced Pager Information** dialog box:
See [“Advanced Pager Information options”](#) on page 673.
- 8 Click **OK** to save the settings in the **Advanced Pager Information** dialog box, and then click **OK** to save the pager configuration settings.

Pager options

You can configure a person recipient to receive notification messages by pager if you have configured the pager notification method.

See [“Configuring a pager for a person recipient”](#) on page 671.

Table 15-16 Pager options

Item	Description
Enable	Activates this notification method for the recipient.
Carrier Phone	Indicates the area code and phone number to access the paging service provider’s modem. The paging service number may be different from the number you enter to manually enter a page.
Country/region name and code	Indicates the name of the country or the region, and the country code in which the pager is located.
Pager Pin	Indicates the pager identification number (PIN). The PIN is provided by the paging service provider. You have a PIN if you use TAP services. In most cases, the PIN is the last seven digits of the pager’s phone number.
Advanced	Enables you to configure additional settings for the pager. See “Advanced Pager Information options” on page 673.
Test	Enables you to test the notification configuration for the recipient.
Enable (Limit the number of notifications sent)	Activates the option.
Notify me a maximum of x times within x minutes	Indicates the total number of notifications you want sent to the recipient for all alerts that are generated within the specified number of minutes. After the specified number of notifications have been sent, additional notifications are not sent until the specified minutes have been reached. The maximum number of minutes that can be set is 1440, which is the number of minutes in a day.
Reset the notification limits after x minutes	Indicates the number of minutes that must be reached before the notification limits are reset. When the time limit has been reached, the number of sent notifications is reset to zero.
Enable (Limit when notifications can be sent)	Activates the option and lets you configure the length of time the recipient is available for notification.

Table 15-16 Pager options (*continued*)

Item	Description
Schedule	Enables you to select the days and times when notifications can be sent to the recipient. See “Scheduling notification for recipients” on page 677.

Advanced Pager Information options

You can configure a person recipient to receive notification messages by pager if you have configured the pager notification method.

See [“Configuring a pager for a person recipient”](#) on page 671.

Table 15-17 Advanced Pager Information options

Item	Description
Password	Indicates the password for the pager, if one is required.
Message Length	Indicates the maximum number of characters you want to use for messages. The paging service provider determines the maximum number.
Retrys	Indicates the number of times you want the paging service provider to retry the page. The paging service provider determines number.
Numeric	Indicates that the pager accepts only numbers.
Alpha-numeric	Indicates that the pager that accepts letters and numbers.
Modem Baud Rate	Indicates the speed of the modem. The speeds that appear are limits set by the paging service; select the appropriate speed regardless of the modem speed rating.
Data bits, Parity, Stop bit	Indicates the communication protocol. In most cases, you should use the Windows default.

Configuring a Net Send recipient

You can configure Net Send to send notification messages to a target computer or user.

If the target computer has Internet pop-up advertisement blocking software installed, the Net Send notification message will not display.

To configure a net send recipient

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure recipients**.
- 4 Click **New**, click **Net Send** in the **Recipient Type** dialog box, and then click **OK**.
- 5 Select the appropriate options from the **Net Send Recipient Properties** dialog box:

See “[Net Send Configuration Properties options](#)” on page 674.
- 6 Click **OK**.

Net Send Configuration Properties options

You can configure Net Send to send notification messages to a target computer or user.

See “[Configuring a Net Send recipient](#)” on page 674.

Table 15-18 Net Send Configuration Properties options

Item	Description
Name	Indicates the name of the recipient that you want to receive the notification.
Target Computer or User Name	Indicates the name of the computer or user to whom you want to send the notification. You should enter a computer rather than a user because the Net Send message fails if the user is logged off the network. If the target computer has Internet pop-up advertisement blocking software installed, the Net Send notification message does not display.
All Computers	Sends the notification to all the computers in the network.
Test	Enables you to test the notification configuration for the recipient.

Table 15-18 Net Send Configuration Properties options (continued)

Item	Description
Enable (Limit the number of notifications sent)	Activates the option.
Notify me a maximum of x times within x minutes	Indicates the total number of notifications you want sent to the recipient for all alerts that are generated within the specified number of minutes. After the specified number of notifications have been sent, additional notifications are not sent until the specified minutes have been reached. The maximum number of minutes that can be set is 1440, which is the number of minutes in a day.
Reset the notification limits after x minutes	Indicates the number of minutes that must be reached before the notification limits are reset. When the time limit has been reached, the number of sent notifications is reset to zero.
Enable (Limit when notifications can be sent)	Activates the option and lets you configure the length of time the recipient is available for notification.
Schedule	Enables you to select the days and times when notifications can be sent to the recipient. See “Scheduling notification for recipients” on page 677.

Configuring a printer recipient

You can select installed printers as a notification method for recipients; however, Backup Exec does not support fax printer devices. Only printers that were configured using the same username and password as the Backup Exec service account can be selected.

To configure a printer recipient

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure recipients**.
- 4 Click **New**, click **Printer** in the **Recipient Type** dialog box, and then click **OK**.

- 5 Select the appropriate options in the **Printer Recipient Properties** dialog box.
 See [“Printer Recipient Properties options”](#) on page 676.
- 6 Click **OK**.

Printer Recipient Properties options

You can select installed printers as a notification method for recipients; however, Backup Exec does not support fax printer devices. Only the printers that were configured using the same user name and password as the Backup Exec service account can be selected.

See [“Configuring a printer recipient”](#) on page 675.

Table 15-19 Printer Recipient Properties options

Item	Description
Name	Indicates the name of the recipient for whom you want to receive the notification. You cannot use a fax printer device to receive the notification.
Target Printer	Indicates the name of the printer to which you want to send the notification message.
Test	Enables you to test the notification configuration for the recipient.
Enable (Limit the number of notifications sent)	Activates the option.
Notify me a maximum of x times within x minutes	Indicates the total number of notifications you want sent to the recipient for all alerts that are generated within the specified number of minutes. After the specified number of notifications have been sent, additional notifications are not sent until the specified minutes have been reached. The maximum number of minutes that can be set is 1440, which is the number of minutes in a day.
Reset the notification limits after x minutes	Indicates the number of minutes that must be reached before the notification limits are reset. When the time limit has been reached, the number of sent notifications is reset to zero.
Enable (Limit when notifications can be sent)	Activates the option and lets you configure the length of time the recipient is available for notification.

Table 15-19 Printer Recipient Properties options (continued)

Item	Description
Schedule	Enables you to select the days and times when notifications can be sent to the recipient. See “Scheduling notification for recipients” on page 677.

Configuring a group recipient

Groups are configured by adding recipients as group members. A group contains one or more recipients and each recipient receives the notification message. Members of the group can be a combination of individual persons, computers, printers, or other groups. In addition, a group can be added to other groups.

To configure a group recipient

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure recipients**.
- 4 Click **New**, click **Group** in the **Recipient Type** dialog box, and then click **OK**.
- 5 In the **Group Name** field, type the group for whom you are configuring the notification.
- 6 To add members to the group, select recipients from the **All Recipients** list, and then click **Add** to move them to the **Group Members** list.
- 7 To remove members from the group, select recipients from the **Group Members** list, and then click **Remove** to move them to the **All Recipients** list.
- 8 When you have completed the group, click **OK**.

Scheduling notification for recipients

During the recipient configuration process, you can enable the **Limit when notifications can be sent** option to select the times of the day and the days of the week the recipient is available to receive the notification messages. You can modify the schedule after the recipient is configured by editing recipient notification properties.

To configure the notification schedule for recipients during recipient configuration

- 1 On the **Recipient Properties** dialog box, under the Limit when notifications can be sent group box, click **Enable** to activate the option.

To access the **Recipient Properties** dialog box, on the navigation bar, click **Alerts**. On the task pane, under **Notification Tasks**, click **Configure recipients**. Click **New** to create a new recipient or select an existing recipient and then click **Properties**.

- 2 Click **Schedule**.

- 3 Do any of the following:

- Clear the **Include work days** check box to exclude Monday through Friday from 8 A.M. to 6 P.M.
- Clear the **Include weeknights** check box to exclude Monday through Friday from 6 P.M. to 8 A.M.
- Clear the **Include weekends** check box to exclude Saturday and Sunday, 24 hours a day.

You can select any combination of **Include work days**, **Include weeknights**, or **Include weekends**, or click any single hour of the chart to select or clear that hour.

- 4 After selecting the days and times you want, click **OK**.

Editing recipient notification properties

You can edit the recipient notification properties at any time and change the recipient information, such as an email address, telephone number, or schedule.

To edit the recipient notification properties

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure recipients**.
- 4 Select the recipient you want to edit, and then click **Properties**.
- 5 Edit the properties for the selected recipient.

You can edit any of the properties except for the recipient name in the **Name** field. To modify the recipient name, you must create a new recipient, and then delete the old one.

- 6 Click **OK**.

Editing recipient notification methods

You can configure new notification methods or edit existing notification methods after you configure recipients.

To edit recipient notification methods

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Assign recipients to alert categories**.
- 4 Under **Recipients**, click **Settings**.
- 5 Edit notification properties for the following types of notification methods:
 - SMTP Configuration.
See [“Configuring SMTP for email or mobile phone text message notification”](#) on page 662.
 - MAPI Configuration.
See [“Configuring MAPI email for notification”](#) on page 664.
 - VIM Configuration.
See [“Configuring VIM email for notification”](#) on page 665.
 - Pager Configuration. Click **Enable** to activate or clear the notification method, and then select a modem from the Configured Modems list.
- 6 Click **OK**.

Removing recipients

You can delete recipients that do not want to receive notification messages; however, the recipient is permanently removed upon deletion. If you want to keep the recipient, but do not want the recipient to receive notifications, clear the **Enable** check box in the recipient properties.

To remove a recipient

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Configure recipients**.
- 4 Select the recipient you want to delete, and then click **Remove**.

Assigning recipients to alert categories for notification

You can assign recipients to alert categories to receive notification messages. When an alert occurs, all the recipients assigned to the alert category receive the notification message. You can also clear a recipient from an alert category and edit properties for the alert categories during the alert notification setup.

To assign a recipient to an alert category for notification

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Assign recipients to alert categories**.
- 4 Under **Alert Categories**, select the alert category to which you want to assign recipients.
- 5 Under **Recipients**, select the recipients you want to receive notification messages for the alert category, or click **Check All** to select all the recipients in the list.
- 6 Click **OK**.

Assign Recipients to Alert Categories options

You can assign recipients to alert categories to receive notification messages. See [“Assigning recipients to alert categories for notification”](#) on page 680.

Table 15-20 Assign Recipients to Alert Categories options

Item	Description
Alert categories	Lists the alert categories to which you can assign recipients.
Properties	Lets you view or change the properties of a selected alert category.
Recipients	Lists the recipients that you can assign to receive alert notifications.
New	Lets you create a new recipient.

Table 15-20 Assign Recipients to Alert Categories options (continued)

Item	Description
Remove	Lets you remove a recipient from the list of recipients that can receive alert notifications. You can dissociate a recipient from an alert category without removing that recipient from the recipient list. See “Stopping alert notification for a recipient” on page 681.
Properties	Lets you view or change the properties for the selected recipient.
Check All	Lets you assign all recipients in the Recipients list to a selected alert category. All recipients receive notifications for the selected alert category.
Settings	Lets you view or change the notification configuration properties.

Stopping alert notification for a recipient

When a recipient no longer needs to receive notifications for an alert category, you can stop the notification.

To stop alert notification for a recipient

- 1 On the navigation bar, click **Alerts**.
- 2 Click **Active Alerts**.
- 3 Under **Notification Tasks** in the task pane, click **Assign recipients to alert categories**.
- 4 Under **Alert Categories**, select the alert category for which you want to stop the notifications for a recipient.
- 5 Under **Recipients**, uncheck the check boxes for the recipients for which you want to stop the notifications.
- 6 Click **OK**.

Sending a notification when a job completes

You can assign recipients to be notified when a job completes. Recipients must be set up before you can set up notification.

To send a notification when a job completes

- 1 Create a new job or edit an existing job.
- 2 In the **Properties** pane, under **Settings**, click **Notification**.
- 3 Select the recipients you want to notify when the job completes.
- 4 To send the job log with the notification to an email address or a printer, check **Include job log with a notification to an email or printer recipient**.
- 5 You can continue selecting other options from the **Properties** pane.

Sending a notification when a selection list is used in a job

Recipients must be set up before you can set up notification.

To send a notification when a selection list is used in a job

- 1 On the navigation bar, click **Job Setup**.
- 2 In the **Backup Selection Lists** pane, select the selection list for which you want to send notifications.
- 3 In the **taskpane**, under **General Tasks**, click **Properties**.
- 4 In the **Properties** pane, under **Source**, click **Selection List Notification**.
- 5 Select the recipients who should receive notification when the selection list is used in a job.
- 6 Click **OK**.

Notification options for jobs

When you set up or edit a job, you can select recipients to receive notification when the job completes. When you set up or edit a selection list, you can select recipients to receive notification when the selection list is used in a job.

Table 15-21 Notification options for jobs

Item	Description
Recipient type	Lists the type of recipient that is available, such as Person, Group, Printer, or Net Send.
Recipient name	Lists the name of the recipient.
Include job log with a notification to an e-mail or printer recipient	Enables Backup Exec to include a copy of the job log with the notification. This option applies only to person recipients who are set up to receive email notification and for printer recipients.
Properties	Lets you view or change the properties of a selected recipient.

About SNMP notification

SNMP (Simple Network Management Protocol) is a method by which a network can be monitored from a central location. SNMP-enabled network applications (like Backup Exec) report to an SNMP Console (a management workstation). The console receives messages (traps) from Backup Exec regarding status and error conditions. A MIB is available in the WINNT\SNMP*language* directory on the Backup Exec installation media that you can load into your SNMP console.

The Object Identifier prefix for Symantec is:

1.3.6.1.4.1.1302

Backup Exec SNMP Traps (messages) have unique object IDs and may include up to four strings.

The following SNMP Trap types are supported:

Table 15-22 SNMP Traps

Trap Type	Object ID	String 1	String 2	String 3	String 4
Product Start	1302.3.1.1.9.1	Backup Exec: Application initializing	machine name	product, version, revision	
Product Stop	1302.3.1.1.9.2	Backup Exec: Application terminating	machine name	product, version, revision	

Table 15-22 SNMP Traps (*continued*)

Trap Type	Object ID	String 1	String 2	String 3	String 4
Job Canceled	1302.3.1.2.8.2	Backup Exec: Job canceled by Operator	machine name	job name	local or remote Operator name
Job Failed	1302.3.1.2.8.1	Backup Exec: Job failed	machine name	job name	detail message
Storage device requires human intervention	1302.3.2.5.3.3	Backup Exec: Storage device requires attention	machine name	job name	detail message
Robotic library requires human intervention	1302.3.2.4.3.3	Backup Exec: robotic library device requires attention	machine name	job name	detail message
Intelligent Disaster Recovery Message	1302.3.1.4.2.1.1	Copy to alternate path failed	machine name	job name	detail message
Intelligent Disaster Recovery Message	1302.3.1.4.2.1.2	Backup complete, update DR disks	machine name	job name	detail message
Backup Exec system error	1302.3.1.1.9.3	The application has encountered an error	machine name	job name	detail message
Backup Exec general information	1302.3.1.1.9.4	Information on normal events	machine name	job name	detail message
Job Success	1302.3.1.2.8.3	The job succeeded	machine name	job name	detail message
Job Success with exceptions	1302.3.1.2.8.4	The job succeeded, but there was a problem	machine name	job name	detail message
Job Started	1302.3.1.2.8.5	The job has started	machine name	job name	detail message
Job Completed with no data	1302.3.1.2.8.6	The job succeeded, but there was no data	machine name	job name	detail message

Table 15-22 SNMP Traps (*continued*)

Trap Type	Object ID	String 1	String 2	String 3	String 4
Job Warning	1302.3.1.2.8.7	The job has a warning	machine name	job name	detail message
PVL Device Error	1302.3.1.5.1.1.1	The device has encountered an error	machine name	job name	detail message
PVL Device Warning	1302.3.1.5.1.1.2	The device has encountered a warning	machine name	job name	detail message
PVL Device Information	1302.3.1.5.1.1.3	Normal device information	machine name	job name	detail message
PVL Device Intervention	1302.3.1.5.1.1.4	Device requires attention	machine name	job name	detail message
PVL Media Error	1302.3.1.5.2.1.1	There is an error with the media	machine name	job name	detail message
PVL Media Warning	1302.3.1.5.2.1.2	There may be a problem with the media	machine name	job name	detail message
PVL Media Information	1302.3.1.5.2.1.3	Normal media information	machine name	job name	detail message
PVL Media Intervention	1302.3.1.5.2.1.4	Media requires attention	machine name	job name	detail message
Catalog Error	1302.3.1.5.3.1.1	There is an error with the catalog	machine name	job name	detail message
Tape Alert Error	1302.3.1.5.4.1.1	There is a TapeAlert error	machine name	job name	detail message
Tape Alert Warning	1302.3.1.5.4.1.2	There is a TapeAlert warning	machine name	job name	detail message
Tape Alert Information	1302.3.1.5.4.1.3	Normal TapeAlert information	machine name	job name	detail message

Table 15-22 SNMP Traps (*continued*)

Trap Type	Object ID	String 1	String 2	String 3	String 4
Database Maintenance Error	1302.3.2.5.5.1.1	There is a database maintenance error	machine name	job name	detail message
Database Maintenance Information	1302.3.2.5.5.1.2	Normal database maintenance information	machine name	job name	detail message
Software Update Error	1302.3.2.5.6.1.1	There is a software update error	machine name	job name	detail message
Software Update Warning	1302.3.2.5.6.1.2	There is a software update warning	machine name	job name	detail message
Software Update Information	1302.3.2.5.6.1.3	Normal software update information	machine name	job name	detail message
Install Update Warning	1302.3.2.5.7.1.1	There is an install warning	machine name	job name	detail message
Install Update Information	1302.3.2.5.7.1.2	Normal Install information	machine name	job name	detail message

See [“Installing and configuring the SNMP system service”](#) on page 686.

Installing and configuring the SNMP system service

In order to receive Backup Exec traps at the SNMP console, you must configure the SNMP system service with the SNMP console's IP address.

SNMP starts automatically after installation. You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

To install the SNMP system service and configure it to send traps to the SNMP console for Windows 2000 and Windows server 2003

- 1 Click **Start**, point to **Settings**, point to **Control Panel**, and then double-click **Add/Remove Programs**.
- 2 Click **Add/Remove Windows Components**.
- 3 In Add/Remove Windows Components, select **Management and Monitoring Tools**, and then click **Details**.

When selecting the component, do not select or clear its check box.

- 4 Select **Simple Network Management Protocol**, and then click **OK**.
- 5 Click **Next**.

Installing the Windows Management Instrumentation performance counter provider

Windows Management Instrumentation (WMI) is an infrastructure through which you can monitor and control system resources. Backup Exec includes performance counter and SNMP providers that can be manually installed and used with WMI.

To install the WMI performance counter provider

- 1 Insert the Backup Exec Installation media.
- 2 At the command prompt, type the following:

```
mofcomp <CD Drive Letter>:\winnt\wmi\backupexecperfmon.mof
```

Installing the Windows Management Instrumentation provider for SNMP

Windows Management Instrumentation (WMI) is an infrastructure through which you can monitor and control system resources. Backup Exec includes performance counter and SNMP providers that can be manually installed and used with WMI.

To use the WMI SNMP provider you must set up SNMP notification.

To install the WMI SNMP provider

- 1 Before you install the SNMP provider included with Backup Exec, you must have the Microsoft SNMP provider installed on your system.

For more information, refer to your Microsoft documentation.

- 2 Insert the Backup Exec Installation media.
- 3 At the command prompt, type the following:

```
mofcomp <CD Drive Letter>:\winnt\wmi\snmp\eng\bkupexecmib.mof
```

Uninstalling the Windows Management Instrumentation performance counter provider

You must uninstall the Windows Management Instrumentation (WMI) performance counter provider and the WMI SNMP provider separately.

To uninstall the WMI performance counter provider

- ◆ At the command line, type:

```
mofcomp <CD Drive  
Letter>:\winnt\wmi\deletebackupexecperfmon.mof
```

Uninstalling the Windows Management Instrumentation provider for SNMP

You must uninstall the Windows Management Instrumentation (WMI) performance counter provider and the WMI SNMP provider separately.

To uninstall the WMI SNMP provider

- ◆ At the command line, type:

```
Smi2smir /d Backup_Exec_MIB
```

Reports in Backup Exec

This chapter includes the following topics:

- [About reports in Backup Exec](#)
- [Viewing the list of available reports](#)
- [Running a report](#)
- [Additional settings for standard reports](#)
- [Available groups for creating reports](#)
- [Running a new report job](#)
- [Saving a report](#)
- [Saving a report to a new location](#)
- [Printing a report from the Backup Exec Report Viewer](#)
- [Printing a report that is saved in PDF format](#)
- [Printing a report that is saved in HTML format](#)
- [Deleting a report from Job History](#)
- [About scheduling report jobs and setting notification recipients](#)
- [About custom reports in Backup Exec](#)
- [Creating a custom report](#)
- [Setting filters for custom reports](#)
- [Copying custom reports](#)
- [Editing custom reports](#)

- [Deleting custom reports](#)
- [Editing application settings for reports](#)
- [Viewing report properties](#)
- [Available reports](#)

About reports in Backup Exec

Backup Exec includes standard reports that show detailed information about your system. When generating most of the reports, you can specify settings that serve as filter parameters or a time range for the data that you want to include in the report. You can then run and view the report immediately, or you can create a new job that saves the report data in the Job History. You can also view general properties for each report.

Backup Exec also provides the following:

- The ability to schedule a report to run at a specified time or to specify a recurring schedule for the report to run.
- The ability to have Backup Exec distribute reports through notification.

To run reports across multiple media servers, you must install the Backup Exec SAN Shared Storage Option, even if you are not operating in a shared storage environment.

Reports can be viewed and printed in the following formats:

- PDF
- HTML
- XML
- Microsoft Excel (XLS)
- Comma Separated Value (CSV)

To properly format integrated Backup Exec reports, you must configure a default printer using the Windows Control Panel Printers applet. This is required even if you do not have a printer attached to your system.

For information on configuring a printer by using the Windows Control Panel Printers applet, see your Microsoft Windows documentation.

See [“Viewing the list of available reports”](#) on page 691.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

See [“About scheduling report jobs and setting notification recipients”](#) on page 698.

See [“Available reports”](#) on page 722.

Viewing the list of available reports

Use the following steps to view the list of available reports.

See [“About reports in Backup Exec”](#) on page 690.

To view the list of available reports

- 1 On the navigation bar, click **Reports**.
- 2 In the Reports pane, click **All Reports**.
- 3 To sort the list of available reports, click the column heading on which you want to sort.

Running a report

When you run a report, you can specify the criteria that is used to determine the items that will be included in the report. The settings, or parameters, available for you to select depend on the type of data that can be included in the report. After the report is generated, only the items that match the criteria appear in the report.

See [“Saving a report”](#) on page 696.

See [“Printing a report from the Backup Exec Report Viewer”](#) on page 697.

See [“About scheduling report jobs and setting notification recipients”](#) on page 698.

See [“Available reports”](#) on page 722.

To run a report

- 1 On the navigation bar, click **Reports**.
- 2 On the **Reports** pane, select the report you want to run.
- 3 In the **task** pane, under **General Tasks**, click **Run report now**.
- 4 If the **Run Report Now Properties** appears, select the appropriate settings, or filter parameters, for the data you want to include in the report.

Only filter parameters that are available for a report appear. Select the appropriate options.

See [“Additional settings for standard reports”](#) on page 692.

5 Click **Run Now**.

The report appears and displays data based on the criteria you set when you ran the report.

6 After you have finished viewing the report, click **OK**.

Backup Exec automatically deletes the report when you close the Report Viewer.

Additional settings for standard reports

You can set additional report settings when you run a report or create a new report. Only settings that are available for a report appear.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

The following table describes the settings you can set for a report:

Table 16-1 Additional settings for standard reports

Item	Description
Media set	Filters the report based on media set names. Media sets include all the media that is inserted into the storage device.
Media server	Filters the report based on media server names. The media server is the server on which Backup Exec is installed. This setting is only available if the SAN Shared Storage Option is available.
Job status	Filters the report based on job status.
Protected server	Filters the report based on specific protected server names. The protected server is the server that is being backed up.
Vault	Filters the report based on specific vault names. A media vault is a virtual representation of the actual physical location of media. See “Media locations and vaults” on page 250.

Table 16-1 Additional settings for standard reports (*continued*)

Item	Description
Ranges	<p>Filters the report based on the time range for the data that you want to include in the report. If range parameters are not available for a report, you will not be able to select the parameter.</p> <p>Range parameters or options available include the following:</p> <ul style="list-style-type: none"> ■ Days. Enables the date filter. <ul style="list-style-type: none"> - Number of days before day report runs. Specifies the number of days prior to the current day to begin the filter process on the data to be included in the report. You can enter a minimum of 0 and a maximum of 32,000 days. - Number of days after day report runs. Specifies the number of days after the current day to begin the filter process on the data to be included in the report. You can enter a minimum of 0 and a maximum of 32,000 days. ■ Hours. Enables the hours filter. <ul style="list-style-type: none"> - Number of hours within time report. Specifies the number of hours either before or after the present hour to filter the data to be included in the report. The time frame depends on the type of report. You can enter a minimum of 0 and a maximum of 32,000 hours. ■ Event count. Enables the event count filter. <ul style="list-style-type: none"> - Maximum number of events to include. Specifies the number of events to include in the report. Events generate alerts and originate from one of the following sources: system, job, media, or device. You can enter a minimum of 0 and a maximum of 32,000 events. Entering a value of zero for the range parameter does not limit the amount of data included in the report; this can result in an extensive report.

Available groups for creating reports

Select a group for which you want to create a report.

See [“Creating a custom report”](#) on page 699.

Table 16-2 Group selections for creating reports

Group	Description
Alerts Group	Includes fields for information such as the alert message text, the alert title, when the alert was created, and the name of the responder.
Device Group	Includes fields for information such as the number of bytes that were read or written, number of hours the device was in use, and the number of errors on the device.
Job Group	Includes fields for information such as the job priority, the job name, the due date, and the policy name.
Job History Group	Includes fields for information such as the backup rate, the device used, errors, and media.
Media Group	Includes fields for information such as the backup set date and time, the backup type, the date allocated and modified, and the media set name.
Policy Group	Includes fields for information such as the job priority, the policy name and description, the selection list name, and the due date.

Running a new report job

You can create a report job that saves the report data in the Job History. You can specify filters and filter ranges. You can also select recipients for notification; however, the report is not included in the notification. Report jobs run immediately and you cannot specify a schedule.

See [“Saving a report”](#) on page 696.

To run a new report job

- 1 On the navigation bar, click **Reports**.
- 2 On the **Reports** pane, select the report for which you want to run a job.
- 3 Under **General Tasks** in the **task** pane, click **New report job**.

- 4 On the **Properties** pane, under **Settings**, click **General** and then type the name for the job in **Job name** and select the Job priority.
If another job is scheduled to run at the same time as this job, the priority you set determines which job runs first.
- 5 On the **Properties** pane, under **Settings**, select the appropriate filter parameters for the data you want to include in the report.
If filter parameters or settings are not available for a report, you will not be able to view the parameters.
See [“Additional settings for standard reports”](#) on page 692.
- 6 To notify recipients when the report job completes and to send the completed report to the recipients, do the following in the order listed.
 - On the **Properties** pane, under **Settings**, click **Notification**.
 - Select the recipients who you want to receive notification when the report job completes.
See [“Configure Recipients options”](#) on page 666.
 - To include a copy of the completed report along with the notification, check **Include job log with a notification to an e-mail or printer recipient**.
- 7 On the **Properties** pane, under **Frequency**, click **Schedule** and then click **Submit job on hold** if you want to submit the job with an on-hold status.
Select this option if you want to submit the job, but do not want the job to run until you change the job’s hold status.
- 8 Click **Enable automatic cancellation**, and then type the number of hours or minutes in the **Cancel job** if not completed within option.
Select this option if you want to cancel the job if is not completed within the selected number of hours or minutes. Backup Exec starts timing the length of time the job takes to run when the job actually begins, not the scheduled time.
- 9 After you have completed all the items you want to set for the new report job, click **Run Now**.
The report is submitted according to the options you selected.

General options for a new report job

When you create a new report job, you can give the job a unique name and you can set the job priority level.

You can set other job options for the report as well.

See [“Additional settings for standard reports”](#) on page 692.

Table 16-3 General options for a new report job

Item	Description
Job name	Indicates a name for the job or accept the default name.
Job priority	Sets the priority level for this job. See “About job priority” on page 195.

Saving a report

Use the following steps to save a report.

See [“Printing a report from the Backup Exec Report Viewer”](#) on page 697.

To save a report

- 1 On the report, click **Save As**.
- 2 When prompted, enter the file name and location where you want to save the report.
- 3 In the **Save as type** box, select a format in which to save the report.
When you save a report in HTML format, a folder is created in location where you save the report. The folder is named with the name that you specify for the report. The folder contains both HTML files and a .GIF image file.
- 4 Click **Save**.

Saving a report to a new location

You can specify a location where a report is saved. Backup Exec also creates a folder, with the same name as the report, in the same location in which the report is saved. The folder contains images and report pages that enable you to view the saved report.

See [“Printing a report from the Backup Exec Report Viewer”](#) on page 697.

To save the report to a new location

- 1 On the navigation bar, click **Job Monitor**.
- 2 Click the **Job List** tab.
- 3 On the **Job History** pane, right-click the report you want to save.

- 4 Click **Properties**.
- 5 Click **Save As**.
- 6 Enter the file name and location where you want to save the report and then click **Save**.

Printing a report from the Backup Exec Report Viewer

You can print reports from a locally-attached printer or a network printer. To print a report, the printer must be configured to print in the landscape mode.

See [“Saving a report”](#) on page 696.

See [“Printing a report that is saved in PDF format”](#) on page 697.

See [“Printing a report that is saved in HTML format”](#) on page 697.

To print a report from the Backup Exec Report Viewer

- 1 On the **Report Viewer**, click **Print**.
- 2 Read the message about printing options and then click **OK**.
- 3 Select a printer from the Windows **Print** dialog box.
- 4 Click **Print**.

Printing a report that is saved in PDF format

Use the following steps to print a report that you saved in PDF format.

Note: You must have the Adobe Reader installed on the computer where you want to print a report that you saved in PDF format.

To print a report that is saved in PDF format

- 1 Navigate to the folder where you saved the report in PDF format.
- 2 Open the report by double-clicking the report's PDF icon.
- 3 From the Adobe Reader menu bar, click **File > Print**.

Printing a report that is saved in HTML format

Use the following steps to print a report that you saved in HTML format.

To print a report that is saved in HTML format

- 1 Navigate to the location where you saved the HTML report.
- 2 Double-click the folder name of the report that you saved.
- 3 Right-click the HTML file named RPT<number>_.htm.
For example, RPT3_.htm
- 4 On the shortcut menu, click **Print**.
- 5 Select a printer from the Windows **Print** dialog box.
- 6 Click **Print**.

Deleting a report from Job History

A report that you create using the **Run report now** option is automatically deleted after you view the report. A report that you create when you select **New report job** is saved in the Backup Exec database until you delete the report from Job History.

See [“Configuring database maintenance”](#) on page 209.

See [“Viewing the properties for completed jobs”](#) on page 572.

See [“Available reports”](#) on page 722.

To delete the report from Job History

- 1 On the navigation bar, click **Job Monitor**.
- 2 Click the **Job List** tab.
- 3 On the **Job History** pane, select the report you want to delete.
- 4 Under **General Tasks** in the task pane, click **Delete**.
- 5 Confirm the job deletion.

About scheduling report jobs and setting notification recipients

You can create a report job and schedule it to run at a specific time or specify a recurring schedule for a report to run.

See [“Running a new report job”](#) on page 694.

See [“Scheduling jobs”](#) on page 356.

See [“Configure Recipients options”](#) on page 666.

You can also assign notification recipients to the report job just as you would for other Backup Exec jobs, such as backups and restores. If you select Include job log with a notification to an e-mail or printer recipient, the report is included with the notification. If this option is not selected, the recipient only gets a message that the report has run.

About custom reports in Backup Exec

You can create reports that contain information to meet the specific requirements of your organization. You choose the data to include in the report, and then determine how the data is filtered, sorted, and grouped. In addition, you can set up pie graphs and bar graphs to graphically represent the report data.

You can customize the look of the reports by doing the following:

- Adding your company logo to the report
- Changing the color of the banner
- Adding text to the footer

See [“Creating a custom report”](#) on page 699.

Creating a custom report

You can create reports that contain information to meet the specific requirements of your organization.

To create a custom report

- 1 On the navigation bar, click **Reports**.
- 2 On the **task** pane, click **New custom report**.
- 3 On the **Custom Report** dialog box, type a name and description for the report.
- 4 If you do not want this report to include the default header and footer settings, uncheck **Use header and footer settings specified in Tools/Options**.
- 5 On the **properties** pane, under **Report Definition**, click **Field Selection**.
- 6 In the **Category** box, select a group for which you want to create a report.

See [“Available groups for creating reports”](#) on page 693.

- 7 Select the fields that you want on the report.

See [“Field options for custom reports”](#) on page 701.

- 8 To adjust the width of the column for a field, do the following in the order listed:
 - Click the field name in the **Fields selected for the report** list.
 - In the **Column width** box, type the new width.
 - Click **Set**.
- 9 Do any of the following:

To set filter criteria for the report See [“Setting filters for custom reports”](#) on page 712.

To group fields for the report Do the following in the order listed:

- On the properties pane, under **Report Definition**, click **Grouping**.
- Complete the appropriate grouping options.
See [“About grouping fields in custom reports”](#) on page 702.

To sort fields for the report Do the following in the order listed:

- On the properties pane, under **Report Definition**, click **Sorting**.
- Complete the appropriate sorting options.
See [“Sorting fields in custom reports”](#) on page 704.

To set graph options for the report Do the following in the order listed:

- On the properties pane, under **Report Definition**, click **Graph Options**.
- Complete the appropriate graphing options.
See [“Setting graph options in custom reports”](#) on page 706.

To preview and test the report Do the following in the order listed:

- On the properties pane, under **Preview**, click **Preview**.
- To test the report, on the **Preview** dialog box, click **Test Report**.

To finish and close the report Click **OK**.

Custom report name and description options

You can give a report that you create a unique name. You can also enter a detailed description of the report.

See [“Creating a custom report”](#) on page 699.

Table 16-4 Custom report name and description options

Item	Description
Name	Indicates a unique name for the report. All custom reports must be named.
Description	Indicates the description of the report.
Use header and footer settings specified in Tools/Options	Enable this option to display header and footer information in custom reports. This option uses the default header and footer settings that you specified for all reports in Reports , under Tools/Options . See “Application settings for reports ” on page 720.

Field options for custom reports

Select the fields that you want to include in the report. Fields are displayed in the order you place them in the **Fields selected for the report** box. All fields are positioned horizontally, from left to right. The first field in the list appears on the left side of the report.

See [“Creating a custom report”](#) on page 699.

Table 16-5 Field selection options

Item	Description
Category	Lets you select fields for a custom report that are based on Backup Exec functionality. Field categories include the following: <ul style="list-style-type: none"> ■ Alerts Group ■ Device Group ■ Job Group ■ Job History Group ■ Media Group ■ Policy Group

Table 16-5 Field selection options (*continued*)

Item	Description
Available fields	<p>Shows the list of available fields for each category.</p> <p>By default, Backup Exec displays only the basic fields for each category. The basic fields include those fields that are most likely to be used in a report. To show all available fields, check Show advanced fields</p> <p>To select consecutive fields, click the first item, press and hold SHIFT, and then click the last item. To select fields that are not consecutive, press and hold CTRL, and then click each item.</p> <p>To move the selected fields to the Fields selected for the report box, click >>.</p>
Fields selected for the report.	<p>Shows the fields that are selected for display on the report.</p> <p>Fields are displayed on the report based on the order in which they appear in the box titled Fields selected for the report. The first field in the list appears on the left side of the report.</p> <p>Click Move UP or Move Down to reposition the fields on the report.</p> <p>To remove a field, double click the item.</p>

About grouping fields in custom reports

You can group a custom report by up to three of the fields that you have chosen for the report. Grouping fields creates sections on the report. For example, if you group by media server, Backup Exec creates a section for each media server that matches the filter criteria. Under each media server’s section, the report displays the data that corresponds to the remaining fields that you selected for the report.

A report must have at least one field that is not grouped. For example, if you select three fields in the report, you can group only two of the fields. If you group all of the fields, no data appears on the report because all of the data is listed in the group section titles. In addition, you must have at least four fields on the report to use all three grouping fields.

After you select a field on which to group the report, you can group the data for that field in ascending or descending order. Ascending order lists numbers from smallest to largest and lists letters in alphabetical order. Descending order lists numbers from largest to smallest and lists letters in reverse alphabetical order. For example, if you group by a date field in ascending order, the report data is grouped by date, starting with the earliest date.

See [“Grouping fields in custom reports”](#) on page 703.

Grouping fields in custom reports

Use the following steps to group fields in custom reports.

See [“About grouping fields in custom reports”](#) on page 702.

To group fields in custom reports

- 1 On the navigation bar, click **Reports**.
- 2 In the tree view, click **Custom Reports**.
- 3 In the reports list, click the report that contains the fields you want to group.
- 4 In the task pane, click **Edit**.
- 5 On the properties pane, under **Report Definition**, click **Grouping**.
- 6 Select the appropriate options.
See [“Grouping options for custom reports”](#) on page 704.
- 7 In the **Group by** box, select the name of the field on which you want to group data.
- 8 Click **Ascending** to group the information in ascending order or click **Descending** to group the information in descending order.
- 9 If you want to group on additional fields, in the **Then group by** box, repeat step 7 and step 8.
- 10 Do any of the following:

To sort fields for the report

Do the following in the order listed:

- On the properties pane, under **Report Definition**, click **Sorting**.
- Complete the appropriate sorting options.
See [“Sorting fields in custom reports”](#) on page 704.

To set graph options for the report

Do the following in the order listed:

- On the properties pane, under **Report Definition**, click **Graph Options**.
- Complete the appropriate graphing options.
 See [“Setting graph options in custom reports”](#) on page 706.

To preview and test the report

Do the following in the order listed:

- On the properties pane, under **Preview**, click **Preview**.
- To test the report, on the **Preview** dialog box, click **Test Report**.

To finish and close the report Click **OK**.

Grouping options for custom reports

You can group report information in ascending or descending order based on the fields that you selected for the report.

See [“About grouping fields in custom reports”](#) on page 702.

Table 16-6 Group options for custom reports

Item	Description
Group by	Groups the report information based on the fields that you select for the report.
Ascending	Groups the report information in ascending order. Ascending order lists numbers from smallest to largest and lists letters in alphabetical order.
Descending	Groups the report information in descending order. Descending order lists numbers from largest to smallest and lists letters in reverse alphabetical order.
Then group on	Lets you group on additional report fields.

Sorting fields in custom reports

You can sort a custom report by up to three of the fields that you have chosen for the report. When you sort on fields, Backup Exec arranges all of the data that matches the sort criteria together in the report. For example, if you sort on the

Media Server field in ascending order, all data for Media Server A displays first, followed by all data for Media Server B, and so on. Ascending order lists numbers from smallest to largest and lists letters in alphabetical order. Descending order lists numbers from largest to smallest and lists letters in reverse alphabetical order.

To sort fields in custom reports

- 1 On the navigation bar, click **Reports**.
- 2 In the tree view, click **Custom Reports**.
- 3 In the reports list, click the report that contains the fields you want to sort.
- 4 In the task pane, click **Edit**.
- 5 On the properties pane, under **Report Definition**, click **Sorting**.
- 6 Select the appropriate sort options.
 See [“Sort options for custom reports”](#) on page 705.
- 7 Do any of the following:

To set graph options for the report

Do the following in the order listed:

- On the properties pane, under **Report Definition**, click **Graph Options**.
- Complete the appropriate graphing options.
 See [“Setting graph options in custom reports”](#) on page 706.

To preview and test the report

Do the following in the order listed:

- On the properties pane, under **Preview**, click **Preview**.
- To test the report, on the **Preview** dialog box, click **Test Report**.

To finish and close the report Click **OK**.

Sort options for custom reports

You can sort report information in ascending or descending order based on the fields that you selected for the report.

See [“Sorting fields in custom reports”](#) on page 704.

Table 16-7 Sort options for custom reports

Item	Description
Sort on	Sorts the report information based on the fields that you select for the report.
Ascending	Sorts the report information in ascending order. Ascending order lists numbers from smallest to largest and lists letters in alphabetical order.
Descending	Sorts the report information in descending order. Descending order lists numbers from largest to smallest and lists letters in reverse alphabetical order.
Then sort on	Lets you sort on additional report fields.

Setting graph options in custom reports

You can include a pie graph or a bar graph in custom reports.

You must select at least two fields on the Field Selection dialog box to create a pie graph, and at least three fields to create a bar graph.

To set graph options in custom reports

- 1 On the navigation bar, click **Reports**.
- 2 In the tree view, click **Custom Reports**.
- 3 In the reports list, click the report for which you want to set graph options.
- 4 In the task pane, click **Edit**.
- 5 On the properties pane, under **Report Definition**, click **Graph Options**.
- 6 In the **Graph type** box, select the type of graph that you want to create. Choices include **Pie** or **Bar**.
- 7 In the **Graph title** box, type the title that you want to display above the graph in the report.
- 8 Complete the options for a pie graph.
 See [“Graph options for custom reports”](#) on page 707.
- 9 Complete the options for a bar graph.
 See [“Graph options for custom reports”](#) on page 707.
- 10 Do any of the following:

To preview and test the report

Do the following in the order listed:

- On the properties pane, under Preview, click **Preview**.
- To test the report, on the Preview dialog box, click **Test Report**.

To finish and close the report Click **OK**.

Graph options for custom reports

You can select to include either a pie or bar graph in a custom report. After you select the graph type, you can select specific options for the graph.

See “[Setting graph options in custom reports](#)” on page 706.

The following table describes the available pie graph options:

Table 16-8 Pie graph options for custom reports

Item	Description
Category field (pie section per value)	Specifies the field for which you want to display sections in the pie chart.
Data field	Specifies the field for which you want to calculate values.
Aggregation function	<p>Selects the way that you want Backup Exec to calculate the values generated for the Data field.</p> <p>Choices include the following:</p> <ul style="list-style-type: none"> ■ Minimum. Calculates the lowest value. This is available for numeric fields only. ■ Maximum. Calculates the highest value. This is available for numeric fields only. ■ Average. Calculates the average value. This is available for numeric fields only. ■ Count. Calculates the number of values. This option is the only available option for non-numeric fields, such as text fields or date fields, but it is also available for numeric fields. ■ Sum. Calculates the sum of the values. This is available for numeric fields only.

The following table describes the available bar graph options:

Table 16-9 Bar graph options for custom reports

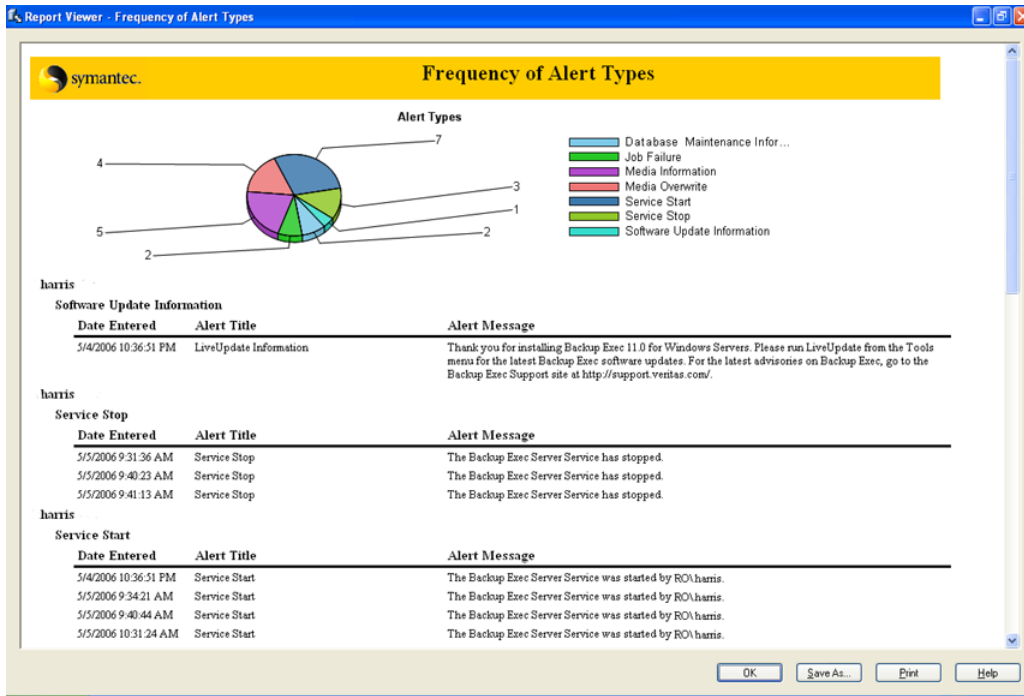
Item	Description
Vertical axis title	Specifies the title that you want to display to the left of the graph. The title will display vertically in the report. There is a 50-character limit.
Series field (bar per value)	Specifies the field that contains the values that you want to display on the horizontal bars of the graph. Backup Exec creates a legend for the values.
Category field (set of series bars per value)	Specifies the field that contains the information for which you want to group information along the left side of the graph.
Data field	Specifies the field for which you want to calculate values.
Horizontal axis title	Specifies the title that you want to display below the graph.
Aggregation function	<p>Specifies one of the following functions to calculate the values generated for the Data field.</p> <p>Choices include the following:</p> <ul style="list-style-type: none"> ■ Minimum. Calculates the lowest value. ■ Maximum. Calculates the highest value. ■ Average. Calculates the average value. ■ Count. Calculates the number of values. This option is available only for non-numeric fields, such as text fields or date fields. ■ Sum. Calculates the sum of the values.

Example graphs for custom reports

This section includes three examples of graphs that you can create in custom reports. In addition, the fields that were used to create the graphs are included.

Review the examples to learn how the fields on the Graph Options dialog box correspond to completed graphs.

Figure 16-1 Example pie graph




The example pie graph, titled Frequency of Alert Types, was created using the following:

Table 16-10 Example pie graph options

Graph option name	Fields selected
Graph Type	Pie
Graph title	Frequency of Alert Types
Category field	Event name
Data field	Event name
Aggregation function	Count

You can create a bar graph.

Figure 16-2 Example bar graph 1



The example bar graph, titled Average Duration of Job per Device per Selection List, was created using the following:

Table 16-11 Example bar graph 1 options

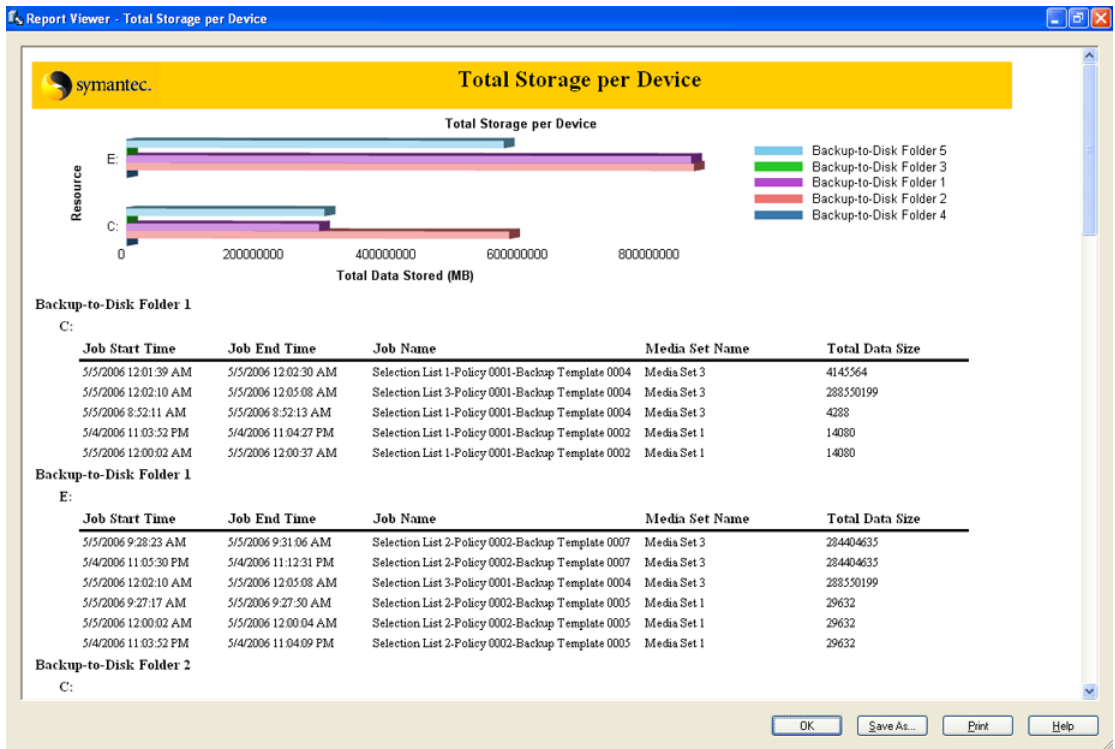
Graph option name	Fields selected
Graph Type	Bar
Graph title	Average duration of job per device per selection list
Vertical axis title	Selection list
Series field	Device name
Category field	Selection list name
Data field	Job duration

Table 16-11 Example bar graph 1 options (*continued*)

Graph option name	Fields selected
Aggregation function	Average
Horizontal axis title	Average duration of a job (seconds)

Compare bar graph 2 with bar graph 1.

Figure 16-3 Example bar graph 2



The example bar graph, titled Total Storage Per Device, was created using the following:

Table 16-12 Example bar graph 2 options

Graph option name	Fields selected
Graph Type	Bar
Graph title	Total Storage Per Device
Vertical axis title	Resource
Series field	Device name
Category field	Resource name
Data field	Total data size
Aggregation function	Sum
Horizontal axis title	Total Data Stored (MB)

Previewing custom reports

Use the preview feature to verify that you created a custom report correctly.

To preview custom reports

- 1 On the navigation bar, click **Reports**.
- 2 In the tree view, click **Custom Reports**.
- 3 In the reports list, click the report that you want to preview and test.
- 4 In the task pane, click **Edit**.
- 5 On the properties pane, under **Preview**, click **Preview**.
- 6 Click **OK**.

Setting filters for custom reports

Use the following steps to set filters for custom reports that you want to create.

To set filters for custom reports

- 1 On the navigation bar, click **Reports**.
- 2 In the tree view, click **Custom Reports**.
- 3 In the reports list, click the report that you want to filter.

- 4 On the task pane, click **Edit**.
- 5 On the properties pane, under **Report Definition**, click **Filters**.
- 6 Create a filter by defining one or more filter expressions.
 See “[Filter expressions for defining custom reports](#)” on page 715.
- 7 Click **Add**.
- 8 Repeat step 6 and step 7 to add more filters.
- 9 To combine sets of filter expressions, do any of the following:

To combine two filter expressions so that both expressions must be true for the result to be true

Click **AND**.

For example, to find all backup jobs that failed, add the following expressions:

- Status = Failed
- Type = Backup

After you set up the expressions, do the following:

- Click **AND** to combine the two expressions.

The combined expression is:

Status = Failed AND Type = Backup

To combine two filter expressions so that one of the expressions must be true for the result to be true

Click **OR**.

For example, to find jobs that either failed or were canceled, add the following expressions:

- Status = Failed
- Status = Canceled

After you set up the expressions, do the following:

- Click **OR** to combine Status = Failed with Status = Canceled.

The combined expression is:

Status = Failed OR Status = Canceled

To combine two filter expressions into a single expression

Click () +

For example, to find backup jobs and restore jobs that failed, add the following expressions:

- Status = Failed
- Type = Backup
- Type = Restore

After you set up the expressions, do the following:

- Use OR to combine Type = Backup with Type = Restore.
- Press and hold Ctrl while you click Type = Backup and Type = Restore.
- Click () + to combine Type = Backup with Type = Restore.
- Use AND to combine Status = Failed with (Type = Backup OR Type = Restore).

The combined expression is:

Status = Failed AND (Type = Backup OR Type = Restore)

To separate two filter expressions that were combined into a single expression

Click () -

For example, if you used () + to combine Type = Backup with Type = Restore, it is displayed on the Filters dialog box like this:
(Type = Backup OR Type = Restore)

To make the combined expression into two individual expressions, do the following:

- Press and hold Ctrl while you click both Type = Backup and Type = Restore.
- Click () -

After you separate the expressions, they are displayed without the parentheses.

10 To change any of the expressions, do the following in the order listed:

- In the Filter criteria box, select the expression that you want to change.
- Click **Edit**.
- In the Filter expression area, edit the expression's values.
- Click **Update**.

11 To remove an expression, select the expression, and then click **Remove**.

12 Do any of the following:

- To group fields for the report

Do the following in the order listed:

 - On the properties pane, under **Report Definition**, click **Grouping**.
 - Complete the appropriate grouping options.
See [“About grouping fields in custom reports”](#) on page 702.

- To sort fields for the report

Do the following in the order listed:

 - On the properties pane, under **Report Definition**, click **Sorting**.
 - Complete the appropriate sorting options.
See [“Sorting fields in custom reports”](#) on page 704.

- To set graph options for the report

Do the following in the order listed:

 - On the properties pane, under **Report Definition**, click **Graph Options**.
 - Complete the appropriate graphing options.
See [“Setting graph options in custom reports”](#) on page 706.

- To preview and test the report

Do the following in the order listed:

 - On the properties pane, under **Preview**, click **Preview**.
 - To test the report, on the **Preview** dialog box, click **Test Report**.

- To finish and close the report

Click **OK**.

Filter expressions for defining custom reports

You can create a filter by defining one or more filter expressions
 See [“Setting filters for custom reports”](#) on page 712.

Table 16-13 Filter expressions for defining custom reports

Item	Description
Show advanced fields	Check Show advanced fields to see all of the fields that are available for filtering. By default, only the most common fields display.
Field name	Select the field on which you want to filter.

Table 16-13 Filter expressions for defining custom reports (*continued*)

Item	Description
Operator	

Table 16-13 Filter expressions for defining custom reports (*continued*)

Item	Description
	<p>Select the appropriate operator for this filter. Operators determine how the field name and the value are linked. The following operators are available in Backup Exec, but the list that displays varies depending on the type of field you selected in Field name:</p> <ul style="list-style-type: none"> ■ = (Equal). The field name must equal the value. ■ <> (Not Equal). The field name must not equal the value. ■ > (Greater Than). The field name must be larger than the value. ■ >= (Greater Than or Equal). The field name must be larger than or equal to the value. ■ < (Less Than). The field name must be smaller than the value. ■ <= (Less Than or Equal). The field name must be smaller or equal to the value. ■ \$ (Contains). The field name contains the text entered in the Value field. ■ NOT\$ (Contains). The field name does not contain the text entered in the Value field. ■ IN LAST. A date or time window that is relative to the time that you create the report. This operator defines dates and times prior to the time when the report is created. This operator is available only for date and time fields. <p>If you enter hours in the Value field you receive more specific results than if you enter days. The Day value is calculated beginning at midnight (00:00) yesterday and ending at the time when the report runs.</p> <p>For example, if you enter 1 day in the Value field and the report runs at 23:59 today, the report includes results for the last 47 hours and 59 minutes. However, if you enter 24 hours, you receive information for exactly 24 hours prior to</p>

Table 16-13 Filter expressions for defining custom reports (*continued*)

Item	Description
	<p>the time when the report runs.</p> <ul style="list-style-type: none"> ■ IN NEXT. A date or time window that is relative to the time that you create the report. This operator defines dates and times after the time when the report is created. For example, to find backup jobs that are scheduled to occur during the next three days, select this operator, and then enter 3 days in the Value field. This is available only for date and time fields.
Value	<p>Type or select the value on which you want to filter. The type of value that you can enter varies depending on the type of field name that you select. For example, if you select Next Due Date in the Field name, Backup Exec displays date and time values.</p>

Copying custom reports

You can create a copy of a custom report, and then modify the copy.

To copy custom reports

- 1 On the navigation bar, click **Reports**.
- 2 In the tree view, click **Custom Reports**.
- 3 In the reports list, click the report that you want to copy.
- 4 In the task pane, click **Copy**.
- 5 In the **Name of copy** box, type a unique name for the copied report.
- 6 Click **OK**.

Editing custom reports

If the report that you want to edit has been run in a previous report job, the changes you make now may affect the appearance of the reports in job history. Symantec recommends that you copy the report and then edit the copy.

To edit custom reports

- 1 On the navigation bar, click **Reports**.
- 2 In the tree view, click **Custom Reports**.
- 3 In the reports list, click the report that you want to filter.
- 4 In the task pane, click **Edit**.
- 5 Change the report settings as needed.
- 6 Click **OK**.

Deleting custom reports

Before you can delete a custom report, you must delete all of the associated job history records.

To delete job history records that are associated with custom reports

- 1 On the navigation bar, click **Job Monitor**.
- 2 In the **Job History** pane, select the job history for the custom report that you want to delete.

The report name is listed in the **Device Name** column.

- 3 On the task pane, under **General Tasks**, click **Delete**.

To delete custom reports

- 1 On the navigation bar, click **Reports**.
- 2 In the tree view, click **Custom Reports**.
- 3 In the reports list, click the report that you want to filter.
- 4 In the task pane, click **Delete**.
- 5 Click **Yes**.

Editing application settings for reports

You can set Backup Exec to display all reports in either HTML or Adobe Portable Document Format (PDF). The default setting is HTML. The format that you select does not affect the format of the reports sent to users with the notification feature.

In addition, you can edit the application settings for the header and footer for all custom reports.

You can do the following:

- Include a logo in the header.

- Choose a color for the banner in the header.
- Include text in the footer.
- Include the time in the footer.

When you choose a color for the banner, you can type the numbers that correspond to the colors (RGB values), or you can select the color from a chart.

To edit application settings for reports

- 1 On the **Tools** menu, click **Options**.
- 2 On the Properties pane, under **Settings**, click **Reports**.
- 3 Complete the appropriate options.
- 4 Click **OK**.

Application settings for reports

You can change the application settings for all Backup Exec reports.

See [“Editing application settings for reports”](#) on page 719.

The following table describes the application settings for reports:

Table 16-14 Application settings for reports

Item	Description
HTML	Specifies that all reports are displayed in HTML format. This is the default setting.
PDF	Specifies that all reports are displayed in Adobe Portable Document Format (PDF).
Maximum number of rows to include in a report	Indicates the maximum number of rows to show in a report. The default is 10,000 rows.
Show all rows	Displays all rows in a report.
Show only the rows that are unique	Displays only the rows that are unique.
Use company logo image file	Uses your company logo in the header of all custom reports.
Image file path	Identifies the path to the logo that you want to use in all custom reports.

Table 16-14 Application settings for reports (*continued*)

Item	Description
Red	Specifies the number that corresponds to the value for red.
Green	Specifies the number that corresponds to the value for green.
Blue	Specifies the number that corresponds to the value for blue.
Colors	Indicates a basic color to use for a custom report banner. You can also create a custom collar for a custom report banner.
Text	Indicates the text that you want to display in the footer of custom reports.
Include time	Includes the time when the report runs in the footer of custom reports.

Viewing report properties

Report properties provide detailed information about each report. The properties can be viewed, but not edited.

To view report properties

- 1 On the navigation bar, click **Reports**.
- 2 On the **Reports** pane, select the report for which you want to view properties.
- 3 In the task pane, under **Report** tasks, click **Properties**.
See “[General properties for reports](#)” on page 721.
- 4 Click **OK** after you have finished viewing the properties.
See “[Running a report](#)” on page 691.

General properties for reports

You can view but not edit properties for each report.

See “[Viewing report properties](#)” on page 721.

The following table describes the report properties:

Table 16-15 General Report Properties

Item	Description
Title	Displays the name of the report.
Description	Describes the type of data that is included in the report.
Category	<p>Specifies the classification for the report.</p> <p>Available report categories include the following:</p> <ul style="list-style-type: none"> ■ Media ■ Media Vault ■ Jobs ■ Devices ■ Configuration ■ Alerts ■ Template
Author	Displays the creator of the report.
Subject	Displays the version of the product for which the report was created.
File name	Displays the file name of the report.
File size	Displays the size of the report.
Creation Date	Displays the date the report was installed on the system.

Available reports

This section provides detailed information about each report available in Backup Exec. The file name of the report, a description, and the information included in the report are listed for each report. The data included in each report will vary depending on the criteria you selected to include in the report.

The following reports are included in Backup Exec:

Table 16-16 Backup Exec Reports

Report Name	Description
Active Alerts	<p>Lists all active alerts chronologically, displaying the most recent alerts first.</p> <p>See “Active Alerts Report” on page 729.</p>

Table 16-16 Backup Exec Reports (*continued*)

Report Name	Description
Active Alerts by Media Server	<p>Lists all active alerts grouped and filtered by media server, displaying the most recent alerts first.</p> <p>See “Active Alerts by Media Server Report” on page 730.</p>
Alert History	<p>Lists all alerts in the alert history chronologically, displaying the most recent alerts first</p> <p>See “Alert History Report” on page 730.</p>
Alert History by Media Server	<p>Lists all alerts in the alert history grouped and filtered by media server, displaying the most recent alerts first.</p> <p>See “Alert History by Media Server Report” on page 731.</p>
Archive Job Success Rate	<p>Displays the number of archive jobs for the protected servers that successfully ran.</p> <p>See “Archive Job Success Rate report” on page 768.</p>
Archive Selections by Archive Rules and Retention Categories	<p>Displays the archive rules and the retention categories that are applied to each archive selection.</p> <p>See “Archive Selections by Archive Rules and Retention Categories report” on page 768.</p>
Application Event Log	<p>Lists all Backup Exec application event logs.</p> <p>See “Application Event Log Report” on page 732.</p>
Audit Log	<p>Lists the contents of the audit logs for selected servers for the specified time period.</p> <p>See “Audit Log Report” on page 732.</p>
Backup Job Success Rate	<p>Lists the success rate for backup jobs run to protect selected servers.</p> <p>See “Backup Job Success Rate Report” on page 733.</p>
Backup Recommendations	<p>Lists any recommendations that can help you to better manage your backups.</p> <p>See “Backup Recommendations Report” on page 733.</p>

Table 16-16 Backup Exec Reports (*continued*)

Report Name	Description
Backup Resource Success Rate	Lists the success rate for backup jobs for specified past number of days for resources on selected servers. See “Backup Resource Success Rate Report” on page 734.
Backup Set Details by Resource	Lists all backup sets that ran within the last 72 hours. The sets are grouped by the server and resource. See “Backup Set Details by Resource Report” on page 734.
Backup Sets by Media Set	Lists all backup sets by media set. See “Backup Sets by Media Set Report” on page 735.
Backup Size by Resource	Lists the backup size for each resource job for up to seven previous runs and then computes the trailing average for up to seven previous runs for each job run. See “Backup Size By Resource Report” on page 736.
Configuration Settings	Lists the contents of the Backup Exec system configuration parameters table. See “Configuration Settings Report” on page 737.
Current Job Status	Details the job queue sorted by status. See “Current Job Status Report” on page 738.
Daily Device Utilization	Lists the percentage of the storage devices' capacity that the media server uses. See “Daily Device Utilization Report” on page 738.
Deduplication device summary	Displays a summary of the deduplication operations for local deduplication storage folders and shared deduplication storage folders. See “Deduplication device summary” on page 739.
Deduplication summary	Displays a deduplication summary for all of the deduplication jobs that run on the Backup Exec media server. See “Deduplication summary” on page 740.

Table 16-16 Backup Exec Reports (*continued*)

Report Name	Description
Device Summary	Lists device usage and error summary for each selected media server. See “Device Summary Report” on page 740.
Device Usage by Policy	Lists all the policies that are targeted to specific drive selections. See “Device Usage by Policy” on page 741.
Error-Handling Rules	Lists all the defined error-handling rules. See “Error-Handling Rules Report” on page 742.
Event Recipients	Lists all events registered by each notification recipient. See “Event Recipients Report” on page 743.
Exchange Mailbox Group Archive Settings	Displays the archive settings that are applied to mailbox groups in each domain. See “Exchange Mailbox Group Archive Settings report” on page 769.
Failed Archive Jobs	Displays what archive jobs failed recently. See “Failed Archive Jobs report” on page 770.
Failed Backup Jobs	Lists all the failed backup jobs sorted by the resource server and time frame. See “Failed Backup Jobs Report” on page 744.
File System Archive Settings	Displays the archive settings that are applied to archive selections for each server. See “File System Archive Settings report” on page 770.
Job Distribution by Device	Lists all the jobs that have been run on each system device during the specified period. See “Job Distribution by Device Report” on page 745.
Jobs Summary	Lists all the jobs that ran within the last 72 hours in chronological order. See “Jobs Summary Report” on page 745.

Table 16-16 Backup Exec Reports (*continued*)

Report Name	Description
Machines Backed Up	Lists all the servers that have been protected by Backup Exec. See “Machines Backed Up Report” on page 746.
Managed Media Servers	Lists the status and configuration for all media servers managed by Backup Exec. See “Managed Media Servers Report” on page 747.
Media Audit	Lists the recent media configuration changes. See “Media Audit Report” on page 748.
Media Errors	Lists the number of errors that occur on all media. See “Media Errors Report” on page 749.
Media Required for Recovery	Lists the media that contain the backup sets for each system backed up on selected servers for the specified time period. This report can be inaccurate if media overwrite settings allow the media to be overwritten. See “Media Required for Recovery Report” on page 749.
Media Set	Lists all the media sets and media used by Backup Exec servers. The current location is given for each media. See “Media Set Report” on page 750.
Media Vault Contents	Lists the media located in each media vault. See “Media Vault Contents Report” on page 751.
Missed Availability Window	Lists all jobs that have missed scheduled availability windows within the specified time range. The jobs are listed in chronological order. See “Missed Availability Report” on page 752.
Move Media to Vault	Lists all media that can be moved to a media vault. The media listed are not currently in a media vault and the media’s append period has expired. See “Move Media to Vault Report” on page 752.

Table 16-16 Backup Exec Reports (*continued*)

Report Name	Description
Operations Overview	<p>Lists past and future operations data for user-set period.</p> <p>See “Operations Overview Report” on page 753.</p>
Overnight Archive Summary	<p>Displays the status of the archive jobs that ran in the last 24 hours.</p> <p>See “Overnight Archive Summary report” on page 771.</p>
Overnight Summary	<p>Lists the results of backup jobs for each resource during the last 24 hours. This report includes backup jobs that were scheduled to run but did not run. Jobs are given a grace period of 24 hours before being marked as past due.</p> <p>See “Overnight Summary Report” on page 755.</p>
Policy Jobs by Resource Summary	<p>Lists all backup sets that were created in the selected period. The sets are grouped by target server and resource.</p> <p>See “Policy Jobs by Resource Summary Report” on page 756.</p>
Policy Jobs Summary	<p>Lists in chronological order all jobs derived from selected policies that have run within the specified time range.</p> <p>See “Policy Jobs Summary Report” on page 757.</p>
Policy Properties	<p>Lists all policies and policy job templates that are defined for the server.</p> <p>See “Policy Properties Report” on page 758.</p>
Policy Protected Resources	<p>Lists job information for each job that is derived from a policy and assigned to protect any part of the named resource.</p> <p>See “Policy Protected Resources” on page 759.</p>
Problem Files	<p>Lists all the problem files reported for jobs. The files are grouped by day and resource.</p> <p>See “Problem Files Report” on page 759.</p>

Table 16-16 Backup Exec Reports (*continued*)

Report Name	Description
Recently Written Media	Lists all media that have been modified in the last 24 hours. See “Recently Written Media Report” on page 760.
Resource Backup Policy Performance	Lists the success rate for policy derived backup jobs. See “Resource Backup Policy Performance Report” on page 761.
Resource Risk Assessment	Lists job information for resources on which the last backup job run on the resource failed. The data is filtered by resource server. See “Resource Risk Assessment Report” on page 761.
Resources Protected by Policy	Lists the policies, templates, and selection lists being used to protect a resource. See “Resources Protected by Policy report” on page 762.
Restore Set Details by Resource	Lists all restore sets that ran within the last 72 hours. The sets are grouped by the server and resource. See “Restore Set Details by Resource Report” on page 762.
Retrieve Media from Vault	Lists all reusable media currently in the specified vault. See “Retrieve Media from Vault Report” on page 763.
Robotic Library Inventory	Lists the contents of slots in robotic libraries attached to media servers. Usage statistics are provided for each piece of media. See “Robotic Library Inventory Report” on page 764.
Scheduled Server Workload	Lists the estimated scheduled workload for the next 24-hour period by server. See “Scheduled Server Workload” on page 765.
Scratch Media Availability	Lists the aging distribution of media. Shows how many media are available for overwrite and when other media will become available for overwrite. See “Scratch Media Availability Report” on page 766.

Table 16-16 Backup Exec Reports (*continued*)

Report Name	Description
Selection Lists	Lists a description, policy name, and job name for protected and unprotected selection lists. See “Selection Lists Report” on page 766.
Test Run Results	Lists the results for the test run jobs set for the selected time period and media servers. See “Test Run Results Report” on page 767.
Vault Store Usage Details	Displays the archives that are in each store and the size of each archive. See “Vault Store Usage Details report” on page 772.
Vault Store Usage Summary	Displays the archived items that are in each vault store and the total size of the vault store. See “Vault Store Usage Summary Report” on page 773.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Active Alerts Report

The Active Alerts report lists all active alerts chronologically, displaying the most recent alerts first. You can limit the number of alerts that appear in the report by entering range parameters for the Event count.

Information displayed in the Active Alerts report is described in the following table.

Table 16-17 Active Alerts Report

Item	Description
Time	Date and time the alert occurred.
Media Server	Name of the media server on which the alert occurred.
Job Name	Name of the job associated with the alert.
Device Name	Name of the device on which the job ran.
Category	Title of the alert, such as Service Start or Job Failed.
Message	Describes the event that caused the alert.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Active Alerts by Media Server Report

The Active Alerts by Media Server report lists all active alerts grouped and filtered by media server, displaying the most recent alerts first. You can limit the amount of data that appears in the report by selecting filter parameters for the Event count or for the Media Server option.

Information displayed in the Active Alerts by Media Server report is described in the following table.

Table 16-18 Active Alerts by Media Server Report

Item	Description
Time	Date and time the alert occurred.
Job Name	Name of the job associated with the alert.
Device Name	Name of the device on which the job ran.
Category	Title of the alert, such as Service Start or Job Failed.
Message	Describes the event that caused the alert.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Alert History Report

The Alert History report lists all the alerts in the Alert History chronologically, displaying the most recent alerts first. You can limit the number of alerts that appear in the report by entering range parameters for the Days or Event count options

Information displayed in the Active History report is described in the following table.

Table 16-19 Alert History Report

Item	Description
Time	Date and time the alert occurred.

Table 16-19 Alert History Report (*continued*)

Item	Description
Received	Time the alert occurred.
Responded	Time the user responded to the alert.
Responding User	User that responded to the alert.
Job Name	The name of the job associated with the alert.
Media Server	Name of the media server on which the alert occurred.
Category	Title of the alert, such as Service Start or Job Failed.
Message	Describes the event that caused the alert.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Alert History by Media Server Report

The Alert History by Media Server report lists all alerts in the alert history grouped and filtered by media server, displaying the most recent alerts first. You can limit the amount of data that appears in the report by selecting filter parameters for the Days, Event count or Media Server option.

Information displayed in the Alert History by Media Server report is described in the following table.

Table 16-20 Alert History by Media Server Report

Item	Description
Media Server	Name of the media server on which the alert occurred.
Time	Date and time the alert occurred.
Received	Time the alert occurred.
Responded	Time the user responded to the alert.
Responding User	User that responded to the alert.
Job Name	Name of the job associated with the alert.
Category	Title of the alert, such as Service Start or Job Failed.
Message	Describes the event that caused the alert.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Application Event Log Report

The Application Event Log report lists all Backup Exec application event logs. Information displayed in the Application Event Log report is described in the following table.

Table 16-21 Application Event Log Report

Item	Description
Number	Number assigned to the event in the Windows Event log.
Event	Type of event that occurred.
Date/Time	Date and time the event occurred.
Source	Source from which the event originated.
Description	Message displayed for the event.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Audit Log Report

The Audit Log report lists the contents of the audit logs for the selected servers for the selected time period. You can limit the amount of data that appears in the report by entering filter parameters for the Media server or Audit Category options and range parameters for the Days and Event count options.

Information displayed in the Audit Log report is described in the following table.

Table 16-22 Audit Log Report

Item	Description
Media Server	Name of the media server on which the audit logs are located.
Category	Category in which the change occurred, such as Logon Account, Alerts, or Job.
Date Entered	Time and date the change occurred.
Message	Description of the change made in Backup Exec.

Table 16-22 Audit Log Report (*continued*)

Item	Description
User Name	User that made the change.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Backup Job Success Rate Report

The Backup Job Success Rate report lists the success rate for backup jobs run to protect selected servers. You can limit the amount of data that appears in the report by entering filter parameters for the Protected server option and range parameters for the Days option.

Information displayed in the Backup Success Rate report is described in the following table.

Table 16-23 Backup Success Rate Report

Item	Description
Server	Name of the server being protected.
Date	Date the backup job was processed.
Total Jobs	Total number of jobs processed by the media server.
Successful	Total number of jobs successfully performed by the media server.
Success Rate	Percentage of successful jobs processed by the media server.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Backup Recommendations Report

The Backup Recommendations report lists any recommendations that can help you to better manage your backups. The recommendations may include better ways to back up specific types of data or suggestions for how to use Symantec products.

The following table describes the information in the Backup Recommendations report.

Table 16-24 Backup Recommendations Report

Item	Description
Media Server	Name of the media server for which the recommendation applies.
Job Name	Name of the job that is associated with the recommendation.
Start Time	Date and time when the job that is associated with the recommendation ran.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Backup Resource Success Rate Report

The Backup Resource Success Rate report lists the success rate for backup jobs for a specific number of days for resources on selected servers. You can limit the amount of data that appears in the report by entering range parameters for the Days option.

Table 16-25 Backup Resource Success Rate Report

Item	Description
Resource	Name of the system being protected.
Date	Date the backup job was processed.
Backup Sets	Total number of backup sets processed by the media server.
Successful	Total number of jobs successfully performed by the media server.
Success Rate	Percentage of successful jobs processed by the media server.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Backup Set Details by Resource Report

The Backup Set Details by Resource report lists all jobs that ran within the specified time range on a selected server. The jobs are grouped by the server and resource. You can limit the amount of data that appears in the report by entering filter

parameters for the Protected server option and range parameters for the Hours option.

Information displayed in the Daily Jobs by Resource report is described in the following table.

Table 16-26 Backup Set Details by Resource Report

Item	Description
Resource	Name of the system being protected.
Start Time	Date and time the operation started.
Duration	Length of time the operation took to process.
Size (MB)	Number of megabytes processed.
Files	Number of files processed.
Directories	Number of directories processed.
MB/Minute	Number of megabytes processed per minute.
Skipped	Number of files skipped during the operation.
Corrupt Files	Number of corrupt files encountered during the operation.
Files in Use	Number of files in use during the operation.
Status	Status of the operation, such as Completed.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Backup Sets by Media Set Report

The Backup Sets by Media Set report lists all the backup sets by media set. You can limit the amount of data that appears in the report by selecting filter parameters for the Media Set option.

Information displayed in the Backup Sets by Media Sets report is described in the following table.

Table 16-27 Backup Sets by Media Sets Report

Item	Description
Media Set	Name of the media set on which the job ran.

Table 16-27 Backup Sets by Media Sets Report (*continued*)

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned bar code label.
Set	Sequential number for backup sets on the media.
Method	Specific type of backup. See “How to choose a backup strategy” on page 270.
Date/Time	Date and time the data was backed up.
Backup Set Description / Source	Describes the data that was backed up and the location of the data.
Directories	Number of directories backed up.
Files	Number of files backed up.
MB	Amount of data backed up in megabytes.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Backup Size By Resource Report

The Backup Size By Resource report lists the backup size for each resource job for up to seven previous policy-based jobs run. It also computes the trailing average, which is the average of the amount of data backed up in the seven previous jobs.

This report only shows jobs created by applying a policy to a resource.

See [“Creating a new policy”](#) on page 522.

See [“About creating jobs using policies and selection lists”](#) on page 544.

You can limit the amount of data that appears in the report by entering filter parameters for the Protected server option.

Information displayed in the Backup Size by Resource report is described in the following table.

Table 16-28 Backup Size by Resource Job Report

Item	Description
Server	Name of the media server where the data for the backup job was located.
Resource	Name of the resource backed up.
Job	Name of the backup job.
Job Date and Time Run	Date and time the backup job was processed.
Backup Size, MB	Amount of data backed up in megabytes.
Trailing Avg, MB	Average amount of data backed up during the seven previous runs.
Difference %	Amount by which the data backed up in the current job differs from the previous backup jobs.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Configuration Settings Report

The Configuration Settings report lists the contents of the Backup Exec system configuration parameters table.

Information displayed in the Configuration Settings report is described in the following table.

Table 16-29 Configuration Settings Report

Item	Description
Parameter Name	Name of the Backup Exec configuration parameter.
Class	Parameters that are associated with the Backup Exec system.
Value	Value of the configuration parameter. Note: The StoreMaintenanceLastrun and StoreMaintenanceRuntime parameters display the date and time in Greenwich Mean Time (GMT).

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Current Job Status Report

The Current Job Status report provides details about the job queue sorted by status. You can limit the amount of data that appears in the report by selecting filter parameters for the Job status.

Information displayed in the Job Queue Status report is described in the following table.

Table 16-30 Job Queue Status Report

Item	Description
Job Status	Displays the job status.
Job	Name of the job.
Next Due Date	Next date and time the job is scheduled to run.
Original Due Date	Original date and time the job was scheduled to run.
Priority	Determines the job priority for which job runs first. If another job is scheduled to run at the same time as this job, the priority you set determines which job runs first. See “About job priority” on page 195.
On Hold	Displays an X if the job is on hold; otherwise, displays a dash (-).
Type	Type of job that was run, such as Backup or Restore.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Daily Device Utilization Report

The Daily Network Device Utilization report lists the percentage of the storage devices’ capacity that the media server uses.

Information displayed in the Daily Network Device Utilization report is described in the following table.

Table 16-31 Daily Network Device Utilization Report

Item	Description
Drive Name	Name of the storage device and the media server where the device is located.

Table 16-31 Daily Network Device Utilization Report (*continued*)

Item	Description
Date	Date the storage device was used.
Jobs	Number of jobs processed by the media server's storage device.
Size (MB)	Number of megabytes processed by the media server's storage device.
Utilization (%)	Percentage of device utilization.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Deduplication device summary

The Deduplication device summary report displays a summary of the deduplication operations for local deduplication storage folders and shared deduplication storage folders.

Table 16-32 Deduplication device summary report

Item	Description
State	Device state, such as online and enabled.
Created	Date media was created.
Total Capacity (MB)	Total capacity of the deduplication storage folder.
Used Capacity (MB)	Capacity presently used by the deduplication storage folder.
Available Capacity (MB)	Remaining capacity of the the deduplication storage folder.
Percent Full	Percentage of storage space that is available in the deduplication storage folder.
Protected Bytes (MB)	Total amount of data that is selected for backup in all jobs using the device before deduplication occurs.
Deduplication Ratio	Ratio of the amount of data before deduplication to the amount of data after deduplication.

Deduplication summary

The Deduplication summary report displays a deduplication summary for all of the deduplication jobs that run on the Backup Exec media server.

Table 16-33 Deduplication summary report

Item	Description
Job Name	Name of the job.
Start Time	Time of day that Backup Exec attempted to start the job.
Duration	Length of time the operation took to process.
Size (MB)	Number of megabytes processed.
MB/Minute	Number of megabytes processed per minute.
Scanned Byte Count (MB)	Total amount of data in megabytes that is selected for backup before deduplication occurs.
Stored Byte Count (MB)	The amount of unique data is stored after deduplication occurs.
Deduplication Ratio	Ratio of the amount of data before deduplication to the amount of data after deduplication.
Status	Status of the operation, such as Completed .

Device Summary Report

The Device Summary report lists all the devices for each selected media server. You can limit the amount of data that appears in the report by selecting filter parameters for the Media Server option.

Information displayed in the Device Summary report is described in the following table.

Table 16-34 Device Summary Report

Item	Description
Server	Name of the server where the device is located.

Table 16-34 Device Summary Report (*continued*)

Item	Description
Device Name	Name of the device, such as the name of the robotic library. This field is left blank for stand-alone drives.
Drive Name	Name of the drive in the robotic library.
Vendor/Product ID	Name of the vendor of the drive, the product ID, and firmware from the SCSI Inquiry string.
SCSI Target	Address of the SCSI Card, SCSI Bus, Target Device ID, and LUN.
State	Device state, such as online.
Created	Date media was created.
Cleaned	Date last cleaning job was run on the drive.
Hours	Hours the device has been in use since the last cleaning job.
Errors	Number of errors occurring since the last cleaning job.
MB	Megabytes read and written since the last cleaning job.
Mounts	Number of mounts occurring since the last cleaning job.
Hours	Total number of hours the device has been in use.
Errors	Total number of errors occurring on the device.
MB	Total number of megabytes read and written to the device.
Mounts	Total number of mounts occurring to the device.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Device Usage by Policy

The Device Usage by Policy report lists all the policies that are targeted to specific drive selections. You can limit the amount of data that appears in the report by selecting filter parameters for the Policy Name.

Information displayed in the Device Usage by Policy report is described in the following table.

Table 16-35 Device Usage by Policy Report

Item	Description
Drive Name	Name of the storage device and the media server where the device is located.
Method	Specific type of backup. See “About backup methods” on page 274.
Policy Name	Name of the policy.
Template Name	Name of the job template.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Error-Handling Rules Report

The Error-Handling Rules report lists all error-handling rules and provides details about each rule. You can limit the amount of data that appears in the report by selecting filter parameters for the Media Server.

Information displayed in the Error-Handling Rules report is described in the following table.

Table 16-36 Error Handling Rules Report

Item	Description
Rule Name	Name of the Error-Handling rule.
Notes	Information entered in the Notes section when the error-handling rule was created.
Job Status	Final job status that activates the rule. Possible statuses are as follows: <ul style="list-style-type: none"> ■ Error ■ Canceled

Table 16-36 Error Handling Rules Report (*continued*)

Item	Description
Error Category	Category of error for which the rule will be applied. Available error categories include the following: <ul style="list-style-type: none"> ■ Device ■ Job ■ Media ■ Network ■ Other ■ Resource ■ Security ■ Server ■ System
Enabled	Displays if the rule is enabled or disabled.
Cancel Job	Displays an X if this option is selected for the error-handling rule. The option cancels all jobs after the maximum number of retries have been attempted.
Pause Job	Displays an X if this option is selected for the error-handling rule. The option enables Backup Exec to pause the job until you can manually clear the error.
Retry Job	Displays an X if this option is selected for the error-handling rule. The option enables Backup Exec to retry the job.
Maximum Retries	Number of times the job is to be retried.
Retry Interval (minutes)	Number of minutes the job is to wait before being retried.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Event Recipients Report

The Event Recipient report lists events registered by each notification recipient. Information displayed in the Event Recipient report is described in the following table.

Table 16-37 Event Recipient Report

Item	Description
Recipient Type	Type of recipient, such as Person, Net Send, Printer, or Group.
Recipient Name	Name of the recipient.
Event Type	Alert category or ad hoc job.
Event Name	Detail for the alert category or ad hoc job.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Failed Backup Jobs Report

The Failed Backup Jobs report lists all the failed backup jobs associated with a policy. The jobs are sorted by the server and specified time frame. You can limit the amount of data that appears in the report by entering filter parameters for the Protected server option and range parameters for the Days option.

Information displayed in the Failed Backup Jobs report is described in the following table:

Table 16-38 Failed Jobs Report

Item	Description
Resource	Name of the system being protected.
Start Time	Date and time the backup job started.
Duration	Length of time the operation took to process.
Job Name	Name of job that failed.
Category	Category for the failed job that may be generated by a system, job, media, or device error.
Error Code	Displays the error code that corresponds to the failure.
Description	Describes the event that caused the error.
Status	Status of the operation, such as Completed.
Device Name	Name of the device on which the job ran.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Job Distribution by Device Report

The Job Distribution by Device report lists the jobs that have been run on each system device during the specified period. It helps determine the device’s job workload. You can limit the amount of data that appears in the report by selecting range parameters for the Days option.

Information displayed in the Job Distribution by Device report is described in the following table.

Table 16-39 Job Distribution by Device Report

Item	Description
Device	Name of the device on which the job ran.
Job Date and Time Run	Date and time the job was processed.
Job	Name of the job that ran on the device.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Jobs Summary Report

The Jobs Summary report lists all jobs that have run within the specified time range. The jobs are listed in chronological order. You can limit the amount of data that appears in the report by selecting range parameters for the Hours option.

Information displayed in the Jobs Summary report is described in the following table.

Table 16-40 Jobs Summary Report

Item	Description
Start Time	Date and time the operation started.
Job Name	Name of the completed job.
Duration	Length of time the operation took to process.
Size (MB)	Number of megabytes processed.
Files	Number of files processed.

Table 16-40 Jobs Summary Report (*continued*)

Item	Description
Directories	Number of directories processed.
MB/Minute	Number of megabytes processed per minute.
Skipped	Number of files skipped during the operation.
Corrupt Files	Number of corrupt files encountered during the operation.
Files in Use	Number of files in use during the operation.
Status	Status of the operation, such as Completed.
Type	Specific type of backup. See “About backup methods” on page 274.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Machines Backed Up Report

The Machines Backed Up report lists all the protected servers and the times they were backed up. You can limit the amount of data that appears in the report by selecting range parameters for the Days option.

Information displayed in the Machines Backed Up report is described in the following table.

Table 16-41 Machines Backed Up Report

Item	Description
Server	Name of the server that was backed up.
Total Backup Count	Total number of backups performed.
Last Backup	Date of the last backup job for this server.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Managed Media Servers Report

The Managed Media Servers report lists status and configuration information for all media servers managed by Backup Exec. You can limit the amount of data that appears in the report by selecting filter parameters for the Media Server option. Information displayed in the Managed Media Servers report is described in the following table.

Table 16-42 Managed Media Server Report

Item	Description
Managed Media Server	Name of the managed media server.
Status	Status of the server. Possible status includes the following: <ul style="list-style-type: none"> ■ Online - available for use. ■ Stalled - not responding immediately to messages ■ No Comm - communications to the server have been lost for some period of time.
Stalled	Time limit used for determining Stalled communications status.
No Comm	Time limit used for determining No Comm communications status.
Catalog Location	Location where server keeps catalog information. Possible locations are as follows: <ul style="list-style-type: none"> ■ Local - the catalog information is kept on the media server itself. ■ CASO - the catalog information is kept on the Central Admin Server.
Logs	When job logs are uploaded from the managed server to the CASO database. Possible upload times are as follows: <ul style="list-style-type: none"> ■ timed basis in seconds ■ schedule time ■ completion of job ■ never

Table 16-42 Managed Media Server Report (*continued*)

Item	Description
History	When job history is uploaded from the managed server to the CASO database. Possible upload times are as follows: <ul style="list-style-type: none"> ■ timed basis in seconds ■ schedule time ■ completion of job ■ never
Status	When status is uploaded from the managed server to the CASO database. Possible upload times are as follows: <ul style="list-style-type: none"> ■ timed basis in seconds ■ schedule time ■ completion of job ■ never
Display Alert	Displays Yes if you have configured an alert to be set if time between server clocks exceed a preset value (maximum time difference tolerance).
Sec	Maximum time difference tolerance in seconds set for server.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Media Audit Report

The Media Audit report lists the recent configuration changes that you made to your media. You can use filter parameters for the Media Server option to limit the amount of data that appears in the report. You can also enter range parameters for the Days or Event count options.

Information displayed in the Media Audit report is described in the following table.

Table 16-43 Media Audit Report

Item	Description
Date Entered	Time and date the change occurred.

Table 16-43 Media Audit Report (*continued*)

Item	Description
Message	Description of the change that was made to the media.
User Name	User that made the change.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Media Errors Report

The Media Errors report lists the number of errors that occur on all media. You can use filter parameters for the Media Set option to limit the amount of data that appears in the report. You can also enter range parameters for the Event count options.

Information displayed in the Media Audit report is described in the following table.

Table 16-44 Media Errors Report

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned bar code label.
Total Mounts	Total number of times this media has been mounted.
Total In Use Hours	Total number of hours that this media has been in use.
Total Errors	Total number of system, job, media, and device error alerts.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Media Required for Recovery Report

The Media Required for Recovery report lists the media that contain the backup sets for each system backed up on the selected media server for the specified time period. However, this report may be inaccurate if media overwrite settings allow the media to be overwritten. You can limit the amount of data that appears in the report by selecting filter parameters for the Protected server option and range parameters for the Days option.

Information displayed in the Media Required for Recovery report is described in the following table.

Table 16-45 Media Required for Recovery Report

Item	Description
Resource	Name of the system that was backed up.
Type	Specific type of backup. See “About backup methods” on page 274.
Date	Date and time the backup job set was created.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned bar code label.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Media Set Report

The Media Set report lists all media sets and media used by Backup Exec servers. Usage statistics are given for each piece of media. You can limit the amount of data that appears in the report by selecting filter parameters for the Media set option.

Information displayed in the Media Set report is described in the following table.

Table 16-46 Media Set Report

Item	Description
Media Set	Name of the media set.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned bar code label.
Media Type	Type of media cartridge, such as 4mm.
Allocated	Date media was allocated to a media set as a result of an overwrite operation.
Modified	Date data was last written to the media.
Location	Location of the media.
Hours	Total number of hours that this media has been in use.

Table 16-46 Media Set Report (*continued*)

Item	Description
Mounts	Total number of times this media has been mounted.
Soft Errors	Number of recoverable read errors encountered.
Hard Errors	Number of unrecoverable read errors encountered.
Write MB	Number of bytes that have been written to this media.
Current MB	Estimate of the number of megabytes currently on this media.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Media Vault Contents Report

The Media Vault Contents report lists all the media in a specified media vault. You can limit the amount of data that appears in the report by selecting filter parameters for the Vault option.

Information displayed in the Media Vault Contents report is described in the following table.

Table 16-47 Media Vault Contents Report

Item	Description
Vault Name	Location of the media.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned bar code label.
Overwrite Protection End Date	Date that data on the media may be overwritten.
Vault Media Rule Move Date	Date media can be moved to vault.
Media Set	Name of media set to which the media belongs.
Vault Media Rule Name	Name of vault media rule.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Missed Availability Report

The Missed Availability report lists all jobs that have missed scheduled availability windows within the specified time range. The jobs are listed in chronological order. You can limit the amount of data that appears in the report by selecting range parameters for the Hours option.

Information displayed in the Missed Availability report is described in the following table.

Table 16-48 Missed Availability Report

Item	Description
Date	Date and time the job was created.
Attempted Start Time	Time Backup Exec attempted to start the job.
Job Name	Name of the job.
Selection List	Name of the selection list for the job.
Begin Time	Date and time for beginning of availability window.
End Time	Date and time for end of availability window.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Move Media to Vault Report

Lists all media that you can move to a media vault.

The media listed are not currently in a media vault and meets one of the following criteria:

- The media has met or exceeded the vault move date specified for the media containing the media.
- The append period has expired, but the overwrite protection period is still current (allocated).

You can limit the amount of data that appears in the report by entering filter parameters for Media Server and range parameters for the Days option.

Information displayed in the Move Media to Vault report is described in the following table.

Table 16-49 Move Media to Vault Report

Item	Description
Media Server	Name of the media server where the data for the backup job was located.
Media Set	Name of the media set.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned bar code label.
Location	Location of the media.
Append Period End Date	Last date that data may be added to the media.
Overwrite Protection End Date	Date that data on the media may be overwritten.
Vault Media Rule Move Date	Date media can be moved to vault.
Vault Name	Name of vault to which media is to be moved.
Vault Media Rule Name	Name of vault media rule.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Operations Overview Report

The Operations Overview report lists details for past and future Backup Exec operations. You can limit the amount of data that appears in the report by entering range parameters for the Days or Event count options.

Information displayed in the Operations Overview report is described in the following table.

Table 16-50 Operations Overview

Item	Description
Job summary for jobs completed in the past x Hours	Details Backup Exec job activity for the specified time period.
Errors	Total number of system, job, media, and device error alerts.
Warnings	Total number of job, media, and device warning alerts.
Information	Total number of system, job, media, and device information alerts.
Attention Required	Total number of alerts that require a response from the user.
Completed (Failed)	Total number of jobs that failed.
Completed (Canceled)	Total number of canceled jobs.
Completed (Success)	Total number of jobs that completed successfully.
Exceptions	Total number of jobs that completed successfully, but may contain one or more skipped files, corrupt files, virus infected files or files in use.
Total Data Backed Up	Total amount of data backed up in MB.
Total Media Used	Total number of media used to back up the completed jobs.
Missed	Total number of missed jobs.
Recovered	Total number of recovered jobs.
Active Jobs	Total number of active jobs.
Scheduled Jobs	Total number of scheduled jobs.
Jobs On Hold	Total number of jobs on hold.
Job Status	The status of the jobs.
Scratch Media	Total number of scratch media available.
Recyclable	Total number of recyclable media available.

Table 16-50 Operations Overview (*continued*)

Item	Description
Imported	Number of imported media (media created by a product other than this installation of Backup Exec).
Allocated	Number of allocated media (media belonging to a user media set).
Total Overwritable Media	Total number of overwritable media available.
Total Appendable Media	Total number of appendable media available.
Media Overwrite Protection Level	Displays level of overwrite protection (Full, Partial, None) assigned to the media.
Online Devices	Total number of online devices.
Offline Devices	Total number of offline devices.
Disabled Devices	Total number of disabled devices.
Paused Devices	Total number of paused devices.
Disabled	Lists the name of the devices that are disabled.
Paused	Name of the paused devices.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Overnight Summary Report

The Overnight Summary report lists the results of backup jobs for each resource during the last 24 hours. This report includes backup jobs that were due to run but did not run. Jobs are given a grace period of 24 hours before being marked as past due. You can limit the amount of data that appears in the report by entering filter parameters for the Protected server option.

Information displayed in the Overnight Summary report is described in the following table.

Table 16-51 Overnight Summary Report

Item	Description
Resource	System being protected.
Type	Specific type of backup. See “About backup methods” on page 274.
Start time	Date and time the operation started.
Status	Status of the operation.
Error Category	Category for the job that may be generated by a system, job, media, or device error.
Media Server	Name of the media server on which the job ran.
Device Name	Name of the device on which the job ran.
Total Tasks	Total number of jobs run within the last 24 hours.
Uncorrected Exceptions	Number of the jobs that fail and were not run again with successful completion. Some of the archive jobs that ran in the past 24 hour period encountered exceptions. You must resolve the exceptions. Otherwise, the jobs that fail because of the exceptions continue to appear during subsequent 24 hour periods until the exceptions are resolved.
Service Level	Percentage of jobs that ran successfully.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Policy Jobs by Resource Summary Report

The Policy Jobs by Resource Summary report lists all of the backup sets that were created within a selected period. The jobs are grouped by target server and resource. You can limit the amount of data that appears in the report by selecting filter parameters for the Protected Server and range parameters for the Hours option.

Information displayed in the Policy Jobs by Resource Summary report is described in the following table.

Table 16-52 Policy Jobs by Resource Summary Report

Item	Description
Policy	Name of the policy.
Start Time	Date and time the operation started.
Duration	Length of time the operation took to process.
Size (MB)	Number of megabytes processed.
Files	Number of files processed.
Directories	Number of directories processed.
MB/Minute	Number of megabytes processed per minute.
Skipped	Number of files skipped during the operation.
Corrupt Files	Number of corrupt files encountered during the operation.
Files in Use	Number of files in use during the operation.
Status	Status of the operation, such as Completed.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Policy Jobs Summary Report

The Policy Jobs Summary report lists all jobs created from selected policies that have run within a specified time range. The jobs are listed in chronological order. You can limit the amount of data that appears in the report by selecting filter parameters for the Policy Name and range parameters for the Hours option.

Information displayed in the Policy Jobs Summary report is described in the following table.

Table 16-53 Policy Jobs Summary Report

Item	Description
Policy	Name of the policy.
Start Time	Date and time the operation started.
Job Name	Name of the completed job.
Duration	Length of time the operation took to process.

Table 16-53 Policy Jobs Summary Report (*continued*)

Item	Description
Size (MB)	Number of megabytes processed.
Files	Number of files processed.
Directories	Number of directories processed.
MB/Minute	Number of megabytes processed per minute.
Skipped	Number of files skipped during the operation.
Corrupt Files	Number of corrupt files encountered during the operation.
Files in Use	Number of files in use during the operation.
Status	Status of the operation, such as Completed.
Type	Specific type of backup. See “About backup methods” on page 274.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Policy Properties Report

The Policy Properties report lists all policies and policy job templates that are defined for the media server.

Information displayed in the Policy Definitions report is described in the following table.

Table 16-54 Policy Definitions Report

Item	Description
Template Name	Name of the job template.
Set Description	Describes the data that was backed up and the location of the data.
Method	Specific type of backup. See “About backup methods” on page 274.
Type	Type of job that will run, such as Backup.
Device	Name of the device on which the job will run.

Table 16-54 Policy Definitions Report (*continued*)

Item	Description
Media Set	Name of the media set on which the job will run.
Overwrite/Append	The media overwrite protection option configured for the backup job template properties.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Policy Protected Resources

The Policy Protected Resources report lists job information for each job derived from a policy and assigned to protect any part of a named resource. You can limit the amount of data that appears in the report by selecting filter parameters for the Protected server option.

Information displayed in the Policy Protected Resources report is described in the following table.

Table 16-55 Policy Protected Resources Report

Item	Description
Resource	System being protected.
Policy	Name of policy.
Job Name	Name of the job.
Next Due Date	Next date and time the job is scheduled to run.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Problem Files Report

The Problem Files report lists all the problem files reported for jobs. The files are grouped by day and resource. You can limit the amount of data that appears in the report by selecting filter parameters for the Protected server option and range parameters for the Days option.

Information displayed in the Problem Files report is described in the following table.

Table 16-56 Problem Files Report

Item	Description
Date	Date the problem file was encountered.
Resource	System on which the problem file is located.
Time	Time the problem file was encountered.
Reason	Error code listed in the job log summary.
File Name	Name of the problem file.
Type	Specific type of file that caused the problem.
Media Server	Name of the server on which the file is located.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Recently Written Media Report

The Recently Written Media report lists all the media that has been modified within the specified period. You can limit the amount of data that appears in the report by selecting range parameters for the Hours option.

Information displayed in the Recently Written Media report is described in the following table.

Table 16-57 Recently Written Media

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned bar code label.
Location	Location of the media, such as the storage vault name or drive name.
Set	Name of backup set.
Date and Time Modified	Date and time media was last modified.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Resource Backup Policy Performance Report

The Resource Backup Policy Performance report lists the success rate for policy backup jobs. You can limit the amount of data that appears in the report by selecting filter parameters for the Protected Server option and range parameters for the Days option.

Information displayed in the Resource Backup Policy Performance report is described in the following table.

Table 16-58 Resource Backup Policy Performance

Item	Description
Policy	Name of policy.
Resource	Name of system being protected.
Date	Date job completed.
Backup Sets	Total number of backup sets processed by the media server.
Successful	Total number of jobs successfully performed by the media server.
Success Rate	Percentage of successful jobs processed by the media server.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Resource Risk Assessment Report

The Resource Risk Assessment report shows job information for resources on which the last backup job that was run on the resource failed. You can limit the amount of data that appears in the report by selecting filter parameters for the Protected server option.

Information displayed in the Resource Risk Assessment report is described in the following table.

Table 16-59 Resource Risk Assessment Report

Item	Description
Resource	System on which the job ran.
Error Text	Describes the event that caused the job to fail.
Start Time	Time the operation started.

Table 16-59 Resource Risk Assessment Report (*continued*)

Item	Description
Job	Name of the job that failed.
Error Category	The category for the failed job that may be generated by a system, job, media, or device error.
Media Server	Name of the media server on which the job ran.
Device Name	Name of the device on which the job ran.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Resources Protected by Policy report

The Resources Protected by Policy report lists the policies, templates, and selection lists being used to protect a resource.

Information displayed in the Resources Protected by Policy report is described in the following table:

Table 16-60 Resources Protected by Policy Report

Item	Description
Policy	Name of the policy.
Resource	The resource that is being protected.
Template Name	The name of the job template contained in the policy that is applied to the resource being protected.
Selection List	The list of resources selected for protection
Job	Name of the job.
Next Due Date	Next date and time that the job is scheduled to run.

Restore Set Details by Resource Report

The Restore Set Details by Resource report lists all restore jobs that ran within the specified time range on a selected server. The jobs are grouped by the server and resource. You can limit the amount of data that appears in the report by entering filter parameters for the Protected server option and range parameters for the Hours option.

Information displayed in the Daily Jobs by Resource report is described in the following table.

Table 16-61 Backup Set Details by Resource Report

Item	Description
Resource	Name of the system being protected.
Start Time	Date and time the operation started.
Duration	Length of time the operation took to process.
Size (MB)	Number of megabytes processed.
Files	Number of files processed.
Directories	Number of directories processed.
MB/Minute	Number of megabytes processed per minute.
Skipped	Number of files skipped during the operation.
Corrupt Files	Number of corrupt files encountered during the operation.
Files in Use	Number of files in use during the operation.
Status	Status of the operation, such as Completed.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Retrieve Media from Vault Report

The Retrieve Media from Vault report lists all reusable media currently in a specified media vault. You can limit the amount of data that appears in the report by selecting filter parameters for the Vault option.

Information displayed in the Retrieve Media from Vault report is described in the following table.

Table 16-62 Retrieve Media from Vault Report

Item	Description
Vault Name	Name of the vault where the media is located.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned bar code label.

Table 16-62 Retrieve Media from Vault Report (*continued*)

Item	Description
Overwrite Protection End Date	Date that data on the media may be overwritten.
Move Date	Date media can be moved to vault.
Media Set	Name of the media set.
Vault Media Rule Name	Name of vault media rule.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Robotic Library Inventory Report

The Robotic Library Inventory report lists the contents of slots in robotic libraries attached to media servers. Usage statistics are provided for each piece of media. You can limit the amount of data that appears in the report by selecting filter parameters for the Media Server option.

Information displayed in the Robotic Library Inventory report is described in the following table.

Table 16-63 Robotic Library Inventory Report

Item	Description
Server	Name of the server where the robotic library is located.
Device Name	Name of the robotic library.
Slot	Sequential number of the slot in the robotic library.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned bar code label.
State	State of operation of the slot: paused, disabled, enabled, offline, or online.
Modified	Date the media in the slot was last accessed.
Write MB	Number of bytes that have been written to this media.

Table 16-63 Robotic Library Inventory Report (*continued*)

Item	Description
Full	Space available on a media; "1" indicates that media is full and "0" indicates that there is space available on the media.
Hours	Total number of hours this media has been in use.
Mounts	Total number of times this media has been mounted.
Append	The time remaining in the media's append period.

See ["Running a report"](#) on page 691.

See ["Running a new report job"](#) on page 694.

Scheduled Server Workload

The Scheduled Server Workload report displays the estimated scheduled workload for a server during the next 24-hour period or a user-defined time period. The report only displays recurring jobs that have already run at least one time, not jobs scheduled to run once. You can use filter parameters for the Media Server option to limit the amount of data that appears in the report. You can also enter range parameters for the Hours option.

Information displayed in the Scheduled Server Workload report is described in the following table.

Table 16-64 Scheduled Server Workload Report

Item	Description
Media Server	Name of the media server that will process the scheduled jobs.
Job	Name of the job scheduled to run.
Next Due Date	Time and day the next job is scheduled to run.
Backup Size, MB	Estimated amount of data in megabytes to be processed during the next 24 hours.
Total Size (MB)	Total amount of data to be processed on the server during the next 24 hours.
Total Size (MB)	Total amount of data to be processed on all media servers.

See ["Running a report"](#) on page 691.

See ["Running a new report job"](#) on page 694.

Scratch Media Availability Report

The Scratch Media Availability report shows the aging distribution of media, how many media are available for overwrite, and when other media will become available for overwrite. You can limit the amount of data that appears in the report by selecting range parameters for the Days option.

Information displayed in the Scratch Media Availability report is described in the following table.

Table 16-65 Scratch Media Availability Report

Item	Description
Category	The media set period configured in media set properties.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned bar code label.
Size (GB)	Capacity of the scratch media available to which data can be written.
Available to Append (GB)	Capacity of scratch media available for append.
Group Total	Total number and capacity of scratch media available to the system.
Media Total	Total number of scratch media available.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Selection Lists Report

The Selection Lists report lists information about protected and unprotected selection lists.

Information displayed in the Resources Protected by Policy report is described in the following table.

Table 16-66 Selection Lists Report

Item	Description
Selection List Name	Name of the selection list.
Selection List Description	Description of the protected selection list.

Table 16-66 Selection Lists Report (*continued*)

Item	Description
Policy Name	Name of the policy. An unprotected selection list does not have a job associated with it.
Job Name	Name of the job. An unprotected selection list does not have a job associated with it.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Test Run Results Report

The Test Run Results report displays the results for the test run jobs set for the selected period and the selected media servers. You can limit the amount of data that appears in the report by selecting filter parameters for the Media Server option and range parameters for the Hours option.

Information displayed in the Test Run Results report is described in the following table.

Table 16-67 Test Run Results

Item	Description
Media Server	Name of the media server on which the job ran.
Job Date and Time Run	Date and time the backup job was processed.
Job Name	Name of the test run job.
Backup Sets	Name of the backup set.
Credential Check	Indicates if the Backup Exec logon account was verified as correct for the resources being backed up.
Backup Size, MB	Size in megabytes of the backup.
Media Type	Type of media used, such as 4mm.
Device Name	Name of the device, such as the name of the robotic library.
Max Needed	Amount of space needed on the media to run the job.
Online	Capacity of media available in the device to which data can be appended.

Table 16-67 Test Run Results (*continued*)

Item	Description
Media Total	Total amount of appendable media available to the system.
Online	Capacity of media available in the device to which data can be written.
Media Total	Total amount of overwritable media available to the system.

See [“Running a report”](#) on page 691.

See [“Running a new report job”](#) on page 694.

Archive Job Success Rate report

The Archive Job Success Rate report displays the number of archive jobs for the protected servers that successfully ran..

Table 16-68 Archive Job Success Rate

Item	Description
Date	Displays the date on which archive jobs ran.
Total Jobs	Displays the total number of archive jobs that have run.
Successful	Displays the total number of archive jobs that are successful.
Success Rate	Displays the success rate of the archive jobs in percentages.

Archive Selections by Archive Rules and Retention Categories report

The Archive Selections by Archive Rules and Retention Categories report displays the archive rules and the retention categories that are applied to each archive selection.

Table 16-69 Archive Selections by Archive Rules and Retention Categories

Item	Description
Archive Rule	Displays the archive rule that you specify to identify the files and mail messages that are eligible for archiving.

Table 16-69 Archive Selections by Archive Rules and Retention Categories
(continued)

Item	Description
Archive Selection	Displays only the archive selection. In the case of an NTFS archive, the network path appears. In the case of an Exchange mailbox archive, the mailbox group appears along with information about the mailbox group selections.
Archive Type	Displays the type of data that you are archiving. Archive types include: <ul style="list-style-type: none"> ■ File System Archive ■ Mailbox Archive
Windows Domain	Displays the Windows domain in which the archived selection resides.
Retention Category	Displays the retention category that applies to the file system selections in the archive job. A retention category specifies the time period for which you want to keep archived items.

Exchange Mailbox Group Archive Settings report

The Exchange Mailbox Group Archive Settings report displays the archive settings that are applied to mailbox groups in each domain.

Table 16-70 Exchange Mailbox Group Archive Settings

Item	Description
Windows Domain	Displays the name of the Windows domain in which the Exchange server belongs.
Mailbox Group	Displays the name of the mailbox group to be archived.
Archive Rules	Displays the archive rule that is used to archive the mailbox group.

Table 16-70 Exchange Mailbox Group Archive Settings *(continued)*

Item	Description
Retention Category	<p>Displays the retention category that applies to the mailbox group selections in the archive job.</p> <p>A retention category specifies the time period for which you want to keep archived items.</p>

Failed Archive Jobs report

The Failed Archive Jobs report displays what archive jobs failed recently.

Table 16-71 Failed Archive Jobs

Item	Description
Start Time	Displays the time that the archive job started.
Duration	Displays the amount of time that the archive job took to run.
Job Name	Displays the name of the archive job.
Category	Displays status of the failed archive job.
Error Code	Displays the error code for the error that caused the archive job to fail.
Description	Displays the description of the error that caused the archive job to fail.
Status	Displays the category for the error that may be generated due to system, job, media, or device issues
Device Name	Displays the name of the storage device that processed the archive job.

File System Archive Settings report

The File System Archive Settings report displays the archive settings that are applied to archive selections for each server.

Table 16-72 NTFS Archive Settings

Item	Description
Server	Displays the name of the Windows server from where the data was archived.
Resource	Displays the path of the resource.
Archive Rules	Displays the archive rule that is used to archive the files.
Vault Store	Displays the name of the vault store where the archived files reside.
Retention Category	<p>Displays the retention category that applies to the file selections in the archive job.</p> <p>A retention category specifies the time period for which you want to keep archived items.</p>

Overnight Archive Summary report

The Overnight Archive Summary report displays the status of the archive jobs that ran in the last 24 hours.

Table 16-73 Overnight Archive Summary

Item	Description
Resource	Displays the name of the server that you are protecting.
Type	Displays the type of job that ran in the last 24 hours.
Start Time	Displays date and the time that the archive operation started.
Status	Displays the status of the archive operation.
Error Category	Displays the category for the error that may be generated due to system, job, media, or device issues.
Media Server	Displays the name of the media server on which the job ran.

Table 16-73 Overnight Archive Summary (*continued*)

Item	Description
Device Name	Displays the name of the device on which the job ran.
Total Tasks	Displays the total number of archive jobs that have run during the preceding 24 hours.
Uncorrected Exceptions	Displays the number of archive jobs that failed because the error condition was never corrected and were not run again with successful completion.
Service Level	Displays the percentage of jobs that ran successfully.

Vault Store Usage Details report

The Vault Store Usage Details report displays the archives that are in each store and the size of each archive.

Table 16-74 Vault Store Usage Details

Item	Description
Vault Store	Displays the name of the vault store where the Backup Exec archives are stored.
Archive Name	Displays the name that the Archiving Option gives to the archive.
Archive Type	Displays the type of data that you are archiving. Archive types include: <ul style="list-style-type: none"> ■ File System Archive ■ Mailbox Archive
Number of Archived Items	Displays number of archived items that are in the vault store.
Total Size (in KB)	Displays the total size of the archived items in the vault store.

Vault Store Usage Summary Report

The Vault Store Usage Summary report displays the archived items that are in each vault store and the total size of the vault store.

Table 16-75 Vault Store Usage Summary

Item	Description
Vault Store	Displays the name of the disk-based vault store where the Backup Exec archives are stored.
Database Name	Displays the name of the vault store database that contains the configuration data and information about each of the archives in the partition.
Vault Store Open Partition	Displays the name of the vault store open partition where the Backup Exec archives are stored.
Vault Store Partition Free Size(in KB)	Displays the amount of available free space in a vault store open partition.
Number of Archives in Vault Store	Displays the total number of existing Backup Exec archives in the vault store.
Total Size (in KB)	Displays the total size in kilobytes of the existing Backup Exec vault store archives.

Disaster preparation and recovery

This chapter includes the following topics:

- [About disaster preparation](#)
- [About key elements of a disaster preparation plan \(DPP\)](#)
- [Returning to the last known good configuration](#)
- [Creating a hardware profile copy](#)
- [About creating an emergency repair disk \(Windows 2000 computers only\)](#)
- [About manual disaster recovery of Windows computers](#)
- [About a manual disaster recovery of a local Windows computer \(includes non-authoritative and authoritative restore of Active Directory for a domain controller\)](#)
- [About a disaster recovery operation of a remote Windows computer \(includes non-authoritative and authoritative restore of Active Directory for a domain controller\)](#)

About disaster preparation

Disaster preparation planning is the implementation of strategies and procedures that will minimize damage in the event a catastrophe destroys your data. While precautions can be taken to minimize the effects of this type of occurrence (UPS devices, password protection, and so forth), unfortunately there is nothing that can safeguard your data 100 percent.

The purpose of a Disaster Preparation Plan (DPP) is to return to an operational status as quickly as possible. Backup Exec is a crucial component of the DPP and this section discusses how to apply this powerful data management tool to your DPP.

The following basic methods are available for disaster recovery:

- Manual recovery. You can manually recover both local and remote Windows computers.
- Automated recovery. Backup Exec's Intelligent Disaster Recovery (IDR) option automates the disaster recovery process for Windows computers.

See [“About the the Intelligent Disaster Recovery Configuration Wizard”](#) on page 1810.

See [“Returning to the last known good configuration”](#) on page 777.

See [“About manual disaster recovery of Windows computers”](#) on page 780.

About key elements of a disaster preparation plan (DPP)

The DPP you put in place with your Backup Exec system should be tailored to your network environment.

While environments vary in different organizations, consider the following elements when creating a comprehensive DPP.

Table 17-1 Key elements of a DPP

Element	Description
Hardware protection	The hardware devices on your network (CPUs, drives, video) are susceptible to damage from many disaster situations. Uninterruptible power supplies (UPS), surge protectors, and security monitoring devices are the equipment most often used today to protect hardware. If you do not already have these items in place, you should consider installing them. The initial investment could be justified many times over in the event of a disaster.

Table 17-1 Key elements of a DPP (*continued*)

Element	Description
The ability to maintain business operations during a disaster period	Make sure that proper precautions are taken by everyone to implement plans for network interruptions. For example, the phones in the sales department won't stop ringing because the server is down, so orders may have to be handwritten until the server is up again. Each department should work out strategies for such occurrences. If the proper precautions are taken, the server can be rebuilt quickly and operations can still continue.
A sound backup strategy.	A well-designed backup strategy that includes a strong media rotation scheme plays a key role in quickly restoring your file server.
Off-site storage of backups.	It is imperative that backed up data be moved off-site regularly. This ensures that if something happens to your facility, all of your backups will not be destroyed. Depending on the importance of your data, you may choose to use several off-site storage facilities. There are companies that provide off-site storage services that pick up and deliver tapes when they are to be rotated.
Effective DPP management	The last element - and possibly the most important - is proper management of your DPP strategy. A person or group of people should be charged with constantly supervising your organization's disaster preparation efforts. Someone should install and maintain hardware protection devices, make sure all departments have a plan if the server goes down temporarily, and make sure that backups are made and rotated off-site regularly. Also, it is a good idea to document your Disaster Preparation Plan for reference purposes.

Backup Exec plays a major role in your DPP by offering an easy, reliable way of backing up and restoring your files. The rest of this chapter describes how to take some precautionary measures to make restoration as straightforward as possible in the event of a disaster.

See [“About selecting data to back up”](#) on page 280.

Returning to the last known good configuration

Changes to the system configuration may keep the system from booting. If you suspect that boot problems are the result of a configuration change, you may be

able to correct the problem by returning to a previous configuration. This method is simple and fast, and in some cases will correct boot problems in a Windows computer. There are slightly different procedures for Windows operating systems. This section includes procedures for each type of computers.

Any changes made to the system since the last time the configuration was saved are lost.

See [“Creating a hardware profile copy”](#) on page 778.

See [“About creating an emergency repair disk \(Windows 2000 computers only\)”](#) on page 779.

To return to a previous configuration

- 1 Restart the system.
- 2 Press <F8> during startup.
- 3 Select one of the following options:

Safe Mode	This option allows you to diagnose and fix system startup problems. For more information, see your Microsoft documentation.
Last Known Good Configuration	This option allows you to return to a previous saved configuration.

Creating a hardware profile copy

Before making a major hardware change, copy the current hardware profile to a new hardware profile and boot into the new profile before adding or changing the hardware. This way, you can return to the previous configuration if something does not work properly.

See [“Returning to the last known good configuration”](#) on page 777.

See [“About creating an emergency repair disk \(Windows 2000 computers only\)”](#) on page 779.

To create a copy of the current hardware profile and make that the preferred boot option

- 1 Right-click the **My Computer** icon.
- 2 Click **Properties** to display the **System Properties** dialog box.
- 3 Click **Hardware**.
- 4 Click **Hardware Profiles**.

- 5 Select the current hardware profile, and then click **Copy**.
- 6 Type the name for the new configuration in the **To** field, and then click **OK**.
- 7 To make the new profile the preferred boot option, select it, and then click the up arrow next to the list box to move the new hardware profile to the top of the box.
- 8 Choose whether Windows is to use the new hardware profile automatically (after a delay) during startup, or if the system should wait indefinitely until the hardware profile is chosen by selecting the appropriate option.
- 9 Click **OK**.

About creating an emergency repair disk (Windows 2000 computers only)

When Windows 2000 Server is installed, the installation program prompts you to create an Emergency Repair Disk (ERD). This disk contains system information that can help get the system running in the event of a disaster. It is important to keep the ERD updated whenever system changes are made. The ERD is only useful if it is kept current.

Whenever a major change is made to the system, make a fresh copy of the ERD before and after the change is made. Major changes include adding, removing, or otherwise modifying hard drives or partitions, file systems, configurations, and so forth. As a general rule, update the ERD before and after the hard drive configuration is changed. The addition of a new component to the server, such as Microsoft Exchange Server or Microsoft SQL Server, and changes from Control Panel, are also situations in which the ERD should be refreshed both before and after the change.

Also remember to make a backup of the ERD; always keep an ERD from at least one generation back. When creating a fresh ERD, use a floppy disk that can be reformatted, because RDISK.EXE, the program that creates the ERD, always formats the floppy disk.

Note: The Emergency Repair Disk is a useful and necessary tool; it is NOT a bootable disk. There is not enough space on the disk for the boot files and the repair information files.

Note: You must not change or delete the systemroot\repair folder because the repair process relies on information saved in this folder.

To create an emergency disk, see your Microsoft documentation.

See [“Returning to the last known good configuration”](#) on page 777.

See [“Creating a hardware profile copy”](#) on page 778.

About manual disaster recovery of Windows computers

If your system is not protected by Backup Exec 2010 Intelligent Disaster Recovery (IDR), you can manually recover a computer.

See [“Running a disaster recovery operation on a remote Windows computer \(includes non-authoritative and authoritative restore of Active Directory for a domain controller\)”](#) on page 786.

If your system is protected by IDR, you should use automated disaster recovery.

See [“About the Intelligent Disaster Recovery Option”](#) on page 1806.

The manual disaster recovery procedures restore your computer’s operating system to its pre-disaster state and restore your data files, except those protected by one of the Backup Exec agents.

You should perform manual disaster recovery in the following situations:

- The Windows operating system has become corrupted and cannot be restored using the Emergency Repair Disks.
- The hard drive containing the Windows operating system has encountered an unrecoverable error that requires reformatting the disk.
- The hard drive that contains the Windows operating system needs to be replaced.

About a manual disaster recovery of a local Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)

This procedure restores your computer’s operating system to a pre-disaster state. It also restores your data files, except for those that are protected by one of the Backup Exec database agents, such as the Exchange Agent or SQL Agent. If any of your data is protected by Backup Exec agents, refer to the section on restoring the data protected by the agent before beginning disaster recovery.

If your system is protected by Backup Exec 2010 Intelligent Disaster Recovery (IDR), you should use IDR for disaster recovery.

See [“About the Intelligent Disaster Recovery Option”](#) on page 1806.

The procedure described in the following section allows you to manually recover a computer not protected by IDR.

A media drive must be attached to the computer that is being recovered.

You will also need the following items:

- A current full backup of the computer to be recovered and any subsequent incremental/differential backups.
- The Windows installation media.
- The Backup Exec installation media.

Note: If you recover a Windows computer that has BitLocker encryption enabled, you must re-enable BitLocker encryption following the restore.

See Microsoft's documentation for more information on BitLocker drive encryption.

See [“Running a manual disaster recovery of a local Windows computer \(includes non-authoritative and authoritative restore of Active Directory for a domain controller\)”](#) on page 781.

See [“Restoring data by setting job properties”](#) on page 605.

See [“Running a disaster recovery operation on a remote Windows computer \(includes non-authoritative and authoritative restore of Active Directory for a domain controller\)”](#) on page 786.

See [“About manual disaster recovery of Windows computers”](#) on page 780.

Running a manual disaster recovery of a local Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)

Use the following steps to manually recover a local Windows computer, which includes non-authoritative and authoritative restore of Active Directory for a domain controller.

To run a manual disaster recovery of a local Windows computer, which includes non-authoritative and authoritative restore of Active Directory for a domain controller

- 1 Install the original version of Windows.

About a manual disaster recovery of a local Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)

This basic Windows installation is necessary to provide Backup Exec with a target to which it can restore the system. The computer name, Windows directory, and the file system (such as NTFS) must be the same as the previous Windows installation. This installation will be overwritten by the backed up version, which will restore your original system configuration, application settings, and security settings.

If you are recovering from an entire hard disk failure, use Windows setup to partition and format the new disk during installation.

Format the partitions with the same file system as before the failure, as follows:

- If the system was in a specific domain or workgroup, do not join the domain or workgroup at this time.
- If you are recovering a domain controller, do not perform the domain controller installation process at this time.

- 2 Install Backup Exec to a directory other than where it was originally installed (this is a temporary installation).

Always log on to Windows using the Administrator account or its equivalent during this procedure.

- 3 Using the Device Configuration Wizard, install the appropriate device driver for the attached media drive.
- 4 Start Backup Exec.
- 5 From the navigation bar, click **Devices**.
- 6 Inventory the media containing the latest full backup of the computer to be recovered.

See [“About inventorying media”](#) on page 448.

- 7 Catalog the media containing the latest full backup of the computer to be recovered. If the subsequent differential/incremental backups are on separate media, catalog those also.

See [“Creating a new catalog”](#) on page 248.

- 8 From the navigation bar, click **Restore**.
- 9 Select all sets from the full and incremental backups that contain logical drives on the hard disk. If differential backup sets are to be restored, select only the last differential set. Make sure you include System State and Shadow Copy components as part of the restore selections.
- 10 On the Properties pane, under **Settings**, click **General**, and then select the following options:

- Restore over existing files
 - Restore security
 - Preserve tree
- 11 On the Properties pane, under **Settings**, click **Advanced**, and then select the appropriate options.
- See [“Advanced options for restore jobs”](#) on page 613.
- If you are restoring a computer that is the only domain controller in the domain or the entire domain is being rebuilt and this is the first domain controller, select the option **Mark this server** as the primary arbitrator for replication when restoring folders managed by the File Replication Service, or when restoring SYSVOL in System State.
- 12 Click **Run Now**.
- 13 If you are restoring a computer that is the only domain controller in the domain or the entire domain is being rebuilt and this is the first domain controller, reboot the computer after the restore job successfully completes.
- Your computer’s operating system is now restored to a pre-disaster state. Your data files have been restored, except those protected by Backup Exec database agents.
- 14 Continue with one of the following:
- If you are performing an authoritative restore go to step 15.
- If you are not performing an authoritative restore the recovery is complete.
- 15 Do the following to change the Backup Exec services to the local system account.
- Right-click My Computer and then select **Manage**.
 - From the left pane of the Computer Management utility, double-click **Services and Applications**.
 - Click **Services**.
 - In the right pane, double-click each Backup Exec service, and from the Log On tab, change Log on as to use Local System account.
 - Close the Computer Management utility.
- 16 Restart the computer.

About a manual disaster recovery of a local Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)

17 Press **F8** during startup.

A menu appears that allows you to diagnose and fix system startup problems.

18 Select **Directory Services Restore Mode**.

19 Launch Backup Exec.

20 From the navigation bar, click **Restore**.

21 Select System State (Windows 2000 and later) or Shadow Copy (Windows Server 2003 and later) components as the restore selections. Run the Restore job.

22 At this point, you can either choose to restore the entire Active Directory, or specific objects from the Active Directory.

Restore the entire Active Directory by performing the following:

- Open a command prompt.
- Type NTDSUTIL and press **Enter**.
- Type Authoritative Restore and press **Enter**.
- Type Restore Database, press **Enter**, click **OK** and then click **Yes**.

See Microsoft's documentation for running NTDSUTIL for Windows Server 2008/2008 R2.

Restore specific objects from the Active Directory by performing the following:

- Open a command prompt.
- Type NTDSUTIL and press **Enter**.
- Type Authoritative Restore and press **Enter**.
- Type Restore Subtree "ou=<OU Name>.dc=<domain name>.dc=<xxx>" (without the quotation marks), and then press **Enter**, where <OU Name> is the name of the organizational unit you want to restore, <domain name> is the domain name the OU resides in, and <xxx> is the top level domain name of the domain controller, such as com, org, or net. You can do this as many times for as many objects you need to restore.

23 Once you have finished restoring Active Directory information, exit NTDSUTIL.

24 Restart the computer.

About a disaster recovery operation of a remote Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)

You can perform a disaster recovery on a remote computer that is attached to the media server. This procedure restores your computer's operating system to its pre-disaster state. It also restores your data files, except those that you protect with a Backup Exec agent.

If any of your data is protected by Backup Exec agents, review the overview of the agents before you begin disaster recovery.

See [“Backup Exec agents and options”](#) on page 80.

If your system is protected by Backup Exec Intelligent Disaster Recovery (IDR), you should use IDR for disaster recovery.

See [“About the Intelligent Disaster Recovery Option”](#) on page 1806.

The procedure described in the following section allows you to manually recover a computer not protected by IDR.

You will need the following:

- A current full backup of the computer to be recovered and any subsequent incremental/differential backups.
- The Windows installation media.

Always log on to Windows using the Administrator account or its equivalent during this procedure.

Note: If you recover a Windows computer that has BitLocker encryption enabled, you must re-enable BitLocker encryption following the restore.

See Microsoft's documentation for more information on BitLocker drive encryption.

See [“Restoring data by setting job properties”](#) on page 605.

See [“About manual disaster recovery of Windows computers”](#) on page 780.

Running a disaster recovery operation on a remote Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)

Use the following steps to run a disaster recovery operation on a remote Windows computer.

To run a disaster recovery operation on a remote Windows computer

- 1 At the remote computer, install the original version of Windows.

This basic Windows installation is necessary to provide Backup Exec with a target to which it can restore the system. The computer name, Windows directory and the file system (such as NTFS) must be the same as the previous Windows installation. This basic installation will later be overwritten by the backed up version, which will restore your system configuration, application settings, and security settings.

If you are recovering from an entire hard disk failure, use Windows setup to partition and format the new disk during installation.

Format the partitions with the same file system as before the failure, as follows:

- If the system was in a specific domain or workgroup, do not join the domain or workgroup at this time.
- If you are recovering a domain controller, do not perform the domain controller installation process at this time.

- 2 At the media server, install the Backup Exec Remote Agent to the remote computer.

See [“About installing the Remote Agent for Windows Systems”](#) on page 139.

- 3 Start Backup Exec.

- 4 From the navigation bar, click **Devices**, and then inventory the media containing the latest full backup of the computer to be recovered.

See [“About inventorying media”](#) on page 448.

- 5 Catalog the media containing the latest full backup of the computer to be recovered. If the subsequent differential/incremental backups are on separate media, catalog those also

See [“Creating a new catalog”](#) on page 248.

About a disaster recovery operation of a remote Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)

- 6 From the navigation bar, click **Restore**, and then select all sets from the full and incremental backups that contain logical drives on the hard disk. If differential backup sets are to be restored, select only the last differential set. Make sure you include System State or Shadow Copy components as part of the restore selections.
- 7 On the Properties pane, under Settings, click **General**, and then select the following options:
 - **Restore over existing files**
 - **Restore security**
 - **Preserve tree**
- 8 On the Properties pane, under **Settings**, click **Advanced**, and then select the appropriate options.

See [“Advanced options for restore jobs”](#) on page 613.

If you are restoring a computer that is the only domain controller in the domain or the entire domain is being rebuilt and this is the first domain controller, select the option Mark this server as the primary arbitrator for replication when restoring folders managed by the File Replication Service, or when restoring SYSVOL in System State.
- 9 Click **Run Now**.
- 10 After the job completes, restart the remote computer.

Your computer’s operating system is now restored to its pre-disaster state. Your data files have been restored, except those protected by Backup Exec database agents.
- 11 Continue with one of the following:
 - If you are performing an authoritative restore go to step 12.
 - If you are not performing an authoritative restore the recovery is complete.
- 12 At the remote server, press **F8** during startup.

A menu appears that allows you to diagnose and fix system startup problems.
- 13 Select **Directory Services Restore Mode**.
- 14 At the media server, start Backup Exec.
- 15 From the navigation bar, click **Restore**.

About a disaster recovery operation of a remote Windows computer (includes non-authoritative and authoritative restore of Active Directory for a domain controller)

- 16 Select System State (Windows 2000 and later) or Shadow Copy (Windows 2003 and later) components as the restore selections.
- 17 From the Properties pane, under **Source**, select **Resource Credentials**.
- 18 Highlight the restore selection for the remote server and click **New**.
- 19 Create a new logon account for this restore job. The account should have administrator privileges on the remote server.
- 20 Select the new logon account and click **OK**.

- 21 Run the Restore job.

At the remote server:

- 22 At this point, you can either choose to restore the entire Active Directory, or specific objects from the Active Directory:

Restore the entire Active Directory by performing the following:

- Open a command prompt.
- Type NTDSUTIL and press **Enter**.
- Type Authoritative Restore and press **Enter**.
- Type Restore Database, press **Enter**, click **OK** and then click **Yes**.

See Microsoft's documentation for running NTDSUTIL on Windows Server 2008/2008 R2.

Restore specific objects from the Active Directory by performing the following:

- Open a command prompt.
- Type NTDSUTIL and press **Enter**.
- Type Authoritative Restore and press **Enter**.
- Type Restore Subtree "ou=<OU Name>.dc=<domain name>,dc=<xxx>" (without the quotation marks), and then press **Enter**, where <OU Name> is the name of the organizational unit you want to restore, <domain name> is the domain name the OU resides in, and <xxx> is the top level domain name of the domain controller, such as com, org, or net. You can do this as many times for as many objects you need to restore.

- 23 Once you have finished restoring Active Directory information, exit NTDSUTIL.
- 24 Restart the computer.

Troubleshooting

This chapter includes the following topics:

- [Troubleshooting hardware-related issues](#)
- [How to get more information about alerts and error messages](#)
- [Troubleshooting backup issues](#)
- [About cluster sizes for NTFS partitions](#)
- [Troubleshooting restore issues](#)
- [How to improve Backup Exec's performance](#)
- [About the Symantec Knowledge Base](#)
- [How to contact Technical Support](#)
- [About Backup Exec diagnostic tools](#)

Troubleshooting hardware-related issues

If you have trouble with your hardware, review the following questions.

Table 18-1 Hardware-related questions

Question	Answer
<p>My drive is not listed in Backup Exec's Devices list. The drive is connected, powered on and recognized in the Windows Device Manager. What should I do?</p>	<p>First, make sure that your devices are supported by Backup Exec. You can find a list of compatible devices at the following URL: http://entsupport.symantec.com/umi/V-269-2</p> <p>If your device is listed on the Hardware Compatibility List, try running Backup Exec's Device Configuration wizard and installing Symantec device drivers.</p> <p>See "About configuring tape devices by using the Tape Device Configuration Wizard" on page 453.</p> <p>The Symantec Device Driver Installation wizard will find and install the most suitable driver for your storage device.</p> <p>Note: The Hardware Compatibility List is frequently updated with newly supported devices.</p>

Table 18-1 Hardware-related questions (*continued*)

Question	Answer
<p>My drive appears as offline. Why?</p>	<p>If the device is offline, this message is displayed. No operations are allowed on the device until it is online again. When the device is online, no message is displayed.</p> <p>Backup-to-Disk folders may go offline in the following situations:</p> <ul style="list-style-type: none"> ■ The drive containing the Backup-to-disk folder is full. ■ The drive containing the backup-to-disk folder is offline. ■ The remote server containing the Backup-to-disk folder is offline. <p>Other storage devices may go offline in the following situations:</p> <ul style="list-style-type: none"> ■ The device was turned off after Backup Exec was started. ■ The device was being used by another application (such as a Windows 2000/XP/Server 2003/2008 backup utility) when Backup Exec was started. ■ The device is removed from the computer. ■ A tape drive failure occurred (check the Event Log to troubleshoot the problem). ■ A tape is stuck in the drive. ■ The firmware of the drive was updated; Backup Exec will behave as if the drive with its old name or identity no longer exists. <p>To place the device online, try the following:</p> <ul style="list-style-type: none"> ■ Check to make sure the device has power and that cables are properly attached. Turn the device on and reboot the server, or stop and restart the Backup Exec services. ■ Stop the utility that is using the device, and then restart the server, or stop and restart the Backup Exec services. You can restart services from Tools > Backup Exec Services. <p>If the drive's firmware has changed, delete the drive and restart Backup Exec services. After the drive appears with its new firmware identity, retarget all jobs that were using the old drive name to the new drive name.</p>

Table 18-1 Hardware-related questions (*continued*)

Question	Answer
<p>I set up bar code rules through the Tools menu by selecting Options, and then selecting Bar Code Rules. However, my bar code rules don't seem to be working. Why?</p>	<p>After setting up bar code rules, you must perform the following two steps in order for the bar code rules to work.</p> <ul style="list-style-type: none"> ■ You must enable the bar code rules for the robotic library by selecting the bar code rules option on the Configuration tab in the robotic library's properties. ■ In addition to setting the bar code rules for each type of media you use, for each drive in your mixed media library you should indicate what type of media can be used and whether that media can be used for read or write operations. <p>See "Bar code rules in mixed media libraries" on page 245.</p>
<p>How do I get the latest device drivers for my hardware?</p>	<p>You can find a list of compatible devices at the following URL: http://entsupport.symantec.com/umi/V-269-2</p>
<p>Backup Exec doesn't detect my robotic library. What could be wrong?</p>	<p>Be sure that Windows operating system properly recognizes the device. This can be verified by checking the Windows Device Manager.</p> <p>See "About robotic libraries in Backup Exec" on page 467.</p>

Table 18-1 Hardware-related questions (*continued*)

Question	Answer
<p>I'm getting an error "Storage device [device] reported an error on a request to read/write data to/from media. Error reported: Data error (cyclic redundancy check)." What should I do?</p>	<p>The cyclic redundancy check (CRC) error can be caused by many factors.</p> <p>The following list contains the most common reasons for this error and potential ways to resolve the problem:</p> <ul style="list-style-type: none"> ■ Contaminated read/write heads of the tape device. Check with the hardware manufacturer for proper cleaning techniques. ■ Bad media. Replace the media. Try a new tape that is certified by the hardware manufacturer. ■ Tape driver. Load the appropriate Backup Exec tape driver. ■ You can find a list of compatible devices at the following URL: http://entsupport.symantec.com/umi/V-269-2 ■ SCSI controller wide negotiation not configured properly. If the device is a wide (68 pin) SCSI device, then wide negotiation may and should be used. If the device is a narrow (50 pin) SCSI device, disable wide negotiation. Use the manufacturer's SCSI setup program to disable wide negotiation on the SCSI controller card. ■ SCSI controller transfer rate is too fast. Use the manufacturer's SCSI setup program to lower the SCSI transfer rate. Check with the controller and backup device manufacturer for the proper configuration for the SCSI transfer rate. ■ SCSI controller synchronous negotiation enabled. Use the manufacturer's SCSI setup program to disable synchronous negotiation on the SCSI controller card. Check with the controller and backup device manufacturer for the proper configuration for SCSI synchronous negotiation. ■ Incorrect termination or bad cables. Verify that the SCSI cable is good and that it is configured to provide proper SCSI termination. Do not mix passive and active termination. ■ Confirm that the tape drive is functioning properly. Check with the tape drive manufacturer for diagnostic software to test the condition of the tape drive hardware. ■ General SCSI problems. Isolate the tape drive on its own controller card or try a different SCSI card.

Table 18-1 Hardware-related questions (*continued*)

Question	Answer
Why does my DLT tape drive pause when cataloging some tapes?	<p>The DLT tape drive maintains internal information about the tape on a tape directory track. The directory track is updated before the tape is ejected from the drive. If the drive is powered off without ejecting the tape first, this information is lost.</p> <p>Re-generating the tape directory information takes several hours to complete, which makes it seem like the drive is hung. Allow sufficient time for the operation to complete and then eject the tape. Normal operation will resume after the directory track has been updated.</p>
A backup to my DLT tape drive is stuck at 99% complete. What should I do?	<p>The backup most likely fails to complete because the Eject media after job completes option is selected on tape drives that require you to manually remove the tape (such as Digital Linear Tape (DLT), Linear Tape-Open (LTO), Travan, and Onstream drives).</p> <p>To remedy this situation, either deselect the Eject media... option or using BEUTILITY, you can configure Backup Exec to set automatic responses to the media alert.</p> <p>See “About configuring tape devices by using the Tape Device Configuration Wizard” on page 453.</p> <p>See “How to get more information about alerts and error messages” on page 794.</p> <p>See “How to improve Backup Exec's performance” on page 797.</p>

How to get more information about alerts and error messages

Backup Exec generates an error message when a condition occurs that is important enough to warrant your attention, or requires that you submit a response. Most alerts and error messages are self explanatory, but there may be times when you need to get more information to resolve a condition.

You can get more information on Backup Exec alert and error messages in the following ways:

- On the alert dialog box, click the link for the Unique Message Identifier (UMI) code, or look in the job log and click the UMI link. This code is a hyperlink to the Symantec Technical Support Web site. You can access the technical notes that are related to the alert.

See [“Linking from the job log to the Symantec Technical Support Web site”](#) on page 577.

- Search the Symantec technical support knowledge base for the error. From the **Help** menu, click **Symantec on the Web**, and then select **Search Knowledge Base**.

See [“About error-handling rules”](#) on page 590.

See [“Troubleshooting hardware-related issues”](#) on page 789.

Troubleshooting backup issues

If you have problems with backing up data, review the following questions.

Table 18-2 Backup questions

Question	Answer
<p>I am unable to back up certain files on my system that are being used by other processes. Why is that?</p>	<p>When Backup Exec encounters a file that is in use by another process, it either skips the file or waits for the file to become available, depending on the Backup open files setting. When Backup Exec is configured to back up open files, it attempts to open the files in a different mode. It locks these files while they are being backed up to prevent other processes from writing to them. This mode should be a last resort to obtaining a backup of open files; in most circumstances, it is more desirable to close applications that leave files open so their files may be backed up in a consistent state.</p> <p>If you want to back up open files on Windows computers, Backup Exec’s Advanced Open File Option provides uninterrupted data protection for network environments.</p>
<p>Why do Backup Exec’s consoles continue to own a storage device even when it’s not running?</p>	<p>Backup Exec is a true client/server application that must always be available to process jobs submitted from both local and remote administrative consoles.</p> <p>Because of the Advanced Device and Media Management functionality, all storage devices attached to the media server are claimed by Backup Exec whenever the server is running. The Advanced Device and Media Management feature in Backup Exec requires constant control of the storage devices in order to perform two important and useful operations: collection of statistics on media and device usage, and media overwrite protection.</p>

Table 18-2 Backup questions (*continued*)

Question	Answer
When performing a local backup, the total number of bytes backed up by Backup Exec does not match the number of bytes displayed by Windows. Why?	<p>This problem may be caused by the type of partition for which the system is formatted.</p> <p>If you have a Windows NTFS compressed partition, Backup Exec displays the uncompressed byte count of the files being backed up while Windows Explorer displays the compressed byte count of the files on the hard drive. For example, a NTFS partition that contains 1 GB of data is compressed by Windows to 500 MB. Backup Exec reports that 1 GB of data was backed up, even though Windows Explorer displays that only 500 MB of compressed data exists on the hard drive.</p> <p>If you have a FAT partition, Backup Exec reports the actual number of bytes of the files being backed up while File Manager reports an inflated amount of disk space. For example, a 2 GB FAT partition has a 32 K cluster size and File Manager displays 1.9 GB of used space. Backup Exec reports that 1.4 GB of data was backed up. Assuming that a 50 MB pagefile.sys is excluded from the backup, there is a 450 MB difference in the number of bytes.</p> <p>Converting to NTFS will regain disk space since it is more efficient and the default cluster size (automatically set by Windows) in NTFS is less than FAT. Windows allows you to specify a cluster size other than the default; however system performance may decrease. For more information, see the Windows documentation.</p> <p>See “About cluster sizes for NTFS partitions” on page 796.</p>

About cluster sizes for NTFS partitions

The following table displays the cluster sizes for NTFS partitions.

Table 18-3 NTFS Partition Cluster Sizes

Partition Size (MB)	Cluster Size
<= 255	512
256 - 511	1024
512 - 1023	2048
1024 - 2047	4096

The following table displays the cluster sizes for FAT partitions.

Table 18-4 FAT Partition Cluster Sizes

Partition Size (MB)	Cluster Size (K)
<= 127	2
128 - 255	4
256 - 511	8
512 - 1023	16
512 - 1023	32

Troubleshooting restore issues

Sometimes the byte counts for data that you restore don't match the byte counts indicated when the data was initially backed up. When data backed up from an NTFS volume is restored to an NTFS volume, the byte count will match between the backup and restore operations. However, when data backed up from a NTFS or FAT volume is restored to a FAT volume, the byte count restored is expected to be less than that backed up. The reason for the discrepancy is that the Windows returns a default ACL (access control list) for FAT data; the stream of data is backed up (and the bytes are counted) but is discarded during a restore (and the bytes are not counted).

See [“Restoring data by setting job properties”](#) on page 605.

See [“Troubleshooting hardware-related issues”](#) on page 789.

How to improve Backup Exec's performance

The following variables can affect throughput performance:

Table 18-5 Variables that affect throughput performance

Item	Description
Hardware	<p>The speed of the disk controller and hardware errors caused by the disk drive, the tape drive, the disk controller, the SCSI bus, or the improper cabling/termination can slow performance.</p> <p>Confirm that the controller is rated for the tape backup hardware and that the SCSI Bios Settings are set properly. Newer models of SCSI Controllers are set to communicate with SCSI Hard Drives by default. Most tape drives can only handle a maximum sync transfer rate (bus speed) of between 3 to 22 MB/sec when utilizing hardware compression. Speed in excess of this will not only affect the ability for data to write to the tape in a continuous stream, but can also potentially damage the tape hardware.</p> <p>In addition, you should also ensure the following:</p> <ul style="list-style-type: none"> ■ Enable disconnect and enable Sync Negotiation is set to NO (in most cases). ■ Initiate Wide Negotiation is set to Yes when the tape device is connected to a 68 pin wide SCSI Cable Connector. ■ Tape drives are not connected to a SCSI Raid Controller.
System	<p>The capacity and speed of the media server performing the backup, or the remote system being backed up significantly impacts performance. System activity during backup also impacts performance.</p> <p>Fragmented disks take a longer time to back up. Heavily fragmented hard disks not only affect the rate at which data is written to tape, but also affect the overall system performance. Fragmented files take longer to back up because each segment of data is located at a different location on the disk, which causes the disk to take longer to access the data. Make sure you defragment disks on a regular basis.</p>

Table 18-5 Variables that affect throughput performance (*continued*)

Item	Description
Memory	<p>The amount of available memory will impact backup speed. Insufficient memory, improper page file settings, and a lack of available free hard disk space will cause excessive paging and slow performance. See “System requirements” on page 114.</p>
File Types	<p>The average file can potentially compress at a 2:1 ratio when hardware compression is used. Higher and lower compression occur depending on the type of files being backed up. Average compression can double the backup speed, while no compression runs the tape device at its rated speed.</p> <p>Image and picture files are fully compressed on disks. Therefore, no hardware compression takes place during the backup causing the tape drive to operate at its native (non-compression) rate of speed. Hardware compression is performed by the tape device and not the backup software.</p>
Compression	<p>Successful compression can increase the tape drive's data transfer rate up to twice the native rate. Some tape drives use the Lempel-Ziv (LZ1) compression algorithm for its superior versatility and efficiency. Compression can be highly variable depending on your input data. Compression algorithms look for repeatable data patterns that can be compacted.</p> <p>Image files from a graphical program like Microsoft Paint, may compress at 4.5:1 or more, while binary files may compress at just 1.5:1. Data that has already been compressed or random data (such as encrypted data or MPEG files) may actually expand by about five percent if you attempt to compress it further. This can reduce drive throughput.</p>
Files	<p>The total number of files on a disk and the relative size of each file impacts backup performance. Fastest backups occur when the disk contains fewer large size files. Slowest backups occur when the disk contains thousands of small files. A large number of files located in the same directory path back up more efficiently than backing them up from multiple directory locations.</p>

Table 18-5 Variables that affect throughput performance (*continued*)

Item	Description
Block Size	<p>Larger block sizes improve the compression ratio, which helps the drive to achieve better throughput and more tape capacity. Make sure that the block and buffer size are set properly. The throughput will increase in proportion to the compression achieved, until the drive's maximum throughput is reached.</p> <p>Some devices (for example, DLT devices) provide better performance when larger block sizes are used. The preferred block size can range from 512 bytes to 64 kilobytes or larger. If you use a device that supports larger block sizes, you can change the device's block size in the Device Configuration tab. However, if the option to change the block size is unavailable, you must configure the device to use a larger size.</p> <p>See the device manufacturer's documentation for help to configure the device.</p>
Network	<p>The backup speed for a remote disk is limited by the speed of the physical connection.</p> <p>The rate at which a remote server's hard disks are able to be backed up depends on the following:</p> <ul style="list-style-type: none"> ■ The make/model of network cards. ■ The network card driver. ■ The mode/frame type configuration for the adapter. ■ The connectivity equipment (hubs, switches, routers, and so on). ■ Windows Settings. <p>Local disk drives on the media server can usually be backed up at a higher rate of speed than backing up remote servers across a network.</p>

Table 18-5 Variables that affect throughput performance (*continued*)

Item	Description
Hardware	<p>The speed of the disk controller and hardware errors caused by the disk drive, the tape drive, the disk controller, the SCSI bus, or the improper cabling/termination can slow performance.</p> <p>Confirm that the controller is rated for the tape backup hardware and that the SCSI Bios Settings are set properly. Newer models of SCSI Controllers are set to communicate with SCSI Hard Drives by default. Most tape drives can only handle a maximum sync transfer rate (bus speed) of between 3 to 22 MB/sec when utilizing hardware compression. Speed in excess of this will not only affect the ability for data to write to the tape in a continuous stream, but can also potentially damage the tape hardware.</p> <p>In addition, you should also ensure the following:</p> <ul style="list-style-type: none"> ■ Enable disconnect and enable Sync Negotiation is set to NO (in most cases). ■ Initiate Wide Negotiation is set to Yes when the tape device is connected to a 68 pin wide SCSI Cable Connector. ■ Tape drives are not connected to a SCSI Raid Controller.

See “[Creating a backup job by setting job properties](#)” on page 332.

About the Symantec Knowledge Base

The Symantec Knowledge Base is a centralized location where you can locate more information about your Symantec products. The Knowledge Base contains information about how to install, upgrade, configure, and use your products. It also contains information about requirements, best practices, and how to troubleshoot problems. The Symantec Knowledge Base is accessible from within Backup Exec.

Note: You must have an active Internet connection to access the Symantec Knowledge Base.

The Knowledge Base uses a keyword-based search technology. It focuses on the important keywords in a search and compares them to other search phrases to provide the best possible results. You can use Boolean search features and expression queries to provide search parameters. For best results, focus on a few keywords that best represent your question.

Searching the Symantec Knowledge Base

When you search the Knowledge Base, a new browser window launches and displays the search results.

To search the Symantec Knowledge Base

- 1 Type your question or keywords in the Search Knowledge Base search box in Backup Exec's upper right corner.
- 2 Click the magnifying glass icon.

How to contact Technical Support

If you have tried everything you can to solve a problem, but still need a resolution, you can contact Technical Support via the Internet with Symantec MySupport or by phone.

You can find a list of phone numbers at the following URL:

<http://entsupport.symantec.com/phonesup>

To expedite the Technical Support process, do the following:

- Know your Backup Exec version and revision number. Locally, the version and build information can be located by selecting the About Backup Exec option from the Help menu.
- Use one of the diagnostic utilities included with Backup Exec to collect information that technical support can use to diagnose your issue.
See “[About the Backup Exec diagnostic application](#)” on page 804.
See “[How to use the Symantec Gather Utility for troubleshooting](#)” on page 808.

About Backup Exec diagnostic tools

The following diagnostic tools help to troubleshoot issues in Backup Exec:

Table 18-6 Backup Exec diagnostic tools

Item	Description
Backup Exec Support Tool	Scans the local computer and generates a report about common issues in your Backup Exec environment. See “About the Symantec Backup Exec Support Tool” on page 803.
Backup Exec diagnostic application	Gathers the pertinent information about a Windows computer for troubleshooting. See “About the Backup Exec diagnostic application” on page 804.
Symantec Gather Utility	Creates and compiles a compressed file that contains the system log files. See “How to use the Symantec Gather Utility for troubleshooting” on page 808.
Gather utility for Linux servers	Creates and compiles a Packet file. The file contains detailed information about installation, diagnostics, and error reports. See “Running the begather utility to troubleshoot Backup Exec components on Linux servers” on page 809.
Backup Exec Debug Monitor	Captures the debug output from Backup Exec and saves it into debug logs. See “Using the Backup Exec Debug Monitor for troubleshooting” on page 810.
Backup Exec debug tool	Generates the diagnostic information about the Backup Exec processes that shut down unexpectedly. See “About the Backup Exec debug tool” on page 810.

About the Symantec Backup Exec Support Tool

The Symantec Backup Exec Support Tool is a diagnostic tool that scans the local computer and generates a report about common issues in your Backup Exec environment. You can view the report and try to resolve the issues before you contact Symantec Technical Support. If you cannot resolve the issues, you can choose to collect data to send to Symantec Technical Support .

The report identifies issues with items such as the following:

- Minimum system requirements.
- Backup Exec service account.
- Backup Exec service account credentials.
- Windows Removable Storage service.

You can find a description of the report and more information about the report at the following URL:

<http://entsupport.symantec.com/umi/V-370-2378-00002>

You can find more information about the Backup Exec Support Tool at the following URL:

<http://entsupport.symantec.com/umi/V-370-2378-00001>

See “[Running the Backup Exec Support Tool](#)” on page 804.

Running the Backup Exec Support Tool

You can use the Backup Exec Support Tool to run a system analysis or collect data for technical support. When you run the Backup Exec Support Tool, it automatically checks for an updated version of the tool. If there is a new version, Backup Exec automatically downloads and runs the new version.

You can also find the latest version of the Backup Exec Support Tool at the following URL:

<http://entsupport.symantec.com/umi/V-370-2378-00003>

See “[About the Symantec Backup Exec Support Tool](#)” on page 803.

To run the Backup Exec Support Tool

- ◆ On the **Tools** menu, click **Support Utilities > Run the Support Tool to collect information before you call Technical Support**.

About the Backup Exec diagnostic application

Backup Exec includes a diagnostic application (Bediag.exe) that gathers information about a Windows computer for troubleshooting purposes. You can run it from the media server, or from a remote computer. This application can be run from within Backup Exec, or it can be run from a command line. The Bediag command line utility is located in the Backup Exec directory on your hard drive (by default, `\Program Files\Symantec\Backup Exec`).

The type of information collected in the bediag.txt file includes the following:

- Account groups, account privileges and environment settings.
- Backup Exec software version and registry information, Backup Exec Agent listing, Windows version information, SCSI hardware configuration, SQL Server information, Driver services information and Windows Services information.
- Server information, supported shared directories and Windows sockets information.

See [“Generating a diagnostic file for troubleshooting”](#) on page 805.

See [“Generating a diagnostic file on a remote media server”](#) on page 808.

See [“Using the command line to generate a diagnostic file for troubleshooting”](#) on page 806.

Generating a diagnostic file for troubleshooting

You can run the Backup Exec diagnostic application to gather information for troubleshooting. Diagnostic information appears in a text file.

To generate a diagnostic file for troubleshooting

- 1 Start Backup Exec.
- 2 On the **Tools** menu, select **Backup Exec Diagnostics**.
- 3 Select the appropriate options.
 See [“Backup Exec Diagnostics”](#) on page 805.
- 4 Click **Run Diagnostics**.
- 5 Click **Close**.

Backup Exec Diagnostics

You select a server and generate a diagnostic file to gather information for troubleshooting.

See [“Generating a diagnostic file for troubleshooting”](#) on page 805.

Table 18-7 Backup Exec Diagnostics options

Item	Description
Server	Displays the name of the media server.
User name	Indicates the user name for an account that has rights on the media server.
Password	Indicates the password for an account that has rights the media server.

Table 18-7 Backup Exec Diagnostics options (*continued*)

Item	Description
Domain	Indicates the domain in which the media server is located.
Select Server	Lets you select a different resource to run the diagnostic application.
View File	Displays the diagnostic information in a text file.
Run Diagnostics	Runs the diagnostic application to gather information for troubleshooting purposes.

Using the command line to generate a diagnostic file for troubleshooting

You can run the Backup Exec diagnostic application from the command line to gather information for troubleshooting.

To use the command line to generate a diagnostic file for troubleshooting

- 1 Launch the command prompt.
- 2 Do one of the following:

To generate a diagnostic file for a media server

From the directory `Program Files\Symantec\Backup Exec\`, type `bediag [switches] servername`.

See [“Command line switches for a diagnostic file”](#) on page 806.

To generate a diagnostic file for a remote server

From the directory `Program Files\Symantec\Backup Exec\`, type `bediag [switches] workstationname`.

See [“Command line switches for a diagnostic file”](#) on page 806.

- 3 Open the "Bediag.txt" from the directory that contains Bediag.exe (by default `Program Files\Symantec\Backup Exec`).

Command line switches for a diagnostic file

You can add the following switches to gather additional information when you generate a diagnostic file for troubleshooting.

Table 18-8 Command line switches for a diagnostic file

Switch	Description
/a	Dumps the Agent List.
/b:[server]	Specifies a Backup Exec media server to poll for service account information.
/c	Dumps the Backup Exec software configuration from the registry.
/app	Dumps the Application Event log.
/sys	Dumps the System Event log.
/bex	Dumps only Backup Exec entries in the Application Event log.
/err	Dumps only error events from any event log.
/recs:n	Dumps only newest n records from given event logs.
	***The bex, err and recs switches must be used in conjunction with the app and/or sys switches.
/o:[file]	Specifies output job log for append.
	***Omitting [file] will send output to the screen.
/h	Dumps SCSI hardware subkey from registry.
/l	Dumps Lotus Notes information.
/n	Dump Windows Socket Network Protocols.
/p	Dumps user privileges.
	Dumps Microsoft SQL Server information.
/s	Dumps information on Services.
/u	Dumps Microsoft update information.
/v	Dumps Server Information.
/w	Dumps Windows version information.
/x	Dumps Microsoft Exchange Server Information.
/?	Displays usage information.

Generating a diagnostic file on a remote media server

You can run diagnostics on a remote media server provided:

- Backup Exec is installed on the remote server.
- Backup Exec services are running.

Diagnostic information appears in a text file.

To generate a diagnostic file on a remote media server

- 1 On the **Tools** menu, click **Backup Exec Diagnostics**.
- 2 Click **Select Server** and select the remote media server on which you want to run the diagnostic utility.
- 3 Select the appropriate options.
See “[Backup Exec Diagnostics](#)” on page 805.
- 4 Click **Run Diagnostics**.
- 5 Click **Close**.

How to use the Symantec Gather Utility for troubleshooting

When troubleshooting an issue with Backup Exec, it may be necessary to review diagnostic logs from the media server. The Symantec Gather Utility simplifies this process by creating and compiling a compressed file that includes various system log files that can be sent to technical support. You can run the Symantec Gather Utility locally, or you can copy it to another computer.

The Symantec Gather Utility runs a diagnostic application called Bediag as part of its collection process. Bediag captures specific log file information. The Gather Utility also provides you with the ability to gather other data by using additional diagnostic tools.

See “[Collecting log file information for troubleshooting](#)” on page 808.

Collecting log file information for troubleshooting

You can use the Symantec Gather Utility to troubleshoot an issue in Backup Exec. After the Symantec Gather Utility collects all of the log file information, you will have the option of viewing all of the data collected by the utility and sending the results via email or FTP. The files gathered contain detailed information regarding installation, diagnostics, and error reporting. Reviewing these logs prior to contacting technical support can reveal the source of the issue. If the solution is not evident based on the gathered logs, please have these logs available when contacting support. The support technician may request an email that contains the log files.

To collect log file information for troubleshooting

- 1 On the **Tools** menu, click **Support Utilities > Run the Gather Utility to collect logs and crash dumps**.
- 2 Check all of the boxes in the **Data to gather (if available)** field.
- 3 Use the default **Output root directory (required)** or specify an alternate one.

If you want to use the default directory Go to step 4.

If you know the name of the directory Type the name of the directory.

If you do not know the name of the directory Click **Browse** to browse to the correct directory.

- 4 Enter your case number with the dashes (for example, 123-456-789).
- 5 If there are additional files you want added to the compressed file enter them now.

If you know the name of the file Type the file name in the Files text box, and then click **Add**.

If you do not know the file name

- Click **Browse** to browse to the correct file.
- Select the file and then click **Open**.
- Click **Add**.

- 6 Click **Gather**.

Running the begather utility to troubleshoot Backup Exec components on Linux servers

The begather utility brings together the files that help you diagnose issues with Backup Exec components on Linux servers. After you run it, the begather utility displays the name of the Packet file that it creates. The files that are gathered contain detailed information regarding installation, diagnostics, and error reporting. Reviewing these files before contacting technical support can reveal the source of the issue. If the solution is not evident based on the gathered files, have the Packet file available when contacting support. The support technician may request an email that contains the Packet file.

Run the begather utility to troubleshoot Backup Exec components on Linux servers

- 1 Log on as root to the Linux server on which the Backup Exec components are installed.

- 2 Navigate to the following directory:

```
/opt/VRTSralus/bin
```

For example:

```
cd /opt/VRTSralus/bin
```

- 3 Start the begather utility.

For example:

```
./begather
```

- 4 Note the location of the Packet file that is displayed on the screen.

Using the Backup Exec Debug Monitor for troubleshooting

The Backup Exec Debug Monitor, or SGMon, is a diagnostic tool that captures debug output from Backup Exec and saves it in debug logs. SGMon debug logs can help you troubleshoot backup issues. Furthermore, debug logs can help Symantec Technical Support diagnose and repair problems.

When you open SGMon, it automatically captures debug data from Backup Exec's services. To collect debug information while SGMon is closed, enable debug log creation outside of SGMon and specify a directory in which to save the logs.

For more information about how to configure the Debug Monitor and read log files, refer to the help within the Debug Monitor.

To use the Backup Exec Debug Monitor for troubleshooting

- ◆ On the **Tools** menu, click **Support Utilities > Run the Debug Monitor for active debugging**.

About the Backup Exec debug tool

Backup Exec includes a debug tool (BEDBG) that generates diagnostic information about the Backup Exec processes that shut down unexpectedly. The diagnostic information helps Symantec Technical Support diagnose and repair the problem. The Backup Exec debug tool runs by default in Backup Exec. The data that the tool gathers is copied into the BEDBG folder which is located in \Program Files\Symantec\Backup Exec.

You can find more information about the Backup Exec debug tool at the following URL:

<http://seer.entsupport.symantec.com/docs/327408.htm>

Using Symantec Backup Exec with Server Clusters

This chapter includes the following topics:

- [About Backup Exec and clusters](#)
- [Requirements for clustering Backup Exec in a Microsoft Cluster Server environment](#)
- [How Backup Exec works in a Microsoft Cluster Server environment](#)
- [Requirements for installing Backup Exec on a Microsoft Cluster Server](#)
- [Installing Backup Exec on a Microsoft Cluster Server](#)
- [Creating device pools for Microsoft Cluster Servers](#)
- [Using checkpoint restart on Microsoft Cluster Server failover](#)
- [Enabling or disabling checkpoint restart](#)
- [Specifying a different failover node](#)
- [Designating a new SAN SSO primary server and central administration server in a Microsoft Cluster Server](#)
- [Configurations for Backup Exec and Microsoft Cluster Servers](#)
- [Using the Central Admin Server Option with Microsoft clusters and SAN SSO](#)
- [About backing up Microsoft Cluster Servers](#)
- [About restoring data to a Microsoft cluster](#)
- [Using Backup Exec with Veritas Cluster Server](#)

- [Requirements for installing Backup Exec with the CASO option on a Veritas Cluster Server](#)
- [Installing Backup Exec with the CASO option on a Veritas Cluster Server](#)
- [Requirements for clustering Backup Exec using Veritas Cluster Server](#)
- [Clustering Backup Exec using Veritas Cluster Server](#)
- [About backing up Veritas Cluster Servers](#)
- [About restoring data to Veritas Cluster Servers](#)
- [About backup job failover with Veritas Cluster Servers](#)
- [Disaster recovery of a cluster](#)
- [Troubleshooting clusters](#)

About Backup Exec and clusters

In a server cluster, Backup Exec can protect data on local disks and shared disks, as well as protect Microsoft SQL and Exchange databases that are configured as virtual server applications; that is, they contain an IP address resource, a Network Name resource, and are displayed on the network with a unique server name (the virtual server name). Clustered servers provide high availability of applications and data to users. In a clustered server, several servers (called nodes) are linked in a network, and run cluster software that allows each node access to the shared disks. If a node becomes unavailable, cluster resources migrate to an available node (called failover). The shared disks and the virtual server are kept available. During failover, users experience only a short interruption in service.

Note: For offhost backups that use the hardware provider in a Microsoft Cluster Server (MSCS) or Veritas Cluster Services environment, the media server and the remote computer must be in different cluster groups. The cluster applications cannot support devices' logical unit numbers (LUNs) that have duplicate signatures and partition layouts, therefore, the snapshots containing the LUNs must be transported to a host, or remote computer, that is outside the cluster.

See [“Installing Backup Exec on a Microsoft Cluster Server”](#) on page 818.

See [“Configurations for Backup Exec and Microsoft Cluster Servers”](#) on page 827.

See [“About backing up Microsoft Cluster Servers”](#) on page 836.

See [“About restoring data to a Microsoft cluster”](#) on page 840.

See [“Requirements for installing Backup Exec on a Microsoft Cluster Server”](#) on page 816.

See [“Disaster recovery of a cluster”](#) on page 850.

See [“Installing Backup Exec with the CASO option on a Veritas Cluster Server”](#) on page 844.

See [“About backing up Veritas Cluster Servers”](#) on page 846.

See [“About restoring data to Veritas Cluster Servers”](#) on page 849.

Requirements for clustering Backup Exec in a Microsoft Cluster Server environment

The following scenarios must be followed if you plan to cluster Backup Exec:

- Symantec highly recommends that you use the default database instance (MSDE) that Backup Exec installs if you plan to cluster Backup Exec.
- Symantec also supports using a remote SQL Server instance to host the Backup Exec database. However, if you plan to use this scenario, review the following: Only one installed instance of Backup Exec can be installed into the remote SQL Server instance on a clustered node. All other installed instances of Backup Exec in the cluster must use the default Backup Exec MSDE database instance.

Note: You must run the Backup Exec cluster wizard on the cluster node that uses the remote SQL Server instance.

If you use **Windows Server 2008** or later and you use a remote, clustered SQL Server instance to host the Backup Exec Database:

- The Backup Exec media server must use the same operating system level that is installed on the computer that hosts the remote SQL Server instance.

If you use **Windows Server 2008** or later and you use the **Backup Exec Utility** to reconfigure the clustered Backup Exec installation or the clustered remote SQL Server instance:

- Run the Backup Exec Utility on a computer that uses the same operating system level at the Backup Exec media server and the computer that hosts the remote SQL Server instance.

How Backup Exec works in a Microsoft Cluster Server environment

When you install Backup Exec into a Microsoft cluster, you install it as a virtual server application. You assign an IP address resource, a Network Name resource (the virtual server name), and a disk resource to Backup Exec.

When a failover occurs, backup jobs that were running are rescheduled. The Backup Exec services are restarted on a designated failover node, and the backup jobs are restarted by default. Backup Exec provides an additional rule for cluster failover restart called Checkpoint Restart. A checkpoint restart option allows backup jobs to continue from the point at which the jobs were interrupted rather than starting the backup over again, making the backups faster and requiring fewer media. If the rule to retry jobs on a cluster failover is enabled, then an additional option can be specified to do a checkpoint restart when retrying the job. Checkpoint Restart is the only property available for the Cluster Failover Rule. You can change the default so that jobs are not restarted.

When the failed server comes back online, the Microsoft cluster can automatically rebalance the workload in a cluster, called failback, by moving cluster groups back to the server that has rejoined the cluster. However, by design, Backup Exec does not failback. The backup jobs will continue to run on the designated failover node. By continuing to run backup jobs on the designated failover node, any further risk of having to restart the jobs again when the failed server rejoins the cluster is avoided. Then, when it is convenient, you can move the Backup Exec cluster group back to the controlling node.

Specific details of how Backup Exec runs in a cluster vary depending on the configuration you use in the cluster.

See [“Configurations for Backup Exec and Microsoft Cluster Servers”](#) on page 827.

See [“Installing Backup Exec on a Microsoft Cluster Server”](#) on page 818.

Requirements for installing Backup Exec on a Microsoft Cluster Server

The following items are required to install Backup Exec on a Microsoft cluster:

- Two-node clusters are supported with Backup Exec on Microsoft Windows 2000 Advanced Server/DataCenter, Windows Server 2003 Enterprise/DataCenter, and Windows Server 2008 R2 Enterprise/DataCenter.

- Four-node clusters are supported with Backup Exec on Microsoft Windows 2000 DataCenter, Windows Server 2003 Enterprise/DataCenter, and Windows Server 2008 R2 Enterprise/DataCenter.
- Up to eight-node clusters are supported with Backup Exec on Microsoft Windows Server 2003 DataCenter.
- Backup Exec clusters can be installed on Windows Server 2003/2008 R2 majority node configurations. However, there must be a shared disk in the configuration in order for Backup Exec to share its database files between nodes. In this type of configuration, if the majority of the cluster nodes fail, then the entire cluster will fail. This configuration normally uses more than two nodes in the cluster configuration.
- The controlling node and designated failover nodes must be online during installation of Backup Exec into the cluster.
- A unique IP address and a unique network name for the Backup Exec virtual server is required during installation.
- During installation of a Backup Exec cluster, the node that runs the installation should own the shared disk. If you use a physical disk resource that belongs to another application, the Backup Exec Cluster Wizard will move all the resources that belong to the other application into the Backup Exec group. It is recommended that Backup Exec not be installed on the cluster quorum.
- An individually licensed copy of Backup Exec, as well as any applicable agents and options, is required for each active node in the cluster as defined in the End User License Agreement. When installing an evaluation version of Backup Exec, a cluster environment is automatically detected and license keys are not required.
- When you install Backup Exec clusters in a SAN SSO configuration, all Backup Exec installations must have the same server configuration. Either all nodes should be database servers or all nodes should be secondary member servers connecting to the same primary.
- All Backup Exec installations into a cluster must either be part of a single cluster group, or be locally installed on each node. If cluster-aware Backup Exec is installed in a cluster as well as a locally installed version of Backup Exec (not cluster-aware), then you cannot log on to the locally installed Backup Exec media server. You can only log on using the Backup Exec virtual server name. To be able to log on to the locally installed Backup Exec media server, you must first use the Cluster Configuration Wizard to uninstall cluster-aware Backup Exec from all the nodes in the cluster.

- Use the same account for Backup Exec services on all nodes in the cluster. If nodes in a cluster use Backup Exec and have different accounts, change the services to use the same account.

See [“Installing Backup Exec on a Microsoft Cluster Server”](#) on page 818.

See [“Configurations for Backup Exec and Microsoft Cluster Servers”](#) on page 827.

Installing Backup Exec on a Microsoft Cluster Server

Symantec does not recommend installing Backup Exec on the same disk that the cluster quorum is installed on. If you have to specify a new drive letter for the quorum disk during a recovery process, Backup Exec will not recognize the new drive and will not run.

See [“Specifying a new drive letter for the cluster quorum disk”](#) on page 841.

Note: By default, failover from the controlling node to a designated node occurs in alphabetical order according to the machine name of each node. To change the order in which failover will occur on the designated nodes, rename the machines.

The Remote Agent is automatically installed on all the nodes in the cluster. If this installation of Backup Exec will be used to back up remote servers outside the cluster, install the Remote Agent on those remote servers as well.

To install Backup Exec on a Microsoft Cluster Server

- 1 Install Backup Exec on all the nodes that you want in the cluster. Use the same installation path for each node.
- 2 From the node that you want to be the active node, start Backup Exec.
- 3 From the **Tools** menu, point to **Wizards**, and then click **Cluster Configuration Wizard**.
- 4 Follow the instructions on the screen.

On the **Virtual Server Information** screen, Backup Exec automatically displays a default name called BKUPEXECVRS for the virtual server. Type a new default name if you do not want to use the default.

- 5 When the Cluster Configuration Wizard completes, create a device pool that contains all the locally attached storage devices on each node to be used when failover occurs. This ensures that jobs can be run on the storage devices that are attached to the failover nodes.

See [“Creating device pools for Microsoft Cluster Servers”](#) on page 821.

- 6 Repeat step 5 for all nodes.

See [“Enabling or disabling checkpoint restart”](#) on page 824.

See [“Configurations for Backup Exec and Microsoft Cluster Servers”](#) on page 827.

See [“Specifying a different failover node”](#) on page 824.

Upgrading Backup Exec on a Microsoft cluster

You can upgrade Backup Exec on the nodes in a cluster without taking the nodes out of the cluster.

You can upgrade to Backup Exec 12 on nodes in a cluster that use the Desktop and Laptop Option. However, you must ensure that each node had been an active host before you begin the upgrade.

Table 19-1 Upgrading Backup Exec on a Microsoft cluster

Step	Action
Step 1	Select a node to upgrade and make that node the active Backup Exec cluster node.
Step 2	Run the Backup Exec installation program on the active node.
Step 3	Move the cluster group to the next node you want to upgrade, and then run the Backup Exec installation program on that node. All of the resources except for the disk should be offline when moved over to each node for upgrade.
Step 4	Repeat step 3 for each node in the cluster.

See [“About upgrading from previous versions of Backup Exec”](#) on page 180.

Installing additional Backup Exec options on a Microsoft cluster

Install additional Backup Exec options on each node of the cluster. For details on installing each option, see the appropriate section in this guide, or in online Help.

Note: If you are using the Advanced Open File Option, set those defaults on each physical node the option is installed on, not on the virtual server. Because the default settings for the static volume can be different on each node, Advanced Open File Option defaults do not fail over.

To install additional Backup Exec options

- 1 On the controlling node, make sure the Backup Exec group is online before you start installing additional options.
- 2 Install the additional options.
See [“Installing additional Backup Exec options to the local media server”](#) on page 123.
- 3 After the installation is complete on the controlling node, use the cluster administrator to move the Backup Exec group to the next appropriate node, and repeat step 2.

Be sure to install the same options with the same settings for each node in the cluster.
- 4 To install the Backup Exec Agent for Oracle Windows or Linux Servers and the Backup Exec Agent for SAP Applications on other nodes, map a drive to the shared disks where Backup Exec is installed on the cluster, and run SETUP.

Uninstalling Backup Exec from a Microsoft cluster

You use the Cluster Configuration Wizard to remove Backup Exec.

To uninstall Backup Exec from a cluster

- 1 From the **Tools** menu, point to **Wizards**, and then click **Cluster Configuration Wizard**.
- 2 Use the wizard to remove cluster-aware Backup Exec from all selected servers.

When unclustering the active node, you can either leave the Backup Exec data on the shared drive or delete it. If you delete the data, you can make the data available on the active node.
- 3 Uninstall Backup Exec from all the nodes.

- 4 After Backup Exec has been uninstalled, move any resource disks from the Backup Exec cluster group to another group, and then delete the Backup Exec cluster group.
- 5 On any node, click **Start**, point to **Settings**, and then click **Control Panel** to uninstall Backup Exec.
- 6 Double-click **Add/Remove Programs**, and then in the list of currently installed programs, select **Symantec Backup Exec (TM) 2010** and click **Change/Remove**.
- 7 Repeat step 5 for all nodes.

Creating device pools for Microsoft Cluster Servers

When Backup Exec is installed on a cluster, it creates default device pools named **All Devices (<Node Name>)** for each node in the cluster. If a node has storage devices, those storage devices are automatically assigned to **All Devices (<Node Name>)**, which is also the default destination device on that node when you create backup or restore jobs. However, to allow jobs to run on the storage devices attached to a failover node after a failover occurs, you must create a device pool that includes the storage devices from all of the nodes. If the cluster is also configured with tape devices on a shared SCSI bus, then add the tape device name used by each node to the device pool. You must also select this device pool as the destination device for all jobs that you want to be restarted.

You can create either a single device pool, or you can create device pools for device or media types so that when jobs fail over they can be restarted on "like" devices and media.

To create a device pool for a cluster

- 1 From the controlling node, open Backup Exec.
- 2 Create a new device pool.
See [“Creating device pools”](#) on page 516.
- 3 Add storage devices and then exit Backup Exec. If there are tape devices on a shared SCSI bus, then add the tape device name used by each node.
See [“Adding devices to a device pool”](#) on page 517.
- 4 Using the cluster administrator, move the Backup Exec resource group to the next appropriate node.
- 5 Open Backup Exec, add storage devices for this node to the previous device pool and then exit Backup Exec. If there are tape devices on a shared SCSI bus, then add the tape device name used by each node.
- 6 Repeat step 4 and step 5 for each node in the cluster.

Using checkpoint restart on Microsoft Cluster Server failover

You can enable or disable checkpoint restart for each backup job run on the cluster (by default, checkpoint restart is enabled). When checkpoint restart is enabled, jobs that were interrupted because of a failover continue from the point of interruption rather than starting over. Files that were already backed up are skipped, and only the remaining files in the job are backed up when the job is restarted. If this option is not selected, jobs are restarted from the beginning.

Checkpoint restart works best for the following file types:

- NTFS
- Exchange mailboxes and public folders
- Exchange 2003 IS with multiple storage groups
- SQL database non-snapshot backups

The following types of files cannot use checkpoint restart:

- System State
- Lotus Domino
- Exchange 2003 IS with one storage group
- NTFS Image sets
- NTFS Snapped volumes
- SQL database snapshot backups
- SQL transaction log backups
- NetWare SMS (the checkpoint restart option should be disabled for NetWare backups using the Remote Agent)

Checkpoint restart is not supported by the following:

- The Advanced Open File Option.
- Microsoft Windows Vista/Server 2008.
- The offhost backup feature in the Advanced Disk-based Backup Option.
- When the option Collect additional information for synthetic backups is selected for the synthetic backup feature in the Advanced Disk-based Backup Option.
- Incremental backups based on the archive bit.

Jobs that are restarted from the point of failover display a status of 'Resumed' in the Job Monitor.

Before using checkpoint restart, review the following:

- If a resource was completely backed up prior to a cluster failover, that resource is skipped upon checkpoint restart, regardless of whether the backup type or file type of that resource is supported by checkpoint restart. This saves media space and backup time.
- If failover occurs in the middle of backing up a resource, the media that was being used at the time of the failover is left unappendable and new media will be requested upon restart. It is recommended that you select an appropriate media overwrite protection level to ensure that media that was used prior to the failover is not overwritten upon restart.
- The data that is backed up upon restart is part of a different backup set than the data that was backed up prior to the failover. Separate catalog backup set entries are created for the data backed up prior to the failover and after the failover.

In addition, if multiple cluster failovers occur during the backup of a given resource, a different backup set is created each time the job restarts. These multiple backup sets allow potential for duplication of backed up data.

It is important to restore the backup sets in the order in which they were backed up. In addition, you should enable the Restore over existing files option when performing a restore operation on these backup sets to ensure that all the data included in the backup set is completely restored.

- If failover occurs during a post-backup verify job, or a pre-backup or post-backup database consistency check job, that job starts at the beginning after failover.
- Entries for full-volume backups that were interrupted by a cluster failover and resumed from the point of failover do not display in the IDR Restore Wizard. However, you can restore these backup sets manually after you make the initial recovery using the IDR Restore Wizard.
- You can enable the checkpoint restart option for a full backup job that backs up and deletes the files. However, if a cluster failover occurs and the job is resumed, the files are not deleted from the source volume after the backup completes.
- If a failover occurs on a clustered managed media server, the job that is recovered resumes on the active cluster node. The job will not be recovered to any other managed media servers outside of the Backup Exec cluster.

Enabling or disabling checkpoint restart

To apply checkpoint restart to backup jobs, make sure that the Error-Handling Rule for Cluster Failover is enabled.

See [“About error-handling rules”](#) on page 590.

To enable or disable checkpoint restart

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Settings**, click **Clusters**.
- 4 Select or clear **Apply CheckPoint Restart (resume jobs from point of failover)**.
Defaults set in Backup Exec remain the same on the failover nodes as they were on the controlling node when failover occurs.
- 5 Enable the Cluster Failover error-handling rule. On the **Tools** menu, select **Error-Handling rules**.
- 6 Select the Cluster Failover rule, and then click **Edit**.
- 7 Verify that the **Enabled** check box is selected.

Specifying a different failover node

You can do the following:

- Change the order in which the nodes fail over.
- Add a failover node to the cluster.
- Remove a failover node from the cluster.

To change the order in which nodes fail over

- By default, in a MSCS cluster, failover from the controlling node to a designated node occurs in alphabetical order according to the machine name of each node. To change the order in which failover will occur on the designated nodes, rename the machines in the order in which they should fail over.
- VCS uses a priority list as the primary method for determining the failover target. To set the priority in VCS, highlight the Backup Exec group in the VCS Cluster Explorer and select **Tools > System Manager**. Modify the priority field to reflect the desired order.

To add or remove a failover node

Before you add a node to the Backup Exec cluster configuration, you must install Backup Exec on it. Cluster services for a node should be online before you add or remove it from the cluster.

If you are removing a node, do not run the cluster configuration wizard from the node you want to remove.

- 1 On the controlling node, on the **Tools** menu, point to **Wizards**.
- 2 Click **Cluster Configuration Wizard**.
- 3 Follow the instructions on the screen to add or remove a node.
- 4 If you have added a failover node, also add any locally attached storage devices that are to be used when failover occurs to the cluster device pool. This ensures that jobs can be run on the storage devices that are attached to the failover nodes.

If you remove some, but not all, nodes in a cluster, an uninstall of Backup Exec results in a password being requested for the virtual server and the services continuing to run. You must remove Backup Exec from all nodes on the cluster.

See [“Uninstalling Backup Exec from a Microsoft cluster”](#) on page 820.

See [“Creating device pools for Microsoft Cluster Servers”](#) on page 821.

See [“Specifying a different failover node”](#) on page 824.

See [“Configurations for Backup Exec and Microsoft Cluster Servers”](#) on page 827.

Designating a new SAN SSO primary server and central administration server in a Microsoft Cluster Server

To designate a new primary SAN Shared Storage Option server or central administration server for a cluster environment, use BEUtility.exe. BEUtility enables you to do various types of configuration and maintenance operations on your Backup Exec media servers.

Note: In a cluster environment, do not use **Change Service Account** in BEUtility.exe.

To change a Backup Exec cluster from a Database Server to a Member Server

- 1 Install the new server as a secondary server with the Library Expansion Option and SAN Shared Storage Option installed.
Make sure connections to the Backup Exec cluster and other member servers are working properly.
- 2 Using the cluster administrator, shut down the Backup Exec cluster services.
Be sure to keep the Disk resource online.
- 3 Move the catalog files from the Backup Exec cluster installation path to the respective installation paths on the new database server.
- 4 Use BEUtility.exe to connect all Backup Exec servers to the new database server and to start all Backup Exec services.
- 5 Stop and restart the Backup Exec Services on the new database server.
- 6 Using the Cluster Administrator, move the Backup Exec resource group to the failover node and make sure services start on that node.
- 7 Use BEUtility.exe to stop and restart the Backup Exec Services on all the member servers of the SAN in order for them to connect to the new database server.

To change a Backup Exec cluster from a central administration server to a managed media server

- 1 Install the new server as a managed media server.
Make sure connections to the Backup Exec cluster and other managed media servers are working properly.
- 2 Using the cluster administrator, shut down the Backup Exec cluster services.
Be sure to keep the Disk resource online.
- 3 Move the catalog files from the Backup Exec cluster installation path to the respective installation paths on the new central administration server.
- 4 Use BEUtility.exe to connect all Backup Exec servers to the new central administration server and to start all Backup Exec services.
- 5 Stop and restart the Backup Exec Services on the central administration server.

- 6 Using the Cluster Administrator, move the Backup Exec resource group to the failover node and make sure services start on that node.
- 7 Use BEUtility.exe to stop and restart the Backup Exec Services on all the managed media servers in order for them to connect to the new central administration server.

See [“Multi-node clusters on a fibre channel SAN with the SAN SSO”](#) on page 832.

Configurations for Backup Exec and Microsoft Cluster Servers

Backup Exec supports various cluster configurations of between two and eight nodes on a fibre channel SAN, with locally attached storage devices, or with storage devices on a shared SCSI bus. You can use any combination of these configurations.

Note: If you install the cluster on a private network, use the Cluster Administrator to enable public communication if necessary.

If you are using a cluster on a fibre channel SAN or with storage devices on a shared SCSI bus and failover occurs, depending on the capability of your various SAN components, media might be orphaned in the tape drive until the failed node becomes active again.

If end-of-job markers were not written to the media before the failover occurred, the media may be marked as unappendable by the Backup Exec engine when the next append backup job is run. The media remains unappendable until it is overwritten (or erased, or the retention period expires, etc.).

If the storage device is a robotic library, you can review the Robotic Library Inventory report to discover if the media was marked unappendable by the Backup Exec engine. If the Full column reports a 3, the Backup Exec engine has marked the media as unappendable.

To add or remove hot-swappable devices in a cluster, run the Hot-swap Device Wizard on all Backup Exec Cluster nodes. If a server is not updated to recognize a new device, any job that is targeted to that device may fail.

See [“About adding or replacing devices by using the Hot-swappable Device Wizard”](#) on page 453.

Examples of various cluster configurations are available.

- See [“Two-node cluster with locally attached storage devices”](#) on page 828.

- See [“Two-node cluster with tape devices on a shared SCSI bus”](#) on page 828.
- See [“Multi-node clusters on a fibre channel SAN with the SAN SSO”](#) on page 832.

Two-node cluster with locally attached storage devices

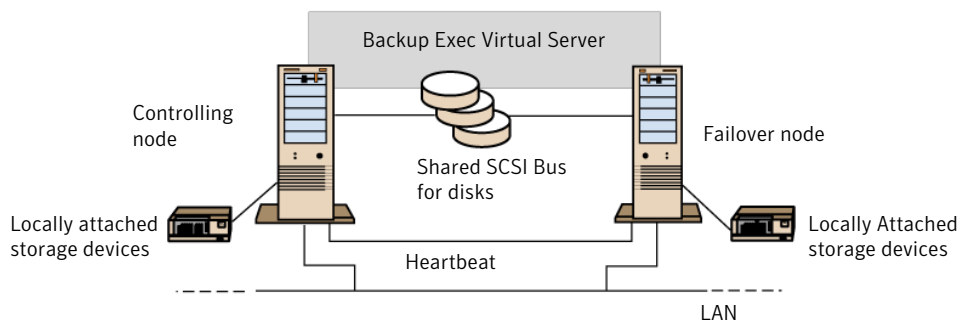
In this configuration, cluster-aware Backup Exec is installed on the controlling node, failover occurs to designated nodes in the cluster, and storage devices are locally attached to each node.

Each node’s locally attached storage devices are automatically assigned to the **All Devices (<Node Name>)** device pool, which is also the default destination device on that node when you create backup or restore jobs. You must create a device pool that includes storage devices on the controlling node and on each failover node in order for jobs to run when failover occurs.

See [“Creating device pools for Microsoft Cluster Servers”](#) on page 821.

To restore data in this configuration, move media to the failover node’s locally attached storage device and reinventory before starting a restore operation.

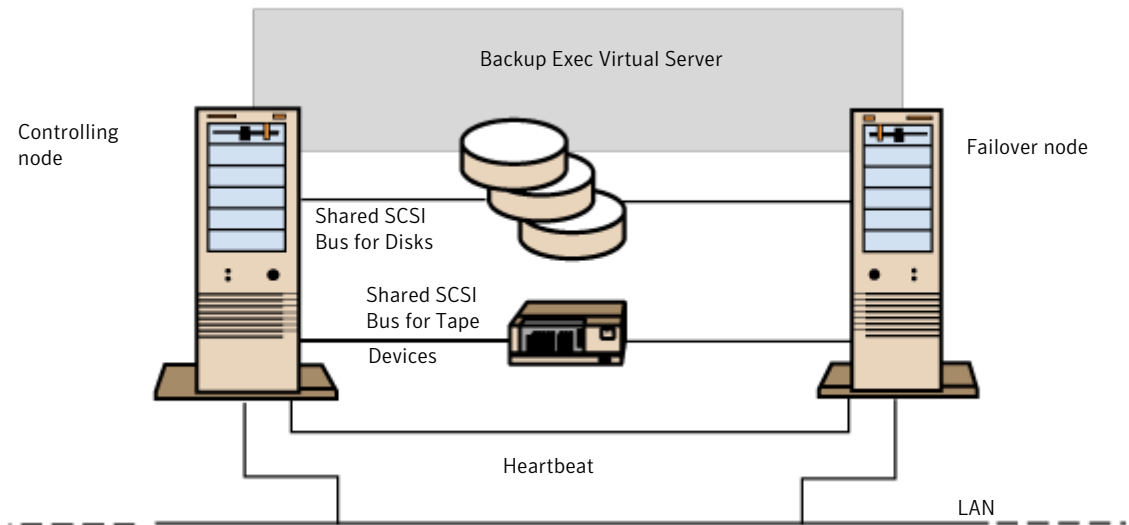
Figure 19-1 Two-node Cluster with Locally Attached Storage Devices



See [“Multi-node clusters on a fibre channel SAN with the SAN SSO”](#) on page 832.

Two-node cluster with tape devices on a shared SCSI bus

In this configuration, cluster-aware Backup Exec is installed on the controlling node, failover occurs to designated nodes in the cluster, and tape devices are attached to a shared SCSI bus that is separate from any shared SCSI bus for disks.

Figure 19-2 Two-node cluster with tape devices on a shared SCSI bus

Because each node creates a unique tape device name for the same device, if the drive is not serialized, this configuration requires you to create a device pool that includes the tape device name used by each node in order for jobs to run when failover occurs.

See [“Creating device pools for Microsoft Cluster Servers”](#) on page 821.

When failover occurs, a SCSI bus reset is issued. Therefore, tape devices and shared drives should not be connected to the same SCSI bus; each should be connected to separate SCSI buses.

See [“Configuring a shared SCSI bus for tape devices”](#) on page 829.

Note: If you are using a serialized tape device in a shared SCSI cluster configuration, media that is orphaned in a device because of a failover will be ejected from the tape device. If you are using a tape device that is not serialized, you need to manually eject the media from the device or reboot the device

See [“Multi-node clusters on a fibre channel SAN with the SAN SSO”](#) on page 832.

Configuring a shared SCSI bus for tape devices

Before configuring a shared SCSI bus for tape devices, please read the following carefully.

To configure tape devices on a shared SCSI bus, you must have SCSI cables, SCSI terminators, a SCSI adapter in each cluster node to provide a shared external bus between the nodes, and at least one tape device on the shared bus.

The tape devices must be connected to a bus that uses the same method of transmission that the device does (single-ended or differential). Only one transmission method can be used on a single SCSI bus, however, if the devices use different transmission methods, you can install a signal converter between the devices. A signal converter converts single-ended SCSI signals to differential SCSI signals.

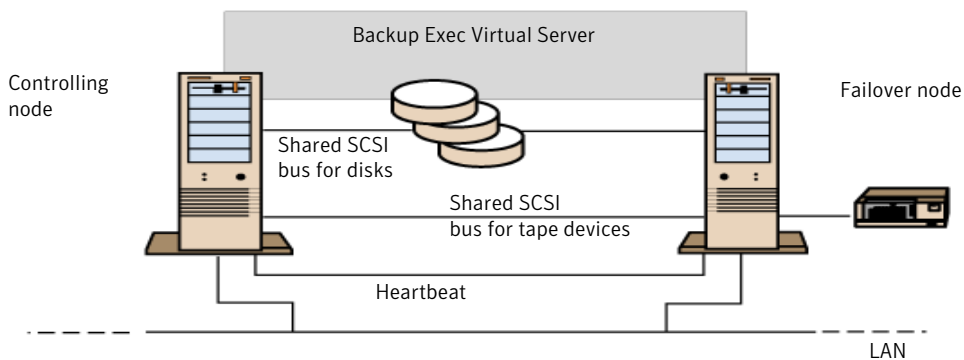
Note: You must use a signal converter to connect single-ended and differential devices in order to avoid hardware damage.

You must terminate the SCSI bus at both ends so that commands and data can be transmitted to and from all devices on the bus. Each SCSI bus must have two terminators and they must be at each end of the segment.

If a tape device is in the middle of the bus, remove any internal termination in that device.

If the tape device is at the end of the bus, and the tape device has internal termination, you can use the device's internal termination to terminate the bus.

Figure 19-3 Example of a shared bus with tape devices at the end of the bus



Following are methods you can use to terminate a bus:

- SCSI adapters. This method is not recommended because if the server is disconnected from the shared bus, or if there is a power supply failure, the bus may not be properly terminated and may be inoperable.
- Pass-through (or feed-through) SCSI terminators. These can be used with SCSI adapters and with some tape devices. If the tape device is at the end of the bus, you can attach a pass-through SCSI terminator to terminate the bus. The internal terminators in the tape device must be disabled. This is a recommended method.

Note: To ensure termination if a power supply failure occurs, turn off the on-board terminators on the SCSI controller (using the host adapter manufacturer's recommended method) and physically terminate the controller with a terminator.

- Y cables. These can be used with some tape devices. If the tape device is at the end of the bus, you can attach a terminator to one branch of a Y cable to terminate the bus. The internal terminators in the tape device must be disabled. This is a recommended method.
- Trilink connectors. These can be used with some tape devices. If the tape device is at the end of the bus, you can attach a terminator to one of the trilink connectors to terminate the bus. The internal terminators in the tape device must be disabled. This is a recommended method.

Besides terminating the bus, Y-cables and trilink connectors also allow you to isolate the devices from the shared bus without affecting the bus termination. You can maintain or remove that device without affecting the other devices on the shared SCSI bus.

To configure a shared SCSI bus for tape devices

- 1 Install the SCSI controllers for the shared SCSI bus.
Make sure that the SCSI controllers for the shared SCSI bus are using different SCSI IDs. For example, on the controlling node, set the SCSI controller ID to 6 and on the failover node, set the SCSI controller ID to 7.
- 2 Prepare the SCSI controllers for the shared SCSI bus. Refer to your SCSI host adapter manufacturer's documentation for details.
Do not have power on to both nodes while configuring the computers, or if both nodes have power on, do not connect the shared SCSI buses to both nodes.
- 3 Connect the shared SCSI tape devices to the cable, connect the cable to both nodes, and then terminate the bus segment using one of the methods discussed in the previous section.

See [“Two-node cluster with tape devices on a shared SCSI bus”](#) on page 828.

Multi-node clusters on a fibre channel SAN with the SAN SSO

In this configuration, one or more clusters are attached to a fibre channel storage area network (SAN), with cluster-aware Backup Exec and the SAN Shared Storage Option (SAN SSO) installed on the controlling node in each cluster. Shared secondary storage devices are attached to the fibre channel, although a single storage device can be shared between one or more clusters. Failover occurs (in alphabetical order of the machine name) to other designated nodes in the cluster.

Note: When using multiple clusters in a SAN SSO environment, it is strongly recommended that the cluster nodes be connected to the storage devices using a fibre switch. If you use a hub rather than a fibre switch, the hub will receive a reset command during a failover event that causes all other components attached to the hub to be disconnected. You can designate any server on the fibre channel SAN as the Shared Storage Option Database server.

You should create a failover device pool for the cluster.

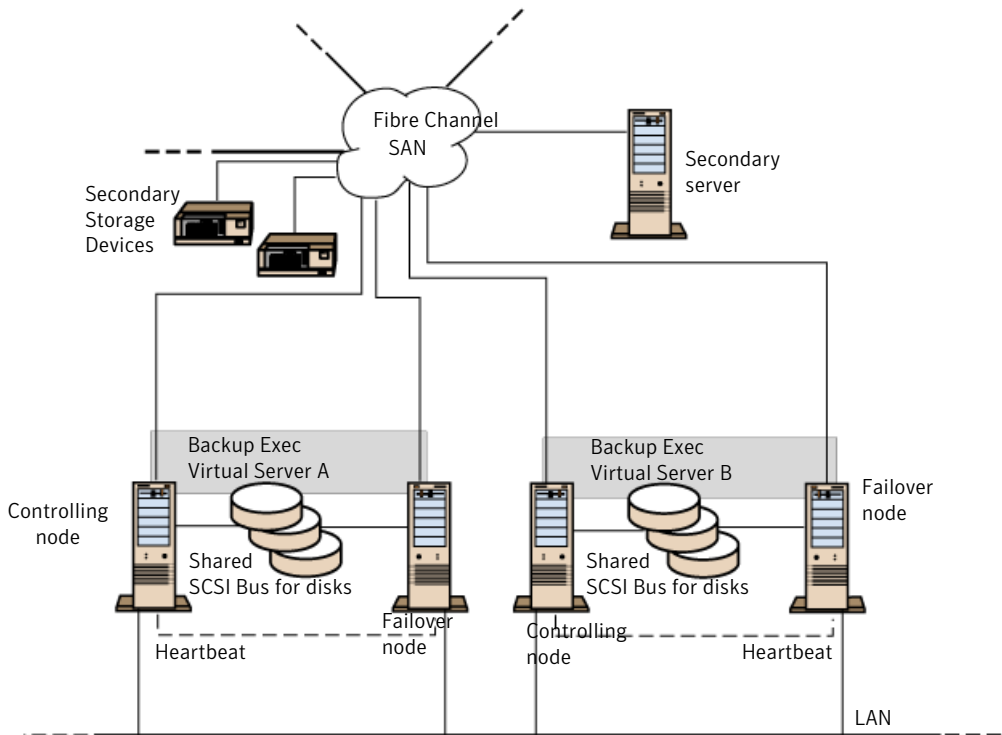
See [“Creating device pools for Microsoft Cluster Servers”](#) on page 821.

This configuration offers increased performance since backups are performed locally instead of over a network. Additionally, centralized media catalogs are available. Because the SAN SSO uses a shared catalog database, a tape that has already been cataloged can be physically moved from one device to another and not have to be recataloged.

Note: The SAN SSO option must be installed on each failover node, with the same settings that were used on the primary node. Either all nodes should be database servers or all nodes should be secondary member servers.

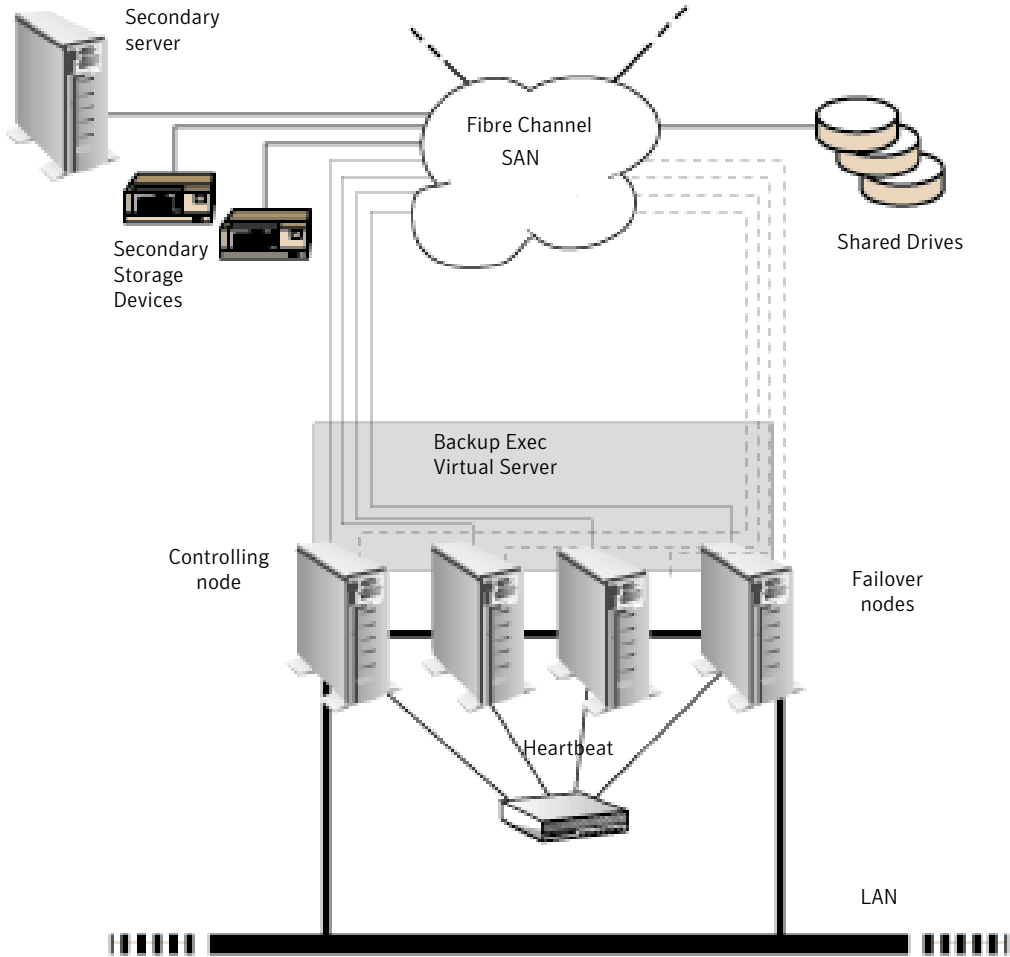
Following are examples of multi-node clusters:

Figure 19-4 Two 2-node Clusters on a Fibre Channel SAN with the SAN SSO



You can have a four-node cluster.

Figure 19-5 Four-node Cluster on Fibre Channel SAN with the SAN SSO



See [“About installing the SAN Shared Storage Option”](#) on page 1994.

See [“Designating a new SAN SSO primary server and central administration server in a Microsoft Cluster Server”](#) on page 825.

Using the Central Admin Server Option with Microsoft clusters and SAN SSO

Managed media servers can be clustered, however, it is not recommended because the central administration server recovers all failed jobs in a distributed job environment.

The following configurations can be used when installing Backup Exec clusters with the Central Admin Server Option (CASO) and SAN SSO.

- Backup Exec cluster with CASO
- Backup Exec cluster with CASO and SAN Shared Storage Option
- Backup Exec cluster with the managed media server configuration
- Backup Exec cluster with the managed media server configuration and SAN Shared Storage Option

To install Backup Exec cluster with CASO

- 1 Install Backup Exec with CASO and any additional options onto your Microsoft cluster nodes.
- 2 From the node that you want to designate as the active node, start Backup Exec.
- 3 From the **Tools** menu, point to **Wizards**, and then click **Cluster Configuration Wizard**.
- 4 Follow the instructions on the screen.
- 5 When the Cluster Configuration Wizard completes, install the managed media server. Use the virtual Backup Exec cluster name when prompted for the central administration server.

To install Backup Exec cluster with CASO and SAN Shared Storage Option

- 1 Install Backup Exec with CASO, the SAN Shared Storage Option, and any additional options onto your Microsoft cluster nodes.
- 2 From the node that you want to designate as the active node, start Backup Exec.
- 3 From the **Tools** menu, point to **Wizards**, and then click **Cluster Configuration Wizard**.
- 4 Follow the instructions on the screen.
- 5 When the Cluster Configuration Wizard completes, install the managed media server. Use the virtual Backup Exec cluster name when prompted for the central administration server and the primary SAN server.

To install Backup Exec cluster with the managed media server configuration

- 1 Install Backup Exec with the managed media server option and any additional options onto your Microsoft cluster nodes.

All nodes that run Backup Exec in the managed media server cluster configuration must access the same central administration server. If the nodes do not access the same central administration server, failovers do not occur properly.

- 2 From the node that you want to designate as the active node, start Backup Exec.
- 3 From the **Tools** menu, point to **Wizards**, and then click **Cluster Configuration Wizard**.
- 4 Follow the instructions on the screen.

To install Backup Exec cluster with the managed media server configuration and SAN Shared Storage Option

- 1 Install Backup Exec with the managed media server option, SAN Shared Storage Option, and any additional options onto your Microsoft cluster nodes.

All nodes that run Backup Exec in the managed media server cluster configuration must access the same central administration server. If the nodes do not access the same central administration server, failovers do not occur properly.

- 2 From the node that you want to designate as the active node, start Backup Exec.
- 3 From the **Tools** menu, point to **Wizards**, and then click **Cluster Configuration Wizard**.
- 4 Follow the instructions on the screen.

About backing up Microsoft Cluster Servers

To protect all data in the cluster, which includes file shares, databases, and the cluster quorum, back up the following:

- Local disks, Shadow Copy Components, and System State on each node. The cluster quorum, which contains recovery information for the cluster and information about changes to the cluster configuration, is included in the System State backup.
See [“Backing up local disks in a Microsoft cluster”](#) on page 837.
- All shared disks, including the data in the Microsoft Cluster Server folder on the Quorum disk.

See [“Backing up shared disks in a Microsoft cluster”](#) on page 838.

- Virtual servers, which may contain data or contain applications such as Microsoft SQL Server or Exchange Server. Use Backup Exec database agents to back up databases.

See [“Backing up database files in a Microsoft cluster”](#) on page 838.

Note: For offhost backup jobs that use the hardware provider, the media server and the remote computer must be in different cluster groups. The cluster applications cannot support devices’ logical unit numbers (LUNs) that have duplicate signatures and partition layouts, therefore, the snapshots containing the LUNs must be transported to a host, or remote computer, that is outside the cluster.

The Command Line Applet can be used with Backup Exec when Backup Exec is installed in a cluster. The only limitation is that you cannot use the Command Line Applet to specify a device for backup. You can use the Command Line Applet to target a device pool, but not a specific device in that pool.

See [“Backing up local disks in a Microsoft cluster”](#) on page 837.

See [“Backing up shared disks in a Microsoft cluster”](#) on page 838.

See [“Backing up database files in a Microsoft cluster”](#) on page 838.

See [“About backing up Windows 2000 and Windows Server 2003/2008 features in a Veritas cluster”](#) on page 847.

See [“Creating a backup job by using the Backup Wizard”](#) on page 331.

See [“Backing up Windows 2008 R2 cluster shared volumes”](#) on page 839.

Backing up local disks in a Microsoft cluster

Select local disks for backup from the physical node to which they are attached.

To back up local disks in a Microsoft cluster

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 From the backup selections pane, expand the domain that contains the nodes, and then select the local disks on each node.

When making backup selections for nodes running Windows 2000, be sure to select System State.

See [“About selecting data to back up”](#) on page 280.

- 4 If you created a device pool for the cluster, select it as the default destination device so that jobs can restart on the failover node if failover occurs.
- 5 Configure the remainder of the settings for the backup job.
- 6 Run the backup job now or schedule it to run later.

Backing up shared disks in a Microsoft cluster

Select shared disks for backup by selecting them from the Microsoft Cluster virtual server or from the Backup Exec virtual server.

To back up shared disks

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the backup selections pane, expand the domain that contains the nodes, and then select either the Microsoft Cluster virtual server or the Backup Exec virtual server. The virtual servers allow your backup jobs to access shared data via any node that controls the disk.
- 4 Select the drive letters that represent the shared disks.
- 5 If you created a device pool for the cluster, select it as the default destination device so that jobs can restart on the failover node if failover occurs.
- 6 Configure the remainder of the settings for the backup job.
- 7 Run the backup job now or schedule it to run later.

To browse clusters in Active Directory Domains, you must enable Kerberos authentication on each virtual cluster. You can enable Kerberos authentication from Microsoft's Cluster Administrator.

Backing up database files in a Microsoft cluster

Select database files for back up from a database icon on a virtual server. If a virtual server contains a database application such as Microsoft SQL Server or Exchange Server, use the appropriate Backup Exec database agent to perform the backup operations; otherwise, only the file system is backed up, not the database files.

To back up database files in a cluster

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections pane, expand the domain that contains the cluster, then expand the virtual server that contains the database files.

- 4 Check the database files.
- 5 Set database-specific defaults. Refer to the specific Backup Exec database agent documentation for details on how to set backup defaults for the database.
- 6 If you created a device pool for the cluster, select it as the default destination device so that jobs can restart on the failover node if failover occurs.
- 7 Configure the remainder of the settings for the backup job.
- 8 Run the backup job now or schedule it to run later.

Backing up Windows 2008 R2 cluster shared volumes

Backup Exec supports backing up and restoring Microsoft Windows 2008 R2 cluster shared volumes.

After it detects each cluster shared volume, Backup Exec places each volume under the cluster name where the shared volume resides. Cluster names appear under **Windows Systems** in the backup selections pane.

Note: You can also add cluster names in **User-defined Selections**.

See [“Backing up Windows 2008 R2 cluster shared volumes”](#) on page 839.

Note: You cannot view properties for cluster shared volumes when you browse Windows Server 2008 R2 clusters from a media server that runs Windows XP/Server 2003. However, cluster shared volume properties can be seen when Backup Exec is installed on a computer that runs Windows Server 2008 or later.

To back up Microsoft Windows 2008 R2 Hyper-V files, Symantec recommends that you use the *Symantec Backup Exec Agent for Microsoft Hyper-V*.

To restore Windows 2008 R2 cluster shared volumes, use normal restore procedures.

See [“About restoring data to a Microsoft cluster”](#) on page 840.

To back up Windows 2008 R2 cluster shared volumes

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 From the backup selections pane, expand **Windows Systems**.
- 4 Select the cluster where the cluster shared volumes reside.
- 5 Select the cluster share volumes that you want to back up.

- 6 If you created a device pool for the cluster, select it as the default destination device.

The device pool that you select as the default destination device ensures that jobs successfully restart on the failover node if a failover occurs.

See [“Creating device pools for Microsoft Cluster Servers”](#) on page 821.

- 7 Select additional backup job options, if appropriate.

See [“Creating a backup job by setting job properties”](#) on page 332.

- 8 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the Properties pane, under **Frequency**, click **Schedule**.
- Set the schedule options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

About restoring data to a Microsoft cluster

For all file restore operations, including redirecting restores, use the normal restore procedures.

See [“Restoring data by setting job properties”](#) on page 605.

When restoring files to shared drives, direct those files to the virtual server or the controlling node of the resource. When restoring individual database files, such as Microsoft SQL Server or Exchange Server, direct those files to the virtual server name of a specific installation of the SQL or Exchange database.

See [“Restoring the cluster quorum for Windows Server 2003/2008 computers to a Microsoft cluster”](#) on page 840.

See [“Specifying a new drive letter for the cluster quorum disk”](#) on page 841.

Restoring the cluster quorum for Windows Server 2003/2008 computers to a Microsoft cluster

The cluster quorum is backed up as part of System State.

You may need to specify a new disk to which the cluster quorum will be restored.

See [“Specifying a new drive letter for the cluster quorum disk”](#) on page 841.

To restore the cluster quorum

- 1 Take the other nodes in the cluster offline.
- 2 On the navigation bar, click the arrow next to **Restore**.
- 3 Click **New Restore Job**.
- 4 On the **Properties** pane, under **Settings**, click **Advanced**.
- 5 Select **Restore Cluster Quorum**.
- 6 Select **Force the recovery of the cluster quorum even if other nodes are online and/or disk signatures do not match** in the following circumstances:
 - If you are not able to take the other nodes in the cluster offline. When this option is selected, the cluster service for any nodes that are online is stopped.
 - If the disk that the cluster quorum previously resided on has been changed. The disk may have been replaced with a new one, or the disk configuration may have been changed so that the cluster quorum now resides on a different disk. This option allows the drive letter of the disk that the cluster quorum was on to remain the same, even if the configuration has changed and the disk signatures contained in the restore media do not match the disk signatures contained in the cluster quorum.
- 7 Select any additional options as needed for this restore job.
See [“Restoring data by setting job properties”](#) on page 605.
- 8 When the restore operation is complete, use the cluster administrator software to restart the cluster service on nodes on which it was stopped.

Specifying a new drive letter for the cluster quorum disk

To use the `clrest.exe` command-line utility to specify a new drive letter for the cluster quorum disk, restore System State, but not the cluster quorum. When System State is restored, the cluster quorum is copied to a default location, `%SystemRoot%\cluster\BackupExec`.

Then, use `clrest.exe` with the `[drive letter]` option to restore the cluster quorum to the quorum drive, which will be assigned the drive letter you specified.

To specify a new drive letter for the cluster quorum disk on Windows 2000

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **Restore Selections** pane, click **System State**.
- 4 On the **Properties** pane, under **Settings**, click **Advanced**.

- 5 Clear the **Restore cluster quorum** option. This option must not be selected.
- 6 Start the restore operation.
During the restore, the cluster quorum files are copied to the default location %SystemRoot%\cluster\BackupExec.
- 7 When the restore has completed, reboot the target node.
- 8 After reboot is complete, run clrest.exe from the command line to restore the cluster quorum from the default location to the quorum disk.

```
clrest path [-f] [drive letter]
```

where

path is the complete path to the cluster quorum; typically, the pathname is %SystemRoot%\cluster\BackupExec. A pathname is required.

[-f] forces the restore to proceed even if other cluster nodes are online and/or the disk signatures do not match. When this option is selected, the cluster service for any nodes that are online is stopped. This option also allows the drive letter of the disk that the cluster quorum was on to remain the same, even if the configuration has changed and the disk signatures contained in the restore media do not match the disk signatures contained in the cluster quorum.

[drive letter] specifies another drive letter for the quorum disk. If you use this option, the drive letter that the cluster quorum resides on will be changed to the drive letter specified. Otherwise, the drive letter that the cluster quorum resides on will stay the same as it was previously.

- 9 After the cluster quorum is restored, use the cluster administrator to bring the other cluster nodes online.

Using Backup Exec with Veritas Cluster Server

If you use Veritas Cluster Server (VCS), there are three possible options for integrating with Backup Exec. The first option consists of managed media servers installed on each of the VCS nodes and a central administration server. You can install the central administration server on any of the nodes or on a system that is not a VCS node. However, the central administration server must be in the same domain as the VCS nodes.

Selections for backup are made on the central administration server, and then based on the backup selections of clustered resources, the central administration server sends the job to the managed media server on which the clustered resource is currently active. The central administration server attempts to balance backup

jobs of clustered resources so that the job is run as a local job. If a failover occurs, backup jobs that were running are rescheduled and then the central administration server restarts the job on the new active node for the failed resource.

The second option is to cluster the Backup Exec application using Veritas Cluster Server, thereby making it highly available. A wizard is provided to guide you through the configuration process.

The third option involves installing the Backup Exec Remote Agent on each of the VCS nodes. The standalone media server can then be installed on any of the nodes or outside of the cluster. In this environment, backup jobs for clustered resources are completed remotely.

Specific details of how Backup Exec runs in a cluster vary depending on the configuration you use in the cluster.

See [“Installing Backup Exec with the CASO option on a Veritas Cluster Server”](#) on page 844.

See [“Installing Backup Exec on a Microsoft Cluster Server”](#) on page 818.

See [“Clustering Backup Exec using Veritas Cluster Server”](#) on page 845.

Requirements for installing Backup Exec with the CASO option on a Veritas Cluster Server

Following are requirements for installing Backup Exec and CASO on a Veritas Cluster Server:

- 32-node clusters are supported with Backup Exec for the following operating systems: Windows Server 2003, Windows Server 2003 Enterprise, and Windows Server 2003 DataCenter.
- An individually licensed copy of Backup Exec, as well as any applicable agents and options, is required for each active node in the cluster as defined in the End User License Agreement. You must enter a license key for each node in the cluster (the cluster must have at least two nodes).
- Storage Foundation for Windows Servers High Availability server components must be installed on the cluster nodes.
- Storage Foundation for Windows Servers High Availability Administrative Console components must be installed on the central administration server.

See [“Installing Backup Exec on a Microsoft Cluster Server”](#) on page 818.

Installing Backup Exec with the CASO option on a Veritas Cluster Server

The Remote Agent is automatically installed on all the nodes in the cluster. If this installation of Backup Exec will be used to back up remote servers outside the cluster, install the Remote Agent on those remote servers as well.

To install Backup Exec with the CASO option on a Veritas Cluster Server

- 1 Install Backup Exec as a managed media server on all the nodes that you want to include in the cluster. All installations must target local drives. Use the same installation path for each node.
- 2 Install the Backup Exec Central Admin Server Option (CASO) on a computer that is included in the same domain as the cluster nodes.

The Central Admin Server Option can be installed on any of the cluster nodes or outside of the cluster.

- 3 If the managed media server contains locally attached drives, create a drive pool that contains all the locally attached storage devices on each node to use when failover occurs. This ensures that jobs can run on the storage devices attached to the failover nodes.
- 4 If the Backup Exec Central Admin Server Option is not installed, then install the Backup Exec Remote Agent for Windows Systems on the local drives of all the nodes in the cluster.

See [“About installing the Remote Agent for Windows Systems”](#) on page 139.

Requirements for clustering Backup Exec using Veritas Cluster Server

Review the following items before you cluster Backup Exec using Veritas Cluster Server:

- 32-node clusters are supported with Backup Exec for the following operating systems: Windows Server 2003, Windows Server 2003 Enterprise and DataCenter, and Windows Server 2003 DataCenter.
- An individually licensed copy of Backup Exec, as well as any applicable agents and options, is required for each active node in the cluster as defined in the End User License Agreement. You must enter a license key for each node in the cluster (the cluster must have at least two nodes).
- Storage Foundation for Windows Servers High Availability server components must be installed on the cluster nodes.

- A shared volume accessible by all nodes in the cluster.
- The shared volume must be part of a dynamic cluster disk group.
- During installation of a Backup Exec cluster, the node that runs the cluster wizard should have exclusive control over the shared volume.
- The shared volume cannot reside on a disk with compression enabled.
- Symantec highly recommends that you use the default database instance (MSDE) that is installed by Backup Exec if you plan to cluster Backup Exec.
- Symantec also supports using a remote SQL Server instance to host the Backup Exec database. However, if you plan to use this scenario, review the following: Only one installed instance of Backup Exec can be installed into the remote SQL Server instance on a clustered node. All other installed instances of Backup Exec in the cluster must use the default Backup Exec MSDE database instance.

Note: You must run the Backup Exec cluster wizard on the cluster node that uses the remote SQL Server instance.

Clustering Backup Exec using Veritas Cluster Server

Clustering Backup Exec makes the application highly available in an active passive configuration. The Backup Exec services can only run on one cluster node at a time. Should the active node go offline, the services and any active jobs will be restarted on another node in the cluster.

To cluster Backup Exec using Veritas Cluster Server

- 1 Install and configure Veritas Cluster Server on each machine that will be a part of the cluster. For more information on installing and configuring Veritas Cluster Server, refer the Storage Foundation for Windows Servers High Availability Edition Administrator's Guide.
- 2 Install Backup Exec to the local drive on each machine that will be part of the Backup Exec service group.
- 3 Create a dynamic disk group and assign a drive letter to the volume that will be used as the shared disk resource. For more information on creating dynamic disk groups, refer the Storage Foundation for Windows Servers High Availability Edition Administrator's Guide.
- 4 Verify that the volume is online and a drive letter is assigned to it only on the node where the Backup Exec Cluster Configuration Wizard will be run.
- 5 On the **Tools** menu, click **Wizards > Cluster Configuration Wizard**.

- 6 At the **Cluster Configuration Wizard** welcome screen click **Next**.
- 7 Enter a name for the Backup Exec cluster group or use the default name.
The cluster group name must not contain spaces.
- 8 The wizard lists the shared location where the Backup Exec application files will be copied. To specify a different location click **Change**.
- 9 Click **Next**.
- 10 Enter a name for the Backup Exec virtual server or use the default name.
- 11 Enter the IP address and subnet mask of the virtual server.
- 12 Click **Next**.
- 13 The wizard will validate the entries.
- 14 Select the nodes that will participate in the cluster. By default, the node on which the cluster configuration wizard was run is included as part of the cluster group.
- 15 Click **Next**.
- 16 Click **Configure** to have the wizard create the service group and move the files to the shared disk.

About backing up Veritas Cluster Servers

To protect all data in the cluster, including file shares and databases, back up the following:

- Local disks and System State on each node
- All shared disks
- Virtual servers, which may contain data or contain applications such as Microsoft SQL Server or Exchange Server. Use Backup Exec database agents to back up databases.

See “[Creating a backup job by using the Backup Wizard](#)” on page 331.

Note: For offhost backup jobs that use the hardware provider, the media server and the remote computer must be in different cluster groups. The cluster applications cannot support devices’ logical unit numbers (LUNs) that have duplicate signatures and partition layouts, therefore, the snapshots containing the LUNs must be transported to a host, or remote computer, that is outside the cluster.

The Command Line Applet can be used with Backup Exec when Backup Exec is installed in a cluster. The only limitation is that you cannot use the Command Line Applet to specify a device for backup. You can use the Command Line Applet to target a device pool, but not a specific device in that pool.

See [“Backing up local disks in a Veritas cluster”](#) on page 848.

See [“Backing up shared disks in a Veritas cluster”](#) on page 848.

See [“Backing up database files in a Veritas cluster”](#) on page 849.

See [“About backing up Windows 2000 and Windows Server 2003/2008 features in a Veritas cluster”](#) on page 847.

About backing up Windows 2000 and Windows Server 2003/2008 features in a Veritas cluster

You must purchase and install the Backup Exec Remote Agent for Windows Systems on all remote Windows 2000 and Windows Server 2003/2008 computers that you want backed up.

Without the Remote Agent, the following Windows 2000 features cannot be correctly backed up:

- Encrypted files
- SIS files
- Disk quota data
- Removable Storage data
- Remote Storage data
- Mount points
- Sparse files
- Windows Management Instrumentation
- Terminal Services
- System State data, including:
 - COM+ Class Registration database
 - Boot and system files
 - Registry
 - Certificate Services database (if the server is operating as a certificate server)
 - Active Directory (if the server is a domain controller)

- SYSVOL (if the server is a domain controller)

Note: You can select System State for backup on a remote computer only when the Remote Agent is installed on the remote computer.

Backing up local disks in a Veritas cluster

Select local disks for backup from the physical node to which they are attached.

Note: If the computer on which you are running a backup using the Advanced Open File Option is in an environment with the Central Admin Server Option and the Veritas Cluster Server installed, and if failover occurs to a VCS node, you must manually clean up the snapshots before restarting the backup on the failover node. Refer to the VSFW documentation for details.

To back up local disks in a Veritas cluster

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 On the backup selection pane, expand the domain that contains the nodes, and then select the local disks on each node.
See [“About selecting data to back up”](#) on page 280.
- 4 If you created a device pool for the cluster, select it as the default destination device so that jobs can restart on the failover node if failover occurs.
- 5 Configure the remainder of the settings for the backup job.
- 6 Run the backup job now or schedule it to run later.

Backing up shared disks in a Veritas cluster

Select shared disks for backup by selecting them from the Veritas Cluster Server virtual server.

To back up shared disks

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 On the **Backup selections** pane, expand the domain that contains the nodes, and then select either the Veritas Cluster Server virtual server. The virtual servers allow your backup jobs to access shared data via any node that controls the disk.

- 4 Select the drive letters that represent the shared disks.
- 5 If you created a device pool for the cluster, select it as the default destination device so that jobs can restart on the failover node if failover occurs.
- 6 Configure the remainder of the settings for the backup job.
- 7 Run the backup job now or schedule it to run later.

Backing up database files in a Veritas cluster

Select database files for back up from a database icon on a virtual server. If a virtual server contains a database application such as Microsoft SQL Server or Exchange Server, use the appropriate Backup Exec database agent to perform the backup operations; otherwise, only the file system is backed up, not the database files.

To back up database files in a Veritas cluster

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Backup Selections** pane, expand the domain that contains the cluster, then expand the virtual server that contains the database files, and then select the databases.
- 4 Set database-specific defaults. Refer to the specific Backup Exec database agent documentation for details on how to set backup defaults for the database.
- 5 If you created a device pool for the cluster, select it as the default destination device so that jobs can restart on the failover node if failover occurs.
- 6 Configure the remainder of the settings for the backup job.
- 7 Run the backup job now or schedule it to run later.

About restoring data to Veritas Cluster Servers

For all file restore operations, including redirecting restores, use the normal restore procedures.

See [“Restoring data by using the Restore Wizard”](#) on page 604.

When restoring files to shared drives, direct those files to the virtual server or the controlling node of the resource. When restoring individual database files, such as Microsoft SQL Server or Exchange Server, direct those files to the virtual server name of a specific installation of the SQL or Exchange database.

About backup job failover with Veritas Cluster Servers

A central administration server configuration installed in a Veritas Cluster Server environment is automatically enabled to failover a job. When you select cluster resources for backup, only a single clustered resource can be included in the backup job. This is necessary to ensure that the central administration server re-delegates the backup job to the designated failover node on the cluster.

If a failover of a clustered resource occurs during backup, the job on the managed media server stops and notification is sent to the central administration server. The central administration server then re-delegates the job to the managed media server that has become the new active node of the failed clustered resource.

Disaster recovery of a cluster

Prepare for recovery by creating a disaster preparation plan.

See [“About key elements of a disaster preparation plan \(DPP\)”](#) on page 776.

Prepare to restore SQL, Exchange, Oracle, and Lotus Domino databases in a cluster after a disaster by reading the sections on preparing for disaster recovery in the appropriate chapters.

In addition to the initial preparation instructions, further action is required to completely protect the Microsoft clusters.

If a disaster occurs, the following information is required to successfully recover the cluster:

- General Cluster Information
 - Cluster name
 - Cluster IP address and subnet mask
 - Cluster node names
 - Node IP addresses
 - Local and shared drive letters and partition scheme
 - Disk signatures
- Cluster Groups
 - Group name
 - Preferred nodes
 - Failover/failback policies
- Cluster Resources

- Resource name
- Resource type
- Group membership
- Possible owners
- Resource dependencies
- Restart and Looks Alive/Is Alive properties
- Resource-related parameters
- Application-specific configuration (SQL Database Character Set)
- If you are recovering a Microsoft Cluster Server, run Dumpcfg.exe from the Microsoft 2000 Resource Kit or and the Clusterrecovery.exe from the Microsoft 2003 Resource Kit to retrieve the disk signatures from the shared disk. The Microsoft 2000 Resource Kit allows you to replace disk signatures.
- If you are recovering a Veritas Cluster Server, run Vmgetdrive.exe to retrieve the disk signatures, disk group, and volume information from the shared disk.

Using IDR to prepare for disaster recovery of a cluster

Backup Exec provides a fully-automated disaster recovery solution called the Intelligent Disaster Recovery (IDR) Option, which allows you to quickly and efficiently recover the nodes that comprise the server cluster after a disaster. Oracle servers and SAP databases cannot be restored using IDR. For more information about disaster recovery for these options, see the appropriate chapters.

See [“About preparing computers for IDR”](#) on page 1808.

Note: To change the setup, use hardware, or a hardware configuration that is different from the original configuration, you must perform a manual recovery.

Recovering nodes on the cluster using IDR

If you used Backup Exec’s Intelligent Disaster Recovery to prepare for a disaster, you can use IDR to recover the nodes to their pre-disaster state.

Note: You must create disaster recovery media for each Windows 2000 and Windows 2003 cluster node. Disaster recovery media is customized for a single computer. You will not be able to use the media interchangeably between the nodes in a cluster.

When recovering both nodes in a cluster, make sure that the drive letters match the original cluster configuration. The scaled-back version of Windows that runs the recovery wizard may detect the hard drives in a different order than what was originally configured under the original version of Windows.

If the original configuration does not match, then to a certain extent, you can control the hard drive numbering scheme that Windows devises.

If you cannot get the IDR Disaster Recovery Wizard to properly detect the hard drive order, you can still manually set up hard drive partitions using the Disk Administrator option within the Disaster Recovery Wizard. After this is done, you can continue with automated restore of your backup media afterward.

Note: After Windows has been installed, you cannot change the system drive's letter. You must restore the system to the same drive letter from which it was backed up.

To recover nodes on the cluster using IDR

- 1** If you are recovering more than one node, disconnect the shared disks. If you are recovering only one node, the shared disks do not need to be disconnected.

If all nodes in the cluster are unavailable and must be recovered, the cluster cannot fail over. Disconnect the shared disks before recovery begins.
- 2** Restore the nodes.

See [“Recovering a computer by using the Intelligent Disaster Recovery Wizard”](#) on page 1831.
- 3** Reconnect the shared drives and bring the nodes online.
- 4** To restore a database to the shared drives, use the appropriate Backup Exec agent.

See [“Disaster recovery of SQL”](#) on page 1312.

See [“Recovering a Lotus Domino server from a disaster”](#) on page 1087.

See [“About restoring and recovering Oracle resources”](#) on page 1341.

Recovering Backup Exec on a Microsoft cluster using IDR

To fully restore a cluster on which Backup Exec is installed, you can use IDR to restore the cluster node and all shared disks or you can rebuild the cluster. To remotely restore the cluster catalog the media that contains the backup sets of the cluster nodes and the shared disk.

To recover Backup Exec on a Microsoft Cluster with IDR:

- 1 Replace all shared disks, if necessary.
- 2 Run the IDR Recovery Wizard on one of the nodes. During this process, use the disk manager to repartition all shared disks to their original configuration. Restore the local disk, system state, and the data files to the shared disk.
- 3 Reboot the server.
The cluster service and all other cluster applications should come online.
- 4 Run the IDR Recovery Wizard on all other nodes. Restore only the local disk and system state.

Recovering the entire cluster using a manual disaster recovery procedure

As part of the manual recovery process, you must reinstall Windows, including the last service pack applied before the failure.

See [“Disaster recovery of SQL”](#) on page 1312.

See [“Recovering a Lotus Domino server from a disaster”](#) on page 1087.

See [“About restoring and recovering Oracle resources”](#) on page 1341.

To recover the entire cluster manually

- 1 On the first node you want to recover, reinstall Windows, including the last service pack applied before the failure.
See [“About manual disaster recovery of Windows computers”](#) on page 780.
- 2 On the other nodes you want to recover, reinstall Windows, including the last service pack applied before the failure.
- 3 Reinstall the cluster services and bring the cluster online.
Do the following:
 - If you are recovering a Veritas Cluster Server, install the Storage Foundation for Windows High Availability server components, which include Volume Manager, and then use Volume Manager to create disk groups and volumes that match the original cluster configuration.
 - If you are recovering a Microsoft Cluster Server, after booting the nodes in a cluster, make sure that the drive letters match the original cluster configuration. If the original configuration does not match, then to a certain extent, you can control the hard drive numbering scheme that Windows devises by using the Disk Administrator.
- 4 Do one of the following:

- If you are recovering a Veritas Cluster Server, re-install Backup Exec. See “[Installing Backup Exec with the CASO option on a Veritas Cluster Server](#)” on page 844.
 - If you are recovering a Microsoft Cluster Server, use the Cluster Wizard to reinstall Backup Exec on the cluster. You must use the same settings used during the initial installation. See “[Installing Backup Exec on a Microsoft Cluster Server](#)” on page 818.
- 5 Catalog the media in the cluster.
 - 6 On the Backup Exec navigation bar on the active node, click **Restore**.
 - 7 In the **Restore selections** pane, select the last full backup sets made of the active node, and then select **System State**.
 - 8 Do one of the following:
 - If you are recovering a Veritas Cluster Server, proceed to step 9.
 - If you are recovering a Microsoft Cluster Server, on the **Properties** pane, under **Settings**, click **Advanced** and then select the **Restore cluster quorum** option (this option must be selected).
 - 9 Start the restore operation.
 - 10 When the restore has completed, reboot the active node.
 - 11 For each node that you need to recover, repeat step 6 through step 10.
 - 12 After all nodes are recovered, restore the Backup Exec data files, and all other data files, to the shared disks.
 - 13 To restore a database to the shared disks, use the appropriate Backup Exec agent.

Restoring the Microsoft Cluster Server data files

To fully recover the cluster, the cluster files in the MSCS folder may need to be restored. If the Quorum disk is still available and has not changed, then you do not have to restore the data files. If the Quorum disk is new, you need to restore the data files to the new Quorum disk. You should disable the cluster disk driver before restoring the data files.

To restore the cluster data files

- 1 Shut down the secondary nodes.
- 2 Start the primary node.
- 3 On the **Computer Management** menu, select **System Tools**. Then select **Device Manager**.

- 4 Right-click the cluster disk driver, and then select **Disable**.
- 5 Click **OK**.
- 6 Reboot the primary node.
- 7 On the Backup Exec navigation bar, click the arrow next to **Restore**.
- 8 Click **New Restore Job**.
- 9 In the **Restore selections** pane, select the most recent backup set of the MSCS folder.
The cluster service should not be running.
- 10 Redirect the restore of the MSCS folder to the designated Quorum disk.
- 11 After the cluster data files have been restored to the Quorum disk, you can enable the cluster disk driver and start the cluster service.
- 12 After the cluster quorum is restored, use the cluster administrator to bring the other cluster nodes online.

Recovering all shared disks in a Microsoft cluster

Recover shared disks using either the Dumpcfg option from the Microsoft 2000 Resource Kit or Cluster recovery from the Microsoft 2003 Resource Kit, which helps automate the recovery process or by performing a manual recovery.

To recover all shared disks using Dumpcfg

- 1 Disable the cluster disk driver on all nodes in order to gain access to the new disk.
- 2 On the Computer Management menu, select **System Tools**. Then select **Device Manager**.
- 3 Right-click the cluster disk driver, and then select **Disable**.
- 4 Replace and then repartition the shared disk. Use Disk Manager to verify that all nodes have access to the same shared disk.
- 5 Run Dumpcfg or Clusterrecovery to replace the disk signature for the Quorum disk.
- 6 Using a remote Backup Exec server, restore the cluster files to the Quorum disk via the node that has access to the disk.
- 7 Enable the cluster disk driver on all nodes.
- 8 On the **Computer Management** menu, select **System Tools**. Then select **Device Manager**.

9 Right-click the cluster disk driver, and then select **Enable**.

10 Reboot all cluster nodes.

To recover all shared disks without using Dumpcfg

1 Uninstall all cluster applications and the cluster software from both nodes.

2 Replace and then use Disk Manager to repartition the shared disk to the previously saved configuration.

3 Reinstall the cluster software.

4 Reinstall the cluster-aware version of Backup Exec on the cluster.

See “[Installing a custom installation of Backup Exec](#)” on page 119.

5 Reinstall additional cluster-aware software applications on the shared disk.

6 Use Backup Exec to restore any data from the catalogs.

Recovering all shared disks in a Veritas cluster

You can recover shared disks in a Veritas cluster using the Veritas Volume Manager.

To recover all shared disks using Volume Manager

1 Use Volume Manger to recreate all shared volumes and disk groups.

2 At a command prompt, type: `vmgetdisk` and press ENTER or RETURN.

The command creates a file named `VmDriveInfo.txt` which contains information about the disk groups and volumes.

3 Use a text editor, such as Notepad to open the `VmDriveInfo.txt` file.

4 From the directory in which Veritas Cluster Server was installed, use a text editor such as Notepad to open the `Main.cf`.

5 Find and replace the GUIDs in `Main.cf` with the GUIDs in the `VmDiskInfo.txt` file for the all the disk groups you want to recover.

Ensure that Lanman and MountV resources start. If you recover SQL or Exchange you cannot start the SQL or Exchange resource, but you can start the Lanman and MountV resources to restore data.

6 Restore the shared information using the virtual server backup.

Recovering Backup Exec in a Microsoft cluster

If you used the IDR option to prepare disaster recovery media for the shared disks, you must use a manual process to recover Backup Exec on a shared disk.

To use a manual process to recover Backup Exec on a shared disk

- 1 Replace the shared disk if necessary, and add that disk to the cluster as a disk resource.
- 2 Reinstall the cluster-aware version of Backup Exec on the cluster using the same information used in the original installation.
See [“Installing a custom installation of Backup Exec”](#) on page 119.
- 3 Use Backup Exec to restore any data from the catalogs.

Troubleshooting clusters

If you experience problems with using Backup Exec in a cluster environment, review the questions and answers in this section.

Table 19-2 Cluster troubleshooting questions and answers

Question	Answer
After I recovered my cluster and all shared disks, the cluster service will not start. Why won't it start and how can I get it started?	<p>The cluster service may not start because the disk signature on the Quorum disk is different from the original signature. If you have the Microsoft 2000 Resource Kit use <code>Dumpcfg.exe</code> or <code>Clusterrecovery</code> from the Microsoft 2003 Resource Kit to replace the disk. For example, type:</p> <pre>dumpcfg.exe /s 12345678 0</pre> <p>Replace 12345678 with the disk signature and replace 0 with the disk number. You can find the disk signature and the disk number in the event log.</p> <p>If you do not have the Microsoft 2000 Resource Kit, you can use <code>-Fixquorum</code> to change the Quorum disk signature.</p> <p>See “Changing the Quorum disk signature” on page 859.</p>

Table 19-2 Cluster troubleshooting questions and answers (*continued*)

Question	Answer
<p>I used the Checkpoint Restart option for my backups. During one of my backups, a Microsoft cluster failover occurred. Multiple backup sets were created. When I try to verify or restore using these backup sets, an "Unexpected End of Data" error occurs on the set that contains the data that was backed up prior to the failover. Why does this occur? Is my data safe?</p>	<p>You received this error because failover occurred in the middle of backing up the resource, therefore the backup set was not closed on the media. However, the objects that were partially backed up in the first backup set were completely backed up again during restart, ensuring data integrity. Therefore, all of the objects on the media for the given backup set should still be restored and verified.</p>
<p>I clustered a primary SAN server with a secondary SAN server. Now the device and media service on the secondary server fails. Why?</p>	<p>This occurs when the secondary server becomes the active node and attempts to connect to the Backup Exec database on the primary server, which is no longer available. To correct this, you must use the Backup Exec Utility (BEUTILITY.EXE) or reinstall the secondary server to be a primary server.</p>

Table 19-2 Cluster troubleshooting questions and answers (*continued*)

Question	Answer
<p>An Advanced Disk Based backup failed due to the application virtual server failover. How do I clean up Veritas Storage Foundation for Windows cluster disk groups and their associated volumes?</p>	<p>If the application virtual server fails when you use Veritas Storage Foundation for Windows (SFW) snapshot provider to perform an advanced disk-based backup, the backup job will fail. The original cluster disk group that the snapped volumes belong to has moved from the primary node to a secondary node and the snapped volumes will not be able to resynchronize with the original volumes.</p> <p>The following is a description of the steps that occur for an advanced disk-based backup:</p> <ul style="list-style-type: none"> ■ The snapped volumes are split from the original volumes. ■ The previously split snapped volumes are placed into a new cluster disk group. ■ The new cluster disk group is removed from the physical node where the production virtual server is currently online and then added to the Symantec Backup Exec media server. ■ The new cluster disk group will eventually be removed from the media server and then added back into the physical node where it previously resided, regardless of where the production virtual server is currently located. ■ The new cluster disk group joins the original cluster disk group if it is located in the same node. ■ The snapped volumes resynchronize with the original volumes. <p>During this process if the production virtual fails over from the currently active node to a secondary node, the new cluster disk group cannot rejoin the original cluster disk group.</p> <p>See “Manually joining two cluster disk groups and resynchronizing volumes” on page 860.</p>
<p>After I performed a manual failover of a Veritas cluster resource, my backup jobs hang. Why won't the backup jobs terminate?</p>	<p>If a manual failover of a Veritas cluster resource occurs, Veritas Cluster Server does not dismount MountV resources if there are open handles. It is recommended that all backup jobs complete before performing a manual failover. If a backup job does hang, you must manually cancel the job before you can complete a manual cleanup process.</p>

Changing the Quorum disk signature

The cluster service may not start because the disk signature on the Quorum disk is different from the original signature. You can change the disk signature.

To change the Quorum disk signature

- 1 Start the cluster service on one node with the -Fixquorum option in the startup parameters.
- 2 Open the Cluster Administrator and right-click the cluster, and then select **Properties**.
- 3 Select the **Quorum** tab.
- 4 In the **Quorum resource** field, select a different disk.
- 5 Click **OK**.
- 6 Stop the cluster services and then restart them without the -Fixquorum option.
You may run the -Fixquorum option as many times as needed to redesignate a Quorum disk signature.
- 7 Bring all other nodes online.

Manually joining two cluster disk groups and resynchronizing volumes

If an Advanced Disk-based backup failed due to the application virtual server failover, you may need to re-join the cluster disk groups.

To manually re-join the two cluster disk groups and resynchronize the volumes

- 1 Import the cluster disk group into the node, if the original cluster disk group is not already imported into the node where the production virtual server is currently online.
- 2 Rejoin the new cluster disk group with the original cluster disk group.
- 3 Snap back the snapped volumes with their original volumes. Ensure that the option to synchronize using the original volume is selected.

If you are not able to import the new cluster disk group into the node where the original cluster disk group is currently located, failover the application virtual server back to its original node before rejoining the two cluster disk groups. For detailed instruction on how to perform SFW operations, consult the Veritas Storage Foundation for Windows user guide.

Using Backup Exec Retrieve

This chapter includes the following topics:

- [About Backup Exec Retrieve](#)
- [How Backup Exec Retrieve works](#)
- [What end users can do with Backup Exec Retrieve](#)
- [Before you install Backup Exec Retrieve](#)
- [Requirements for installing Backup Exec Retrieve on a Web server](#)
- [Requirements for using Backup Exec Retrieve on end users' computers](#)
- [Upgrading from Backup Exec Retrieve that runs under Backup Exec System Recovery Manager 8.5](#)
- [Installing Backup Exec Retrieve](#)
- [About configuring Backup Exec Retrieve](#)
- [Editing application settings for Backup Exec Retrieve](#)
- [Uninstalling Backup Exec Retrieve](#)
- [Troubleshooting Backup Exec Retrieve](#)

About Backup Exec Retrieve

Backup Exec Retrieve provides a Web-based method for end users to search, browse, preview, and retrieve archived files and email. The files and email must be in shared folders to which the end users have permission. End users save retrieved files or emails to a location that they specify. Backup Exec Retrieve is not designed to restore system databases or other system-level files. Likewise,

end users cannot delete, change, move, or rename files by using Backup Exec Retrieve.

You can configure Backup Exec Retrieve to let end users retrieve their own data. Backup Exec Retrieve works with the following data sources:

- Backup Exec Archiving Option
- Backup Exec Continuous Protection Server (CPS)
- Backup Exec Desktop and Laptop Option (DLO)
- Backup Exec System Recovery Manager

Backup Exec Retrieve uses Windows security and Internet browser download features and is limited to Windows server platforms.

How Backup Exec Retrieve works

The following table describes a typical use-case scenario for installing, configuring, and using Backup Exec Retrieve. It includes information for both the administrator and the end user .

End users log on with their domain credentials. The system restricts their access so that they can only retrieve the files to which they originally had access. For example, if a file server is protected, users likely only have access to the files that are located inside shared folders. Backup Exec Retrieve also allows Exchange email retrieval. In such cases, mailbox permissions and mailbox folder permissions control the access.

Table 20-1 How Backup Exec Retrieve works

Process order	Role	Task or Process
1	Administrator	Optional - Installs Backup Exec and Backup Exec Archiving Option on the Backup Exec media server.
2	Administrator	Installs Backup Exec Retrieve on a Web server or on the Backup Exec media server that is on a Web server.

Table 20-1 How Backup Exec Retrieve works (*continued*)

Process order	Role	Task or Process
3	Administrator	<p>Configures Backup Exec Retrieve with the locations of any of the following indexers that can be included in searches by end users:</p> <ul style="list-style-type: none"> ■ Media server (for the files and the emails that the Backup Exec Archiving Option archives). ■ Backup Exec System Recovery Manager server (for files in recovery points that Backup Exec System Recovery creates). ■ Continuous Management Service server (for the files that Continuous Protection Servers back up). ■ Desktop and Laptop Option server (for the files that the Backup Exec Desktop and Laptop Option backs up).
4	Administrator	<p>Notifies the end users of the following Backup Exec Retrieve Web address so that they can search their own files and emails.</p> <p>https://<Backup Exec Retrieve Web server name>/BERetrieve</p> <p>You must also provide the following information to the end users:</p> <ul style="list-style-type: none"> ■ If the standard Windows security alert screens are displayed to the user, inform the user to click Yes or OK to continue. ■ If a message appears to the user with information about the security certificate, it is related to SSL certificates. Have the user click Yes to continue.

Table 20-1 How Backup Exec Retrieve works (*continued*)

Process order	Role	Task or Process
5	End user	<ul style="list-style-type: none"> ■ Launches Backup Exec Retrieve using the Web address that the administrator supplies. ■ The browser may prompt the end user to download and install Microsoft Silverlight, which is a necessary component to use Backup Exec Retrieve. Your organization may not permit end users to download files from the Web. In such cases, the administrator must deploy Silverlight to the end users' computers before end users can use Backup Exec Retrieve. See "About deploying the Silverlight run time in your organization" on page 869. ■ If Silverlight is already installed on the user's computer, the browser immediately displays the Backup Exec Retrieve logon screen. ■ At the logon screen, the end users type their user name, password, and domain. The credentials that a user specifies here determines what data they are authorized to view and retrieve.
6	End user	Selects to search, browse, or view recent activity for files or email.
7	End user	Submits a query. Backup Exec Retrieve displays the results for the end user to page through. The user can click a related link to drill down into the information.
8	End user	Retrieves the selected file or email and saves it locally on their computer or elsewhere.

What end users can do with Backup Exec Retrieve

Backup Exec Retrieve lets end users search, browse, or retrieve their own files and emails from a Web-based user interface.

The following table describes the most frequently used tasks in Backup Exec Retrieve.

Table 20-2 What you can do with Backup Exec Retrieve

Task	Description
Basic search	<p>Lets end users find the following information:</p> <ul style="list-style-type: none"> ■ Files with text in the file name or in the content. Support is also included for file system wildcards. ■ Email messages with text in the Subject, Content, From and To fields.
Advanced search	<p>Lets end users find the following:</p> <ul style="list-style-type: none"> ■ Files based on file name, file content, file folder, or specified date range. ■ Email messages based on the text in the following locations: <ul style="list-style-type: none"> ■ Subject field ■ Subject and email contents ■ From and To fields ■ A date range
Recent activity	<p>Lets end users find recently archived, deleted, or edited files or email messages.</p>
Browse folders	<p>Lets end users navigate through shared folders to which they have permissions. Also lets them view the files and email that were backed up or archived.</p> <p>When a backed up file is located, the end user can view all stored versions of that file. Versioning does not apply to archived files or email messages. Review the date, time, and file size to determine the version of the file that you want to retrieve.</p>
Preview	<p>Lets end users do the following:</p> <ul style="list-style-type: none"> ■ Open an abbreviated preview of email messages or some files in the list of search results. ■ Preview an entire email message or some files in the list of search results.

Table 20-2 What you can do with Backup Exec Retrieve (*continued*)

Task	Description
Show versions	Lets end users view all versions of backed up files. (Does not apply to archived files or email messages.)
Retrieve files and emails	Lets end users retrieve and save a file or an email message by using the Web browser's Save As dialog box. Email messages are saved with a .Msg file extension. Users can then open the file in Microsoft Outlook.

Before you install Backup Exec Retrieve

Installation procedures might vary, depending on your work environment and how you want to install Backup Exec Retrieve. You must install the software on a Web server. If your Backup Exec media server is also a Web server, you can install Backup Exec Retrieve on that computer as well.

You can have multiple installations of Backup Exec in a network domain.

During the installation, you may be prompted to install the latest version of Microsoft's .NET Framework and Microsoft's Internet Information Services with ASP.NET.

You must have Administrator rights, or be able to use an account that has Administrator privileges, to install Backup Exec Retrieve.

Before you install the product, make sure that the computer meets the specified requirements.

See [“Requirements for installing Backup Exec Retrieve on a Web server”](#) on page 866.

See [“Installing Backup Exec Retrieve”](#) on page 869.

See [“Uninstalling Backup Exec Retrieve”](#) on page 876.

Requirements for installing Backup Exec Retrieve on a Web server

To install Backup Exec Retrieve, the computer on which you install the software must meet minimum requirements:

See [“Before you install Backup Exec Retrieve”](#) on page 866.

Table 20-3 Requirement for installing Backup Exec Retrieve on a Web server

Component	Requirement
Processor	Intel Pentium 4 CPU 2.0 GHz or faster
Operating system	<p>The following operating systems are supported:</p> <ul style="list-style-type: none"> ■ Windows Server 2008 R2 ■ Windows Server 2008 R2 (x64) ■ Windows Server 2008 ■ Windows Server 2008 (x64) ■ Windows Server 2003 R2, SP 2 or later ■ Windows Server 2003 R2, SP 2 or later (x64) ■ Windows Server 2003, SP1 ■ Windows Server 2003, SP 1 (x64)
RAM	2 GB
Available disk space	45 MB
Software	<p>The following is installed on the computer on which you want to install Backup Exec Retrieve:</p> <ul style="list-style-type: none"> ■ Microsoft .NET Framework 3.5 SP1. At a minimum, Microsoft .NET Framework 3.5 SP 1 is required to run Backup Exec Retrieve. <p>Note: The latest version of the .NET Framework is automatically installed if it does not already exist or if an older version exists. This note applies only when you install Backup Exec Retrieve for the first time.</p> <ul style="list-style-type: none"> ■ Microsoft Internet Information Services (IIS) 6.0 or later. ■ Microsoft ASP.NET

Requirements for using Backup Exec Retrieve on end users' computers

For end users to use Backup Exec Retrieve, client computers must meet the following requirements:

Table 20-4 Requirements for using Backup Exec Retrieve on end users' computers

Component	Requirement
Operating system	<p>The following operating systems are supported:</p> <ul style="list-style-type: none"> ■ Windows XP SP2 or later ■ Windows Vista
Software	<p>The following software is required:</p> <ul style="list-style-type: none"> ■ Microsoft Internet Explorer 7.0 or later. Other browsers should also be compatible, but certain features and user interface layouts may vary. ■ Microsoft Silverlight 3.0 or later plug-in If Silverlight is not detected when a user accesses the Backup Exec Retrieve Web page, the browser offers to install it. Your organization may not permit end users to download files from the Web. In such cases, the administrator must deploy Silverlight to the end users' computers before end users can use Backup Exec Retrieve. See "About deploying the Silverlight run time in your organization" on page 869.
Internet	Internet access is required.
Web address	<p>End users must know the following Web server address, and their own user name and password to access Backup Exec Retrieve.</p> <p>https://<Backup Exec Retrieve Web server name>/BERetrieve</p> <p>If the standard Windows security alert screens are displayed to the user, inform the user to click Yes or OK to continue.</p> <p>If a message appears to the user with information about the security certificate, it is related to SSL certificates. Have the user click Continue to this Web site to continue.</p>

About deploying the Silverlight run time in your organization

If your organization does not permit end users to download files from the Web, you may need to deploy Silverlight to your organization's users' computers. Silverlight is a necessary component to use Backup Exec Retrieve.

Functioning as an administrator, you can download the latest version of Silverlight from the following Web site:

<http://www.microsoft.com/silverlight/downloads.aspx>

You can deploy Silverlight across your network by using any one of the following methods:

- Windows software update services
- Microsoft System Center Configuration Manager
- Group policy

For guidance on deploying the Silverlight run time in your own organization, see the following:

<http://www.microsoft.com/silverlight/resources/technical-resources/>

Upgrading from Backup Exec Retrieve that runs under Backup Exec System Recovery Manager 8.5

This version of Backup Exec Retrieve replaces instances of Backup Exec Retrieve that run under Backup Exec System Recovery Manager 8.5. You should uninstall your existing version of Backup Exec Retrieve before you install the current version.

Use the Microsoft Windows **Add or Remove Programs** utility to uninstall older versions of Backup Exec Retrieve.

Installing Backup Exec Retrieve

You install Backup Exec Retrieve from the Backup Exec installation media browser. Following the installation, you must configure Backup Exec Retrieve by adding the data sources that end users can search for their files and email.

See “[About configuring Backup Exec Retrieve](#)” on page 871.

To install Backup Exec Retrieve

- 1 Log on to your Web server (or the Backup Exec media server if it has Web server capabilities).

You must use either the Administrator account or an account with administrator privileges.

- 2 Insert the Backup Exec installation media into the appropriate drive of the computer.
- 3 Do one of the following:

If the installation starts automatically Continue to the next step.

If the installation does not start automatically

Do the following in the order listed:

- On the Windows desktop, click **Start > Run**.
- Type: <media drive letter>:\Setup.exe.
For example, e:\setup.exe.
- Continue with next step.

- 4 Click **Backup Exec Retrieve 2010**.
- 5 In the **Welcome** panel, click **Next**.
- 6 In the **License** panel, read the license agreement, and then click **I accept the terms of the license agreement**.
- 7 Click **Next**.
- 8 In the **Environment Check** panel, review the results of the Environment Check . For each requirement, the following results are displayed:

Check mark

The requirement and the recommendations are met.

X

The requirement is not met. You cannot continue with the installation until the requirement is met.

Click the associated link for additional information.

- 9 Click **Next**.

Backup Exec Retrieve program files are installed in the default path that appears on the **Options** page.

10 In the **Destination** panel, do one of the following:

To change the folder where the Backup Exec Retrieve files are installed	Click Change to select a new folder. The default path is C:\Program Files\Symantec\Backup Exec Retrieve\
To accept the default	Continue to the next step.

11 Click **Next** to begin the installation.

During the installation of Backup Exec Retrieve, you may be prompted to install Symantec LiveUpdate or Microsoft .NET Framework 3.5. In such cases, follow the on-screen prompts to complete those specific installations.

12 In the **Complete** panel, click **Finish**.

About configuring Backup Exec Retrieve

To allow end users to search and retrieve their own files and email, you must first configure Backup Exec Retrieve. During the configuration, you point to repositories of user data that are available in the enterprise. These repositories are known as data sources. These data sources contain backup copies of files or archives of files and email messages.

For example, you can add a Backup Exec media server on which the Archiving Option is installed, Continuous Protection Servers, or Backup Exec System Recovery Manager servers. You add these various data sources to Backup Exec Retrieve which in turn provide retrieval data to end users.

When you delete a data source, end users can no longer search that repository for their data using Backup Exec Retrieve.

You must have local administrator rights to add, edit, or delete data sources from the Backup Exec Retrieve console.

See [“Adding a data source”](#) on page 871.

See [“Editing a data source”](#) on page 873.

See [“Deleting a data source”](#) on page 874.

Adding a data source

You can add the following items to the list of available data sources that hold the end user's files and email:

- Backup Exec media server on which the Archiving Option is installed

- Continuous Protection Servers
- Desktop and Laptop Option servers
- Backup Exec System Recovery Manager servers

You must have local administrator rights to add, edit, or delete data sources from Backup Exec Retrieve.

See [“About configuring Backup Exec Retrieve”](#) on page 871.

To add a data source

- 1 Log on to the Backup Exec Retrieve Web server as a local administrator.
- 2 On the **Windows Start** menu, click **All Programs > Symantec Backup Exec Retrieve > Backup Exec Retrieve Configuration Console**.
- 3 In the **Symantec Backup Exec Retrieve Configuration** dialog box, click **Add**.
- 4 Set the **Add Data Source** options.
See [“Add or edit data source options”](#) on page 872.
- 5 Click **OK**.

Add or edit data source options

When you add or edit a data source, you must specify account credentials to ensure proper authentication to the server.

See [“Adding a data source”](#) on page 871.

See [“Editing a data source”](#) on page 873.

Table 20-5 Add or edit data source options

Item	Description
Data Type	Identifies the data source type that you want Backup Exec Retrieve to connect to.
Name or IP address	Indicates the name or the IP address of the data source that you want Backup Exec Retrieve to connect to.
User name	Indicates the user name for an account that has authorized access to this data source. This option does not apply to the Archiving Option data type.

Table 20-5 Add or edit data source options (*continued*)

Item	Description
Password	Indicates the password for this account. The software encrypts the password but does not display it. This option does not apply to the Archiving Option data type.
Confirm password	Confirms the password for this account. This option does not apply to the Archiving Option data type.
Domain	Indicates the domain name for this account (if applicable to the type of data source you selected). This option does not apply to the Archiving Option or the Backup Exec System Recovery Manager data types.

Editing a data source

You can edit the configuration settings of each data source that you have added to Backup Exec Retrieve.

You must have local administrator rights to add, edit, or delete data sources from the Backup Exec Retrieve console.

See [“About configuring Backup Exec Retrieve”](#) on page 871.

To edit a data source

- 1 Log on to the Backup Exec Retrieve Web server as a local administrator.
- 2 On the **Windows Start** menu, click **All Programs > Symantec Backup Exec Retrieve > Backup Exec Retrieve Configuration Console**.
- 3 In the **Symantec Backup Exec Retrieve Configuration** dialog box, select the data source that you want to change.
- 4 Click **Edit**.
- 5 Set the **Edit Data Source** options.
See [“Add or edit data source options”](#) on page 872.
- 6 Click **OK**.

Deleting a data source

When you delete a data source, indexing of any new files or email messages continues to occur. However, Backup Exec Retrieve no longer searches the indexes. As a result, end users cannot search for those new files or emails.

You must have local administrator rights to add, edit, or delete data sources from the Backup Exec Retrieve console.

See [“About configuring Backup Exec Retrieve”](#) on page 871.

To delete a data source

- 1 Log on to the Backup Exec Retrieve Web server as a local administrator.
- 2 On the **Windows Start** menu, click **All Programs > Symantec Backup Exec Retrieve > Backup Exec Retrieve Configuration Console**.
- 3 In the **Symantec Backup Exec Retrieve Configuration** dialog box, select a data source name or type in the table.
- 4 Click **Delete**.
- 5 Click **Yes** to confirm the deletion of the data source.

Editing application settings for Backup Exec Retrieve

You can use the application settings for Backup Exec Retrieve that are set during installation, or you can edit the application settings. You can also launch the Backup Exec Retrieve configuration console from within Backup Exec.

To edit application settings for Backup Exec Retrieve

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Settings**, click **Backup Exec Retrieve**.
- 3 Select the appropriate settings.
See [“Application settings for Backup Exec Retrieve”](#) on page 874.
- 4 Click **OK**.

Application settings for Backup Exec Retrieve

Backup Retrieve must be installed and configured before you can set it up to work with Backup Exec. You can configure Backup Exec Retrieve and Backup Exec to allow end users to retrieve the data that was backed up with the Archiving Option. Backup Exec Retrieve works with the following data sources:

- Backup Exec Archiving Option

- Backup Exec Continuous Protection Server (CPS)
- Backup Exec Desktop and Laptop Option (DLO)
- Backup Exec System Recovery Manager

See [“Editing application settings for Backup Exec Retrieve”](#) on page 874.

Table 20-6 Application settings for Backup Exec Retrieve

Item	Description
Enable Backup Exec Retrieve to let end users retrieve their data	<p>Enables or disables Backup Exec Retrieve to work with Backup Exec.</p> <p>For the Backup Exec Archiving Option, check this checkbox and provide the name of the Web server on which Backup Exec Retrieve is installed and configured.</p> <p>See “How Archiving Option end users retrieve archived data by using Backup Exec Retrieve” on page 1432.</p> <p>If you clear this check box later, then all of the existing links to the Backup Exec Retrieve URL in the archived folders and mailboxes are removed.</p>
Backup Exec Retrieve Web server	Type in the name of the server on which Backup Exec Retrieve is installed and configured. A URL is created from this server name, and full URL is listed under Backup Exec Retrieve URL for user retrieval . End users can access the URL with their Web browsers and logon using the appropriate permissions to browse and retrieve their data.
Add Data Sources	Lets you configure the Backup Exec Retrieve Web server for other data sources. This option launches the Backup Exec Retrieve Configuration console. For example, if you want to allow CPS and DLO end users to retrieve their own data, you can launch this console and add them as data sources. (Appropriate credentials are required.)
Automatically add this media server as an Archiving Option data source	Adds this Backup Exec media server as a data source for the Backup Exec Archiving Option. You can also use Add Data Sources to manually add this media server or another computer as the Backup Exec Archiving Option data source.
Backup Exec Retrieve URL for user retrieval	Lists the URL for the current Backup Exec Retrieve Web server. Provide this URL to the end users so they can retrieve their data. End users see this URL next to their email messages that have been archived from Microsoft Outlook.

Table 20-6 Application settings for Backup Exec Retrieve (*continued*)

Item	Description
Copy to Clipboard	Copies the Backup Exec Retrieve Web server URL to your clipboard. You can paste the URL in an email message notifying the end users about retrieving their data with Backup Exec Retrieve.

Uninstalling Backup Exec Retrieve

You can use the Microsoft Windows **Add and Remove** utility to uninstall Backup Exec Retrieve.

To uninstall Backup Exec Retrieve

- 1 From the Windows server on which Backup Exec Retrieve is installed, click **Start**, point to the **Control Panel**, and then click **Add or Remove Programs**.
- 2 Click **Backup Exec Retrieve**, and then click **Remove**.
- 3 Click **Yes** to proceed with the uninstall.

Troubleshooting Backup Exec Retrieve

For help to resolve the problems that you might encounter with Backup Exec Retrieve, you can check the following information.

Table 20-7 Troubleshooting Backup Exec Retrieve

Problem	Description
Users get a secure sockets layer (SSL) certificate warning in their browser when they go to the Backup Exec Retrieve Web site.	Backup Exec Retrieve uses a Secure Sockets Layer certificate to protect communication between the client and server. Web browsers may warn users of problems with the site's security certificate. You can continue to the Web site. For more information on resolving the warning, click http://entsupport.symantec.com/umi/V-367-2-1 .

Table 20-7 Troubleshooting Backup Exec Retrieve (*continued*)

Problem	Description
<p>Selecting a file causes an error in the application that launches for viewing the file</p>	<p>When you select a file in Microsoft Internet Explorer, the file is saved to the Internet cache. Then, the program that is required to open it is launched. If the Do not save encrypted pages to disk option is selected in Internet Explorer, then the file is not saved to the Internet cache. However, the associated application is launched but it cannot open the file. To fix this issue, deselect the option in Internet Explorer (the option is the default on Windows 2003).</p> <p>To deselect the option to save encrypted pages to disk</p> <ul style="list-style-type: none"> ■ In Internet Explorer, click the Tools menu, and then click Internet Options. ■ Click Advanced, and then scroll down to the Security section. ■ Uncheck Do not save encrypted pages to disk. ■ Click Apply > OK.
<p>Users cannot log on with a local account</p>	<p>Backup Exec Retrieve usually runs on a separate computer from the original computer that was backed up. The only authority for local accounts and passwords is the original computer, which might not be available when you use Backup Exec Retrieve. You must use a domain account to log on and retrieve files using Backup Exec Retrieve. Backup Exec Retrieve makes every effort to capture the domain users and domain groups that are part of local groups. Therefore, if your domain account is a member of the Local Administrators group on your computer, you can retrieve your files.</p>
<p>Users cannot see all of my data (indexing take too long)</p>	<p>Indexing is a resource-intensive process. Performance expectations vary significantly depending on your hardware, network configuration, and data change rate. Initial indexing for a new storage location that contains many base recovery points is the most intense load, in the range of minutes per image. A smaller load for subsequent incremental recovery points is in the range of seconds per incremental. Recovery points, archives, and files that not indexed do not display in search results. If you find that indexing takes too long, you might need to distribute the load across additional Indexing Servers.</p>

Table 20-7 Troubleshooting Backup Exec Retrieve (*continued*)

Problem	Description
<p>Unable to download a file when you use the computer name or IP address of Backup Exec Retrieve in Internet Explorer</p>	<p>If an end user cannot download a file from Backup Exec Retrieve in Internet Explorer, enable automatic downloads in the browser.</p> <p>To enable automatic downloads in Internet Explorer</p> <ul style="list-style-type: none"> ■ In Internet Explorer, click Tools > Internet Options. ■ In the Security tab, click Custom Level. ■ In the Security Settings - Internet Zone page, scroll to Downloads > Automatic prompting for downloads. ■ Click Enable. ■ Click OK, and then click Yes to confirm the change. ■ Click OK to return to Internet Explorer.
<p>Clicking the Backup Exec Retrieve in Internet Explorer results in a prompt to add the site to the trusted sites list.</p>	<p>If Enhanced Security is enabled in Windows, you are prompted to add the Backup Exec Retrieve URL to Internet Explorer's trusted sites list. If you continue without adding the URL to the trusted sites list, you are prompted to install Silverlight, even if it already installed. If you try to install Silverlight again, the installation fails. Symantec recommends that you add the Backup Exec Retrieve URL to Internet Explorer's trusted sites list.</p> <p>To add the Backup Exec Retrieve URL to Internet Explorer's trusted sites list:</p> <ul style="list-style-type: none"> ■ In Internet Explorer, click Tools > Internet Options. ■ In the Security tab, click Trusted sites. ■ Click Sites, and then on the Trusted sites page, add the Backup Exec Retrieve URL. ■ Click Add, then click Close. ■ Click OK to return to Internet Explorer.

Symantec Backup Exec Active Directory Recovery Agent

This appendix includes the following topics:

- [About the Active Directory Recovery Agent](#)
- [Requirements for the Active Directory Recovery Agent](#)
- [About installing the Active Directory Recovery Agent](#)
- [How the Active Directory Recovery Agent works](#)
- [How Granular Recovery Technology works with Active Directory and ADAM/AD LDS backups](#)
- [Editing defaults for Active Directory and ADAM/AD LDS backup and restore jobs](#)
- [Backing up Active Directory](#)
- [Backing up ADAM/AD LDS](#)
- [Active Directory Recovery Agent backup job options](#)
- [About restoring individual Active Directory and ADAM/AD LDS objects](#)
- [About recreating purged Active Directory and ADAM/AD LDS objects](#)
- [Resetting the Active Directory computer object and the computer object account](#)

About the Active Directory Recovery Agent

The Symantec Backup Exec 2010 Active Directory Recovery Agent (ADRA) is installed as a separate, add-on component of Backup Exec 2010.

With ADRA you can use Granular Recovery Technology (GRT) to restore individual Active Directory objects and attributes without performing an authoritative or non-authoritative full restore. You can also restore individual Active Directory Application Mode (ADAM) and Active Directory Lightweight Directory Services (AD LDS) objects and attributes.

See [“Requirements for the Active Directory Recovery Agent”](#) on page 880.

See [“About installing the Active Directory Recovery Agent ”](#) on page 881.

See [“How the Active Directory Recovery Agent works”](#) on page 882.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 888.

Requirements for the Active Directory Recovery Agent

Review the following requirements before you restore individual objects and attributes using the Active Directory Recovery Agent:

- You must have a full mode backup of the ADAM/AD LDS or the Windows System State (where Active Directory is installed).
- You must use one of the following Windows operating systems on the computer where Active Directory is in use:
 - Windows XP Professional x64 Edition
 - Windows 2000 Server with Service Pack 4.
ADRA does not support reanimation of objects from the Active Directory Deleted Objects container on a Windows 2000 domain controller. Symantec recommends that you use the Remote Agent on a Windows 2003 domain controller to run GRT restore jobs of deleted objects. Deleted objects can only be restored using an agent on a Windows 2000 domain controller if the Recreate deleted object check box is checked. This check box is located on the Restore Job Properties dialog box after you select the Microsoft Active Directory node under Settings.
 - Windows Server 2003 with Service Pack 1 or later
 - Windows Server 2003 R2
 - Windows Server 2008
 - Windows Server 2008 R2

- You must use a version of the Windows operating system that supports minifilter drivers on the media server that runs the restore job. Minifilter drivers are supported in the following Windows operating systems:
 - Windows 2000 with both Service Pack 4 and the Windows 2000 Rollup Patch 1 installed
 - Windows Server 2003 with Service Pack 1 or later installed.
 - Windows Server 2003 R2
 - Windows Server 2008
 - Windows Server 2008 R2
- You must run the Backup Exec Remote Agent for Windows Systems on the computer where Active Directory is installed.
- You must designate a location on the media server disk where Backup Exec can temporarily place the objects and attributes that are being restored when you restore from tape.
- Make sure you select the option Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from Active Directory backups. Individual attributes and properties cannot be restored from full Active Directory and ADAM/AD LDS backups if you do not select this option during backup.

Note: You cannot restore individual objects and attributes from Active Directory backups for a Read-Only Domain Controller (RODC). You should do GRT backups and restores of the Active Directory to a writable centralized datacenter domain controller.

See [“About installing the Active Directory Recovery Agent ”](#) on page 881.

See [“How the Active Directory Recovery Agent works”](#) on page 882.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 888.

About installing the Active Directory Recovery Agent

ADRA is installed locally as a separate, add-on component of Backup Exec 2010.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

How the Active Directory Recovery Agent works

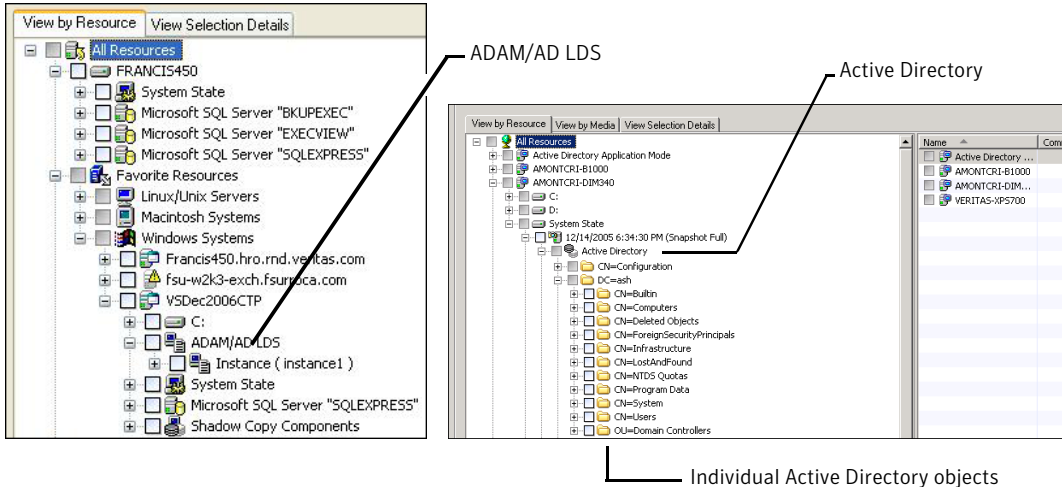
ADRA works with backups of the Windows System State (where Active Directory is installed) and ADAM/AD LDS.

When you back up the Windows System State, the Active Directory is included in the backup job, because Active Directory is a component of Windows System State.

You can also use ADRA to restore individual ADAM/AD LDS objects and attributes. If multiple ADAM/AD LDS instances are backed up, each instance appears under the Active Directory Application Mode node.

The following figure shows ADAM/AD LDS and Active Directory.

Figure A-1 View by Resource view - ADAM/AD LDS and Active Directory



ADRA also lets you restore tombstoned objects from the Active Directory Deleted Objects container in the following situations:

- Their tombstone lifetimes have not passed.
- They have not been purged from the Deleted Objects container.
- You are restoring to a Windows Server 2003/2008/2008 R2/XP Professional x64 Edition system.

Symantec recommends that Active directory and ADAM/AD LDS backups be backed up to a backup-to-disk folder before you back them up to tape. This strategy provides you with shorter backup windows. It also lets you administer Active Directory or ADAM/AD LDS without requiring the individual cataloging of the backed up objects and properties.

When you back up any Windows Active Directory or ADAM/AD LDS application database directly to tape, objects and properties that are added or deleted during the backup will not match the individual objects and properties that are available for restore from the backup set. The back up of the database is a snapshot backup of the live Active Directory or ADAM/AD LDS database and the cataloging of the individual Active Directory or ADAM/AD LDS objects occurs after the snapshot is performed. Since the catalog operation catalogs objects and properties from the live Active Directory or ADAM/AD LDS database, object and property changes can occur after the snapshot was taken.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 888.

See [“About recreating purged Active Directory and ADAM/AD LDS objects”](#) on page 892.

How Granular Recovery Technology works with Active Directory and ADAM/AD LDS backups

Granular Recovery Technology (GRT) lets you restore individual objects and attributes from Active Directory and ADAM/AD LDS backups without performing an authoritative or non-authoritative full restore. To restore individual items, you must enable the Granular Recovery Technology feature when you create a backup job. You should review the requirements for a GRT-enabled backup before you configure it.

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“Recommended devices for backups that use Granular Recovery Technology”](#) on page 324.

See [“About requirements for jobs that use Granular Recovery Technology”](#) on page 325.

Editing defaults for Active Directory and ADAM/AD LDS backup and restore jobs

You can edit the default settings for all Active Directory and ADAM/AD LDS backup and restore jobs. You also can override these defaults when you set up Active Directory and ADAM/AD LDS backup and restore jobs.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 888.

See [“About recreating purged Active Directory and ADAM/AD LDS objects”](#) on page 892.

To edit defaults for Active Directory and ADAM/AD LDS backup and restore jobs

- 1 On the **Tools** menu, click **Options**.
- 2 In the task pane, under **Job Defaults**, click **Microsoft Active Directory**.
- 3 Select the default backup and restore options for the Active Directory Recovery Agent.

See [“Microsoft Active Directory default options”](#) on page 884.

- 4 Click **OK**.

Microsoft Active Directory default options

You can edit the default settings for Active Directory and ADAM/AD LDS backup and restore jobs.

See [“Backing up Active Directory”](#) on page 885.

See [“Backing up ADAM/AD LDS”](#) on page 886.

See [“Editing defaults for Active Directory and ADAM/AD LDS backup and restore jobs”](#) on page 883.

Table A-1 Microsoft Active Directory default options

Item	Description
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from Active Directory backups (not supported for Read-Only Domain Controllers)	Enables the restore of individual items from full backups of the Active Directory or ADAM/ AD LDS. Ensure that you meet the requirements for Granular Recovery Technology. See “About requirements for jobs that use Granular Recovery Technology” on page 325.
Perform consistency check before backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider (Windows Server 2008)	Checks snapshots for data corruption. This option applies only to snapshots that are performed by the Microsoft Volume Shadow Copy Services (VSS).

Table A-1 Microsoft Active Directory default options (*continued*)

Item	Description
Continue with backup if consistency check fails	Enables the backup job to continue even if the consistency check fails. You may want the job to continue if a backup of the database in its current state is better than no backup at all. Or you may want the job to continue if you back up a large database that may have only a small problem.
Recreate deleted objects that cannot be restored from the Active Directory Deleted Objects container	<p>Attempts to recreate deleted objects if both of the following events occurred:</p> <ul style="list-style-type: none"> ■ The objects' tombstone lifetimes have passed. ■ The objects were purged from the Active Directory Deleted Objects container. <p>You must use this option to restore deleted objects on a computer that runs Windows 2000.</p> <p>See “About recreating purged Active Directory and ADAM/AD LDS objects” on page 892.</p>

Backing up Active Directory

Use the following steps to back up Active Directory.

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux Servers is installed.

To backup Active Directory

- 1 On the navigation bar, click the down arrow next to **Backup**.
- 2 Click **New Backup Job**
- 3 On the **View by Resources** tab, under **All Resources**, expand the name of the computer that contains the Active Directory that you want to back up.
- 4 Click **System State**.
- 5 In the task pane, under **Settings**, click **Microsoft Active Directory**.

- 6 Select the backup options you want to use.
See “[Active Directory Recovery Agent backup job options](#)” on page 887.

- 7 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

Backing up ADAM/AD LDS

Use the following steps to back up ADAM/AD LDS.

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux Servers is installed.

To back up ADAM/AD LDS

- 1 On the navigation bar, click the down arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **View by Resources** tab, expand **Favorite Resources**.
- 4 Expand **Windows Systems**.
- 5 Expand the name of the computer where ADAM/AD LDS is installed.
- 6 Select the backup options you want to use.
See “[Active Directory Recovery Agent backup job options](#)” on page 887.
- 7 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

Active Directory Recovery Agent backup job options

Select the appropriate Active Directory Recovery Agent backup job options.

See “[Backing up Active Directory](#)” on page 885.

See “[Backing up ADAM/AD LDS](#)” on page 886.

Table A-2 Options for Active Directory Recovery Agent backup jobs

Item	Description
Full - Back up entire Active Directory	Due to limitations within Microsoft Active Directory, only the full backup method can be used to back up Active Directory. Incremental and differential backup methods are not available.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from Active Directory backups (not supported for Read-Only Domain Controllers)	Enables the restore of individual items from full backups of the Active Directory or ADAM/ AD LDS. Ensure that you meet the requirements for Granular Recovery Technology. See “ About requirements for jobs that use Granular Recovery Technology ” on page 325.
Perform consistency check before backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider (Windows Server 2008)	Checks snapshots for data corruption. This option applies only to snapshots that are performed by the Microsoft Volume Shadow Copy Services (VSS). If corrupt data is found and this option is not selected, the job fails.

Table A-2 Options for Active Directory Recovery Agent backup jobs (*continued*)

Item	Description
Continue with backup if consistency check fails	Continues the backup job even if the consistency check fails. You may want the job to continue if a backup of the database in its current state is better than no backup at all. Or you may want the job to continue if you back up a large database that may have only a small problem.

About restoring individual Active Directory and ADAM/AD LDS objects

Before starting the restore job, you should review information on finding and viewing specific data to restore, as well as on details on restore options and restore jobs.

See [“About restoring data”](#) on page 599.

When you restore Active Directory and ADAM/AD LDS objects from tape, you must specify an on-disk staging location where the objects will be placed prior to being restored. The staging location must be a path on a local NTFS volume on the media server running the restore job and the Backup Exec service account must also have access to it.

Note: If you previously defined a default staging location in the option, **Path on an NTFS volume that is local to the media server for temporary storage of restore data** under **Tools > Options > Restore**, you can override the default by specifying an alternate staging location for each Active Directory and ADAM/AD LDS restore job by entering a path in the Advanced node found under **Settings** on the **Restore Job Properties** pane.

System volumes should not be used as a staging location because of the potentially large file sizes that are created on the disk specified in the staging location path.

Because restoring objects from tape requires the creation of a staging location, restoring from tape requires more time than if you are restoring from disk.

By default, ADRA restores deleted Active Directory or ADAM/AD LDS objects from the Active Directory Deleted Objects container if their tombstone lifetimes have not passed.

When objects in Active Directory are deleted, they are removed from their current Active Directory or ADAM/AD LDS container, converted into tombstones, and then placed in the Active Directory Deleted Objects container where their tombstone lifetime is monitored. After their tombstone lifetime passes, the tombstones are purged from the Active Directory Deleted Objects container, which permanently deletes the objects from the Active Directory and ADAM/AD LDS databases.

Following are requirements for backup and restore operations when an Active Directory or ADAM/AD LDS backup is enabled for the restore of individual items:

Table A-3 Requirements for backup and restore operations for Active Directory or ADAM/AD LDS

Item	Description
If the destination device for the backup job is a backup-to-disk folder	Backup-to-disk folders provide the most efficient method of storage for GRT-enabled backups. You must create a temporary hard disk staging location on a local NTFS volume to restore individual items from GRT-enabled backups on tape. The data is first copied from tape to the temporary staging location before it can be restored. As such, a restore from tape takes more time. For best results, you should specifically select the backup-to-disk folder you want to use for your GRT-enabled backup jobs when you set them up.
If you create full backups	The full job templates must be in a policy, and must have a backup-to-disk folder as the destination device. If you run only a full backup of the Active Directory or ADAM/AD LDS, the full job template does not have to be in a policy. See “Creating a new policy” on page 522.
If you restore individual items from an Active Directory or ADAM/AD LDS backup set that is on a device other than a backup-to-disk folder	Backup Exec must temporarily stage the entire database to a path on an NTFS volume on the media server to extract individual items. You must specify this path.

When restoring Active Directory user objects, you must reset the object’s user password and then re-enable the object’s user account. For ADAM/AD LDS user objects, you must reset the object’s user password and then re-enable the object’s user account. For Active Directory user objects, use the Microsoft Active Directory Users and Computers application. For ADAM/AD LDS user objects, use ADSI Edit. For Active Directory computer objects, you must reset the object’s account.

See [“Resetting the Active Directory computer object and the computer object account”](#) on page 895.

ADRA does not support reanimation of objects from the Active Directory Deleted Objects container on a Windows 2000 domain controller. It is recommended that individual restores of deleted objects be done by a Backup Exec Remote Agent on a Windows 2003 domain controller, if one exists in the same domain. If a Windows 2003 domain controller is not available in the domain, deleted objects can only be restored using an agent on a Windows 2000 domain controller if the Recreate deleted object check box is checked.

Note: Some objects in the Active Directory Configuration Partition node cannot be reanimated from the Active Directory Deleted Objects container. However, recreated objects may not be recognized by some applications.

For more information, see your Microsoft Active Directory documentation.

See [“About inventoring media”](#) on page 448.

See [“Creating a new catalog”](#) on page 248.

See [“Restoring individual objects from an Active Directory backup”](#) on page 890.

See [“About recreating purged Active Directory and ADAM/AD LDS objects”](#) on page 892.

See [“Restoring individual objects from an ADAM/AD LDS backup”](#) on page 891.

See [“Resetting the Active Directory computer object and the computer object account”](#) on page 895.

Restoring individual objects from an Active Directory backup

Use ADRA to restore individual objects from Active Directory.

See [“Resetting the Active Directory computer object and the computer object account”](#) on page 895.

See [“Restoring individual objects from an ADAM/AD LDS backup”](#) on page 891.

See [“About recreating purged Active Directory and ADAM/AD LDS objects”](#) on page 892.

To restore individual objects from an Active Directory backup

- 1 On the navigation bar, click the down arrow next to **Restore**.
- 2 Click **New Restore Job**.

- 3 In the **View by Resource** pane, double-click a backup set that contains the most recent System State backup.

If you want to restore Active Directory objects from a previous backup, select the appropriate backup set.
- 4 Double-click **System State**.
- 5 Double-click the most recent System State snapshot.
- 6 Double-click **Active Directory**.
- 7 In the **Results** pane, select the appropriate object or objects.
- 8 If you are restoring from tape, do the following:
 - On the task pane, under **Settings**, click **Advanced**.
 - If you have not set a default temporary staging location, type a path in the box titled **Path on an NTFS volume that is local to the media server for temporary storage of restore data**.
- 9 Click **Run Now** to start the restore job or select other restore options from the task pane.

Any Active Directory or ADAM/AD LDS object or property that is selected for restore will overwrite existing objects and properties, even if you selected **Skip if file exists** or **Overwrite the file on disk only if it is older** on the **General Restore Job Properties** dialog box.
- 10 If you restored a a deleted user object, use the Microsoft Active Directory Users and Computers application to reset the object's user password and re-enable the object's user account. If you restored a computer object, you must reset the account for it.

Restoring individual objects from an ADAM/AD LDS backup

Use ADRA to restore individual objects from ADAM/AD LDS.

See [“Restoring individual objects from an Active Directory backup”](#) on page 890.

See [“About recreating purged Active Directory and ADAM/AD LDS objects”](#) on page 892.

To restore individual objects from an ADAM/AD LDS backup

- 1 On the navigation bar, click the down arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** pane, double-click **Active Directory Application Mode**.

- 4 Double-click the appropriate ADAM/AD LDS instance.
- 5 Double-click the appropriate backup set.
- 6 In the **Results** pane, select the appropriate object or objects.
- 7 If you are restoring from tape, do the following:
 - On the task pane, under **Settings**, click **Advanced**.
 - If you have not set a default temporary staging location, type a path in the box titled **Path on an NTFS volume that is local to the media server for temporary storage of restore data**.
- 8 Click **Run Now** to start the restore job or select other restore options from the **Properties** pane.

Any Active Directory or ADAM/AD LDS object or property that is selected for restore will overwrite existing objects and properties, even if you selected **Skip if file exists** or **Overwrite the file on disk only if it is older** on the **General Restore Job Properties** dialog box.
- 9 If you restored a deleted user object, use the ADSI Edit application to reset the object's user password and re-enable the object's user account.

About recreating purged Active Directory and ADAM/AD LDS objects

You can attempt to recreate deleted objects if their tombstone lifetimes have passed and the objects have been purged from the Active Directory Deleted Objects container.

However, you should be aware of the following:

- Most applications will not recognize a recreated object since recreated objects are not identical to the original deleted object. Recreated objects are assigned new global unique identifiers (GUIDs) and security identifiers (SIDs) that cannot be identified by the applications that created the original object.
- Attributes created by the Windows operating system cannot be recreated when a purged object is recreated. Hence, objects that rely on attributes set by the operating system will not be recognized by Windows when the objects are recreated.

See [“Recreating purged Active Directory objects”](#) on page 893.

See [“Recreating purged ADAM/AD LDS objects”](#) on page 894.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 888.

Recreating purged Active Directory objects

You can attempt to recreate deleted Active Directory objects after they have been purged from the Active Directory Deleted Objects container by restoring the object from a prior Active Directory backup.

See [“Recreating purged ADAM/AD LDS objects”](#) on page 894.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 888.

See [“Resetting the Active Directory computer object and the computer object account”](#) on page 895.

To recreate purged Active Directory objects

- 1 On the navigation bar, click the down arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** pane, double-click a backup set that contains the most recent System State backup.
If you want to restore Active Directory objects from a previous backup, select the appropriate backup set.
- 4 Double-click **System State**.
- 5 Double-click the most recent System State snapshot.
- 6 Double-click **Active Directory**.
- 7 In the **Results** pane, select the appropriate objects.
- 8 On the task pane, under **Settings**, click **Microsoft Active Directory**.
- 9 Check **Recreate deleted objects that cannot be restored from the Active Directory Deleted Objects container**.
- 10 If you are restoring from tape, do the following:
 - On the task pane, under **Settings**, click **Advanced**.
 - If you have not set a default temporary staging location, type a path in the box titled **Path on an NTFS volume that is local to the media server for temporary storage of restore data**.

- 11 Click **Run Now** to start the restore job or select other restore options from the **Properties** pane.

Any Active Directory or ADAM/AD LDS object or property that is selected for restore will overwrite existing objects and properties, even if you selected **Skip if file exists** or **Overwrite the file on disk only if it is older** on the **General Restore Job Properties** dialog box.

- 12 Use the Microsoft Active Directory Users and Computers application to reset the object's user password and re-enable the object's user account.

Recreating purged ADAM/AD LDS objects

You can attempt to recreate deleted ADAM/AD LDS objects after they have been purged from the Active Directory Deleted Objects container by restoring the object from a prior ADAM/AD LDS backup.

See [“Resetting the Active Directory computer object and the computer object account”](#) on page 895.

To recreate purged ADAM/AD LDS objects

- 1 On the navigation bar, click the down arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** pane, double-click **Active Directory Application Mode**.
- 4 Double-click the appropriate ADAM/AD LDS instance.
- 5 Double-click the appropriate backup set.
- 6 In the **Results** pane, select the appropriate objects.
- 7 On the task pane, under **Settings**, click **Microsoft Active Directory**.
- 8 Check **Recreate deleted objects that cannot be restored from the Active Directory Deleted Objects container**.
- 9 If you are restoring from tape, do the following:
 - On the task pane, under **Settings**, click **Advanced**.
 - If you have not set a default temporary staging location, type a path in the box titled **Path on an NTFS volume that is local to the media server for temporary storage of restore data**.

- 10 Click **Run Now** to start the restore job or select other restore options from the Properties pane.

Any Active Directory or ADAM/AD LDS object or property that is selected for restore will overwrite existing objects and properties, even if you selected Skip if file exists or Overwrite the file on disk only if it is older on the General Restore Job Properties dialog box.

- 11 Use the ADSI Edit application to reset the object's user password and re-enable the object's user account.

Resetting the Active Directory computer object and the computer object account

In Active Directory, computer objects are derived from user objects. Some attributes that are associated with a computer object cannot be restored when you restore a deleted computer object. The attributes can only be restored if the attributes were saved through schema changes before the computer object was originally deleted. Because computer object credentials change every 30 days, the credentials from the backup may not match the credentials that are stored on the actual computer.

Note: To reset a computer object, you must use the Microsoft Active Directory Users and Computers application.

For more information on resetting a computer object, see your Microsoft Active Directory Users and Computers application documentation.

If a computer object's **userAccountControl** attribute was not preserved before the object was deleted, you must reset the object's account after you restore the object.

See ["Recreating purged ADAM/AD LDS objects"](#) on page 894.

To reset the Active Directory computer object account

- 1 Remove the computer from the domain.
- 2 Re-join the computer to the domain. The SID for the computer remains the same since it is preserved when you delete a computer object. However, if the object's tombstone expires and a new computer object is recreated, the SID is different.

Symantec Backup Exec Advanced Disk-based Backup Option

This appendix includes the following topics:

- [About the Advanced Disk-based Backup Option](#)
- [About installing the Advanced Disk-based Backup Option](#)
- [About the synthetic backup feature](#)
- [What you can back up with synthetic backup](#)
- [Requirements for synthetic backup](#)
- [Methods for creating a synthetic backup](#)
- [About true image restore](#)
- [Enabling backups for true image restore](#)
- [About true image catalogs](#)
- [About restoring a backup set enabled for true image restore](#)
- [Selecting backup sets that are enabled for true image restore](#)
- [Troubleshooting tips for true image restore](#)
- [About offhost backup](#)
- [Configuring a GRT-enabled offhost backup for Exchange resources](#)
- [About restoring offhost backup data](#)

- [Troubleshooting the offhost backup](#)

About the Advanced Disk-based Backup Option

The Advanced Disk-based Backup Option (ADBO) is installed as a separate, add-on component of Backup Exec.

The Advanced Disk-based Backup Option provides the following features:

- **Synthetic backup.** This feature uses a policy to enable a full backup to be assembled, or synthesized, from a baseline and subsequent incremental backups that are also contained in a policy.

The benefits of using a synthetic backup include the following:

- **A reduced backup window** since the synthetic backup can be scheduled outside of the time-critical backup window.
- **Reduced network traffic** since the synthetic backup does not need to access the network.
- **True image restore.** This feature enables Backup Exec to restore the contents of directories to what they were at the time of any full or incremental backup. Restore selections are made from a view of the directories as they existed at the time of the particular backup. Files that were deleted before the time of the backup are not restored. In true image restore, only the correct versions of files are restored from the appropriate full or incremental backups that contain them. Previous versions are not unnecessarily restored and then overwritten.
- **Offhost backup.** This feature enables the backup operation to be processed on a Backup Exec media server instead of on the remote computer, or host computer. Moving the backup from the remote computer to a media server enables better backup performance and frees the remote computer as well.

See [“About the synthetic backup feature”](#) on page 899.

See [“About true image restore”](#) on page 912.

See [“About offhost backup”](#) on page 919.

About installing the Advanced Disk-based Backup Option

ADBO is enabled on the media server when you enter the ADBO license key.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

About the synthetic backup feature

The synthetic backup feature eliminates the need to perform recurring full backups for supported remote resources. A policy created for the synthetic backup feature enables the synthetic backup to be assembled from a full backup (called a baseline) and subsequent incremental backups that are also contained in the policy.


The resulting synthetic backup then becomes the new baseline, so only incremental backups are required until the next synthetic backup is created. The synthetic backup is as current as the last incremental backup that it contains.

The components of the policy for synthetic backup are as follows:

- **Baseline backup.** The first backup to run that is associated with the synthetic backup. The baseline backup runs one time only, and backs up all of the files on the selected resources when it runs.
- **Recurring incremental backups.** Subsequent backups that back up files that are changed after the baseline backup.
- **Recurring synthetic backups.** The process that combines the data from the baseline backup and the incremental backups to form a synthesized full backup of the selected resources. This synthesized full backup becomes a new baseline backup, which can then be combined with subsequent incremental backup sets to form the next synthesized full backup.


The baseline backup runs from the synthetic backup policy.

Figure B-1 Baseline backup




Incremental backups run from the synthetic backup policy.

Figure B-2 Incremental backup



The synthetic backup runs from the policy and is assembled from the baseline and incremental backups.

Figure B-3 Synthetic backup



Synthetic backup can only be created within a policy. You can use the Policy Wizard to create a policy that will contain the necessary job templates for the synthetic backup feature, or you can copy the example policy for synthetic backup and then modify it to suit your specific needs, or you can manually create a policy, and then add the necessary job templates.

For all of the associated backup templates in a policy, you can also use a **Duplicate Backup Set** template to create a multi-stage backup strategy for backing up data to disk and then copying it to tape.

See [“About duplicate backup set templates”](#) on page 548.

See [“Best practices for synthetic backup”](#) on page 902.

See [“Methods for creating a synthetic backup”](#) on page 904.

See [“About collecting additional information for synthetic backup and true image restore”](#) on page 904.

What you can back up with synthetic backup

Only file system resources are supported for synthetic backup.

Supported resources include common file system objects, such as volumes, drives, and folders. Do not include database resources or other unique resources in the selection list.

Backup Exec will not create synthetic backup jobs when the selection list that is associated with the synthetic backup policy contains unsupported resources.

See [“Requirements for synthetic backup”](#) on page 901.

See [“Methods for creating a synthetic backup”](#) on page 904.

Requirements for synthetic backup

Before you create a synthetic backup, review the following information:

- Synthetic backup and the associated templates can only be created in policies.
- In a policy that contains a synthetic backup, if an encryption key is used, all of the associated templates must use the same encryption key. The encryption key should not be changed after the policy has been created. The encryption key that is selected in the associated templates is automatically applied to the synthetic backup template.
- In a policy that contains a synthetic backup, incremental backups must use backup-to-disk folders or virtual tape libraries as destination devices. You cannot save the policy if one of these devices is not available.
- To display the example policy for synthetic backup, you must have the following:
 - A backup-to-disk folder
 - A virtual tape library

- A license key for the Advanced Disk-based Backup Option
See [“About creating a synthetic backup by copying the example policy”](#) on page 906.
See [“Re-creating example policies”](#) on page 528.
- The option **Collect additional information for synthetic backup and for true image restore** must be selected for backup templates for incremental and full backup jobs created for synthetic backup. This option is on the **General** page of the **Backup Job Template** properties.
See [“Creating a synthetic backup by adding templates to a policy”](#) on page 907.
This option also enables true image restore for backup sets.
See [“About true image restore”](#) on page 912.
- If the baseline backup job was written to tape, and if you also want to write the synthetic backup job to tape, two tape drives are required: one to mount the source job on (the baseline backup) and one to mount the destination job on (the synthetic backup job).

Following are limitations when running synthetic backup:

- Only file system resources are supported for synthetic backup.
See [“What you can back up with synthetic backup”](#) on page 901.
- If the Central Admin Server Option is installed, the synthetic backup job template and any associated full and incremental job templates must be run on destination devices that can all be accessed by the media server that runs the synthetic backup job.
See [“Requirements for duplicate backup data and synthetic backup jobs in CASO”](#) on page 1551.
- The option **Checkpoint Restart** is not supported when the option **Collect additional information for synthetic backup and for true image restore** is selected.
See [“Using checkpoint restart on Microsoft Cluster Server failover”](#) on page 822.
See [“About collecting additional information for synthetic backup and true image restore”](#) on page 904.
See [“About creating a synthetic backup by copying the example policy”](#) on page 906.
See [“Creating a synthetic backup by adding templates to a policy”](#) on page 907.

Best practices for synthetic backup

The following best practices are for using the synthetic backup feature:

- Use synthetic backups to back up file system resources. Do not include database backups in synthetic backups. Synthetic backup jobs are not created if the backup selection list contains any resources that are not supported.

Note: Synthetic backup is not supported for a remote resource that is in a different time zone than the media server.

- Do not select the option **Use the Microsoft Change Journal if available** if a volume contains hard links, or if you enable Single Instance Storage. Backup Exec detects that these files have been modified and performs backups without using the Change Journal. In this situation, backups for which the Change Journal option is enabled can require more time.
- Copy the example policy for a synthetic backup that Backup Exec provides, and then customize it. The example policy contains the default settings for synthetic backup.
See [“About creating a synthetic backup by copying the example policy”](#) on page 906.
- Use the template rules to ensure that the baseline backup job and the recurring incremental jobs do not run at the same time. You can select the following template rule to ensure that the baseline backup and the recurring incremental backups do not run at the same time:
If start times conflict, <Template A> will start and upon completion, starts <Template B>.
See [“Setting template rules”](#) on page 542.
- To automatically copy the backup data to tape, add a **Duplicate Backup Set** template to the synthetic backup policy. The **Duplicate Backup Set** template provides automatic duplication of backup sets.
- If you use an encryption key in a synthetic backup policy, use the same encryption key for all of the associated templates. Do not change the encryption key after you create the policy.
- Create a full backup template for the baseline backup. This configuration is useful if the baseline backup uses a different destination device or runs on a different schedule than the recurring incremental backups.

See [“About the synthetic backup feature”](#) on page 899.

About collecting additional information for synthetic backup and true image restore

All of the backup job templates that are created for synthetic backup and true image restore must have the option **Collect additional information for synthetic backup and for true image restore** selected. This option can be selected on the **General** page on the backup job template properties when a policy is being created.

This option specifies that Backup Exec collects the information required to detect files and directories that have been moved, renamed, or newly installed since the last backup, and then includes those files and directories in the backup jobs. This option also lets Backup Exec keep track of deleted files so that they are not included in true image restore of the corresponding backup sets, and are not included in the backup sets created by synthetic backup.

If this option is not selected, Backup Exec skips these files and directories if their archive bits are unchanged. If this option is selected, Backup Exec compares path names, file names, modified times, and other attributes with those from the previous full and incremental backups. If any of these attributes are new or changed, then the file or directory is backed up.

For synthetic backup, the first backup that is associated with the synthetic backup always backs up all of the files, even if it is an incremental backup. Backup Exec starts collecting the additional information with this first backup, but does not compare it to any previous backup.

See [“About the synthetic backup feature”](#) on page 899.

See [“Methods for creating a synthetic backup”](#) on page 904.

See [“Adding a backup template to a policy”](#) on page 530.

See [“About true image restore”](#) on page 912.

Methods for creating a synthetic backup

A synthetic backup can only be created in a policy.

You can use the following methods to create a policy that contains the necessary job templates for a synthetic backup:

- Use the **Policy Wizard**.
See [“Creating a synthetic backup by using the Policy Wizard”](#) on page 905.
- Make a copy of the synthetic backup policy example and then modify the job templates to suit your specific needs.
See [“About creating a synthetic backup by copying the example policy”](#) on page 906.

- Manually create a policy, and then add the necessary job templates for the synthetic backup feature.
See [“Creating a synthetic backup by adding templates to a policy”](#) on page 907.

Before creating a synthetic backup, review the requirements for synthetic backups.

See [“Requirements for synthetic backup”](#) on page 901.

See [“About the synthetic backup feature”](#) on page 899.

See [“Best practices for synthetic backup”](#) on page 902.

Creating a synthetic backup by using the Policy Wizard

You can use the **Policy Wizard** to help you create all of the templates necessary for a synthetic backup.

You can set up the policy to use the following:

- A weekly synthetic backup with daily incremental backups.
- A monthly synthetic backup with a weekly synthetic or incremental backup and daily incremental backups.

Note: You must select a backup-to-disk folder or a virtual tape library as the destination device for the incremental backups. Otherwise, you cannot save the policy.

To create a synthetic backup by using the Policy Wizard

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Policy Tasks**, click **New policy using wizard**.
- 3 On the **Welcome** page, click **Next** to continue, and follow the instructions in the wizard to complete the policy.

All of the necessary job templates for synthetic backup will be created in the new policy.

See [“About the synthetic backup feature”](#) on page 899.

See [“What you can back up with synthetic backup”](#) on page 901.

See [“Requirements for synthetic backup”](#) on page 901.

See [“Best practices for synthetic backup”](#) on page 902.

See [“About creating a synthetic backup by copying the example policy”](#) on page 906.

See [“Creating a synthetic backup by adding templates to a policy”](#) on page 907.

About creating a synthetic backup by copying the example policy

Backup Exec provides example policies that contain standard settings for different tasks. The example policy for synthetic backup contains the standard settings and job templates that are required to run a synthetic backup. You can copy this example policy, rename it, and change the times when the templates are scheduled to run.

See [“Using an example policy”](#) on page 526.

Note: The example policy for a synthetic backup only appears if you have a backup-to-disk folder or a virtual tape library. You can create a backup-to-disk folder, and then re-create the example policies.

See [“Re-creating example policies”](#) on page 528.

The example policy for synthetic backup contains the following templates:

- **Baseline Backup** - the backup template that creates the baseline backup job. This job only needs to run once. You can run additional baseline backups whenever you prefer, although the backup window is reduced by running a synthetic backup instead of a full backup.
- **Incremental Backup** - the backup template that creates the subsequent incremental backup jobs and runs after the baseline backup job runs.
- **Synthetic Backup** - the backup template that creates the synthetic backup job.

The example policy displays the template rules that you can use to set the order in which the templates should run. All template rules are optional, but they can help you make sure that the templates run in the proper sequence.

In the example policy, a full backup template is included to act as the baseline. This template should be the first template to run, and it only needs to run once. The template that runs first is called Template A.

The following rules were added:

- **<Template A> must complete at least once before any other templates will be allowed to start.** Baseline Backup is selected as <Template A>.
- **If start times conflict, <Template A> will start and upon completion, starts <Template B>.** Incremental Backup is selected as <Template A> and Synthetic Backup is selected as <Template B>.
- **Run <Template A> only once.** Baseline Backup is selected as <Template A>.

See [“What you can back up with synthetic backup”](#) on page 901.

See [“Requirements for synthetic backup”](#) on page 901.

See [“Setting template rules”](#) on page 542.

See [“About the synthetic backup feature”](#) on page 899.

See [“Best practices for synthetic backup”](#) on page 902.

Creating a synthetic backup by adding templates to a policy

Creating a new policy for a synthetic backup involves choosing a name and description for the policy, adding the necessary job templates for synthetic backup to the policy, and setting up relationships between the templates. After you set up all of the templates for a synthetic backup in a policy, you can combine the policy with a selection list to create jobs.

See [“About creating jobs using policies and selection lists”](#) on page 544.

To create a synthetic backup by adding templates to a policy

- 1 On the navigation bar, click **Job Setup**.
- 2 On the task pane, under **Policy Tasks**, click **New policy**.
- 3 Type a policy name and description for this synthetic backup policy, and then click **New Template**.

The **Template Selection** dialog box appears.

- 4 Select **Backup Template**, and then click **OK**.
- 5 On the **Properties** pane, under **Settings**, click **General**, and specify either of the following backup methods for the baseline.
 - Select **Full - Back up files - Using archive bit (reset archive bit)** to add an optional baseline full backup template.
 - Select **Incremental - Back up changed files since last full or incremental - Using archive bit (reset archive bit)** to add a recurring, incremental backup template.
- 6 Select the **Collect additional information for synthetic backup and for true image restore** option.
- 7 (Optional) Select the **Use the Microsoft Change Journal if available** option.

See [“About using the Windows NTFS Change Journal to determine changed files”](#) on page 280.

If you selected an incremental backup method, then under **Destination**, click **Device and Media**.

- 8 Select a backup-to-disk folder or a virtual tape library as the destination device.

9 On the **Properties** pane, under **Settings**, click **Network and Security**.

In a policy that contains a synthetic backup, if an encryption key is used, all of the associated templates must use the same encryption key. The encryption key should not be changed after the policy has been created. The encryption key that is selected in the associated templates is automatically applied to the synthetic backup template.

See [“About encryption”](#) on page 415.

If the Central Admin Server Option (CASO) is installed, an option displays to allow managed media servers to use any network interface to access remote agents.

See [“Enabling managed media servers to use any available network interface card”](#) on page 1540.

10 Under **Frequency**, click **Schedule** and set the scheduling options you want to use.

See [“Schedule properties for a template”](#) on page 532.

If you added the optional full backup template for the baseline backup in step 5, you must configure it to be the first backup template to run.

The baseline backup only needs to run once. You can run additional baseline backups whenever you prefer, although the backup window is reduced by running a synthetic backup instead of a full backup.

If you added an incremental backup template, you must configure it to be a recurring job so that the first instance can become the baseline backup.

11 Select other options as appropriate, and then click **OK**.

See [“Adding a backup template to a policy”](#) on page 530.

In a policy that contains a synthetic backup, backup templates that create incremental backup jobs must have backup-to-disk folders as destination devices.

12 Do one of the following:

- If you added the optional full backup template for the baseline backup in step 5, then continue with the next step to create a backup template for an incremental backup job.
- If you added a recurring incremental backup template, go to step 18 to add the synthetic backup template

13 On the **New Policy** dialog box, click **New Template**, and then on the **Template Selection** dialog box, select **Backup Template** again, and then click **OK**.

- 14 On the **Properties** pane, under **Settings**, click **General**, and select the backup method **Incremental - Back up changed files since last full or incremental - Using archive bit (reset archive bit)**.
- 15 Select the **Collect additional information for synthetic backup and for true image restore** option.
- 16 (Optional) Select the option **Use the Microsoft Change Journal if available**.
See [“About using the Windows NTFS Change Journal to determine changed files”](#) on page 280.
- 17 Under **Frequency**, click **Schedule** and set the scheduling options you want to use.
See [“Schedule properties for a template”](#) on page 532.
- 18 Select other options as appropriate, and then click **OK**.
See [“Adding a backup template to a policy”](#) on page 530.

In a policy that contains a synthetic backup, backup templates that create incremental backup jobs must have backup-to-disk folders as destination devices.
- 19 On the **New Policy** dialog box, click **New Template**, select **Synthetic Backup Template**, and then click **OK**.
- 20 On the **Properties** pane, under **Destination**, select **Device and Media**, and complete the appropriate options.
See [“Device and media options for backup jobs and templates”](#) on page 339.
- 21 On the **Properties** pane, under **Settings**, click **General**, and complete the appropriate options.
See [“General options for synthetic backup templates”](#) on page 911.
- 22 On the **Properties** pane, under **Settings**, click **Advanced**, and complete the appropriate options.
See [“Advanced options for synthetic backup templates”](#) on page 911.
- 23 If you want Backup Exec to notify someone when the backup job completes, on the **Properties** pane, under **Settings**, click **Notification**.
See [“Sending a notification when a job completes”](#) on page 682.
- 24 Do one or both of the following:
 - Set scheduling options
See [“Schedule properties for a template”](#) on page 532.
 - Create template rules

See [“Creating template rules to run job templates for synthetic backup”](#) on page 910.

Creating template rules to run job templates for synthetic backup

Set template rules to run the jobs for the synthetic backup in the correct sequence.

To create template rules to run job templates for synthetic backup

- 1 Review the documentation about template rules.
See [“Setting template rules”](#) on page 542.
- 2 On the **New Policy** dialog box, under **Template rules**, click **New Rule**.
- 3 On the **Template Rule Properties** dialog box, click the template rules drop down list, and then do one of the following:
 - If you created a backup template for a full backup job to run as the baseline backup, select **Run <Template A> only once**. Click the **Template A is:** drop down list and select the name of the template that you want to run as the baseline backup, and then click **OK**.
 - If you created a backup template for a recurring incremental backup job to run as the baseline backup, go to step 5.
- 4 Click **New Rule** again to add another rule.
- 5 On the **Template Rule Properties** dialog box, click the template rules drop down list, and then select **<Template A> must complete at least once before any other templates will be allowed to start**.
- 6 Click the **Template A is:** drop down list, select the template name of the baseline backup, and then click **OK**.
- 7 Click **New Rule** again to add another rule.
- 8 On the **Template Rule Properties** dialog box, click the template rules drop down list, and then select **If start times conflict, <Template A> will start and upon completion, starts <Template B>**.
- 9 Click the **Template A is:** list, and then select the template name of the incremental backup.
- 10 Click the **Template B is:** list, and then select the template name of the synthetic backup, and then click **OK**.
- 11 On the **New Policy** dialog box, click **OK**.

See [“Creating a synthetic backup by adding templates to a policy”](#) on page 907.

General options for synthetic backup templates

General options for a synthetic backup template provides information about the job template.

See [“Creating a synthetic backup by adding templates to a policy”](#) on page 907.

Table B-1 General options for synthetic backup templates

Item	Description
Template name	Displays the name for this job template.
Backup set description	Displays a description of the information you are backing up.
Preferred source device	Displays the device used as the destination device for the original backup job.

Advanced options for synthetic backup templates

Advanced options for a synthetic backup template provide information about verify operations and compression types for the job.

See [“Creating a synthetic backup by adding templates to a policy”](#) on page 907.

Table B-2 Advanced options for synthetic backup templates

Item	Description
Verify after job completes	Lets Backup Exec automatically perform a verify operation to make sure the media can be read after the backup has been completed. Verifying all backups is recommended.

Table B-2 Advanced options for synthetic backup templates (*continued*)

Item	Description
Compression type	Displays one of the following: <ul style="list-style-type: none"><li data-bbox="534 369 1198 546">■ None<p>This option copies the data to the media in its original form. If the data was backed up using software compression, then it is copied in its software compression form. Using some form of data compression can help expedite backups and preserve storage media space.</p><p>Hardware data compression should not be used in environments where devices that support hardware compression are used interchangeably with devices that do not have that functionality. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive.</p><li data-bbox="534 730 1198 847">■ Hardware [if available, otherwise none]<p>This option uses hardware data compression (if the storage device supports it). If the drive does not feature data compression, the data is backed up uncompressed.</p>

About true image restore

True image restore enables Backup Exec to restore the contents of directories to what they were at the time of any full or incremental backup. Restore selections in backup sets are made from a view of the directories as they existed at the time of the particular backup. Files that were deleted before the time of the backup are not restored. In true image restore, only the correct versions of files are restored from the appropriate full or incremental backups that contain them. Previous versions are not unnecessarily restored and then overwritten.

When you restore from backup sets that were enabled for true image restore, you do not have to manually select from the previous full backup, and then select the incremental backups one after another. The files that you need are automatically selected from the appropriate backups.

Backup Exec starts collecting the true image restore information beginning with the next full or incremental backup run by the policy after the option is enabled. The true image restore information is collected regardless of whether any files were actually changed.

For true image restore, Backup Exec also collects the information required to detect files and directories that have been moved, renamed, or newly installed from a tar or zip archive. Those files and directories are included in true image

restore incremental backups for this policy. Depending on how the files were packaged and how they were installed, some newly installed files are not backed up by normal incremental backups. With true image restore enabled, Backup Exec compares path names with path names from the previous full or incremental backup. If a name is new or changed, the file or directory is backed up.

The following are examples where using true image restore backs up files that would not otherwise be backed up:

- A file named C:\pub\doc is moved to or installed in C:\spec\doc. Here, the archive bit of files and subdirectories inside that directory is unchanged but C:\pub\doc is new in the C:\spec\ directory and is backed up.
- A directory named C:\security\dev\ is renamed as C:\security\devices\. Here, the archive bit of files and subdirectories inside that directory is unchanged but C:\security\devices\ is a new directory and is backed up.

The following table lists the files backed up in the C:\user\doc directory during a series of backups between December 1, 2009 and December 4, 2009:

Table B-3 Example table of files backed up because true image restore is enabled

Day	Type of Backup	Files Backed Up in C:\user\doc	Files Backed Up in C:\user\doc	Files Backed Up in C:\user\doc	Files Backed Up in C:\user\doc	Files Backed Up in C:\user\doc	Files Backed Up in C:\user\doc
December 1, 2009	Full	file1	file2	dirA\fileA	dirB\fileB	file3	
December 2, 2009	Incremental	file1	file2	dirA\fileA	-----	-----	
December 3, 2009	Incremental	file1	file2	dirA\fileA	-----	-----	
December 4, 2009	Incremental	file1	file2	-----	-----	-----	file4

Note: Dashes (-----) indicate that the file was deleted prior to this backup.

Assume that you are going to restore the December 4, 2009 version of the C:\user\doc directory.

If you perform a regular restore of the full backup set followed by a regular restore of subsequent incremental backup sets, the restored directory contains all files and directories that ever existed in C:\user\doc from December 1, 2009 (last full backup) through December 4, 2009.

For example, the following files and directories are included:

- file1
- file2
- dirA\fileA
- dirB\fileB
- file3
- file4

If you perform a true image restore of the December 4, 2009 backup, the restored directory has only the files and directories that existed at the time of the incremental backup on December 4, 2009.

The following list includes the files and directories that existed.

- file1
- file2
- file4

Backup Exec does not restore any of the files that were deleted prior to the December 4, 2009 incremental backup.

The restored directory does not include the dirA subdirectories, even though they were backed up on December 4, 2009. Backup Exec does not restore these directories because they did not exist at the time of the incremental backup, which was the reference for the true image restore.

A true image restore preserves files that are currently in the directory but were not present when the backup was completed. Assume you created a file named file5 after the incremental backup occurred on December 4, 2009, but before doing the restore.

In this case, the directory contains the following files after the restore:

- file1
- file2
- file4
- file5

See [“About collecting additional information for synthetic backup and true image restore”](#) on page 904.

See [“Best practices for true image restore”](#) on page 915.

See [“About true image catalogs”](#) on page 916.

See [“About restoring a backup set enabled for true image restore”](#) on page 917.

See [“Troubleshooting tips for true image restore”](#) on page 918.

Requirements for true image restore

Following are requirements for true image restore:

- Backup Exec must be installed on the media server.
- The Backup Exec Remote Agent for Windows Systems or the Remote Agent for Linux or UNIX Servers must be installed on any remote computers that you want to back up.
- The Advanced Disk-based Backup Option (ADBO) must be installed on the media server.
- The backup sets must be created by a policy that contains full and incremental job templates in which the option **Collect additional information for synthetic backup and for true image restore** is enabled.

You can use true image restore to back up the following resources only:

- File system data.
- Windows System State.

See [“About collecting additional information for synthetic backup and true image restore”](#) on page 904.

See [“About true image catalogs”](#) on page 916.

See [“About restoring a backup set enabled for true image restore”](#) on page 917.

See [“Troubleshooting tips for true image restore”](#) on page 918.

Best practices for true image restore

The following best practices are for true image restores:

- Do not select the option **Use the Microsoft Change Journal if available** when you create a backup template and select the option **Collect additional information for synthetic backup and true image restore** if a volume has the following:
 - Many hard links.
 - Single Instance Storage enabled.
 - Junction Points that were created with Linkd.exe.
- Avoid creating numerous incremental backups between full backups.
- Run weekly synthetic full or regular full backups.

See [“About true image restore”](#) on page 912.

See [“Troubleshooting tips for true image restore”](#) on page 918.

Enabling backups for true image restore

You can enable backups for true image restore.

Note: When the Central Admin Server Option (CASO) is installed, Backup Exec ensures that full and incremental backups that have the **Collect additional information for synthetic backup and for true image restore** option enabled are sent to a device that is accessible by the same media server. If it is not possible to send the full and incremental backups to devices that are accessible by the same media server, the policy cannot be created. You are prompted to change the policy and resubmit it. Additionally, a true image restore operation is delegated to the media server that has access to the device that contains the selected backup set.

Note: Symantec recommends that you set up a policy with a minimum of a weekly full backup and a daily incremental backup. You can add other templates as needed.

See [“Best practices for synthetic backup”](#) on page 902.

Enabling backups for true image restore

- 1 Create a policy that includes templates for a weekly or monthly full backup, and a daily incremental backup.
See [“Creating a new policy”](#) on page 522.
- 2 In each template's backup properties, under **Settings**, click **General**.
- 3 Select **Collect additional information for synthetic backup and for true image restore**.

About true image catalogs

Catalogs contain information about objects that were backed up during a backup job and that are contained in the backup set created from that job. True image catalogs for incremental backups contain additional information about all selected files and directories that were on the volume at the time of the backup job and about the latest backed-up versions of those objects. This additional information creates a true image of the entire volume as of the time of the backup job, even though the incremental backup job backed up only the changed files. In addition,

true image catalogs track deleted files, so files that were deleted before the time of the incremental backup are not restored.

Caution: If you delete a true image catalog, you can no longer perform true image restore for the backup sets contained in that catalog.

See [“About collecting additional information for synthetic backup and true image restore”](#) on page 904.

See [“About true image restore”](#) on page 912.

See [“Troubleshooting tips for true image restore”](#) on page 918.

About restoring a backup set enabled for true image restore

If backup sets are enabled for true image restore, you can choose restore selections from a view of the volume as it existed at the time of the selected backup.

You can also choose true image restore selections from a duplicate backup set. Then, if the backup set on the disk becomes unavailable, you can make true image restore selections from the duplicate backup set on the tape. Backup Exec automatically selects the most suitable duplicate backup sets that are available.

Backup Exec uses the following order of preference to choose the most suitable duplicate backup sets:

- A backup set in a backup-to-disk folder.
- A backup set that is on a tape already in a drive, or in a slot of a robotic library.
- Any other known duplicate copy.

Note: Only backup sets created by Backup Exec version 12.5 or later are supported for true image restore. When a previous version of the Remote Agent is used, you can still select backup sets to be restored from a true image view, but the restore operation performs traditional restores of backup sets. The restore operation starts from the prior full backup and processes incremental backups forward to the selected backup set. Deleted and renamed files are also restored. A message in the job log informs you that true image restore was not performed, but that all appropriate backup sets were automatically selected.

A unique icon represents the true image backup sets. Each true image backup set shows the entire selection list as it appeared on disk when the backup ran. Different

icons distinguish between the objects backed up in the viewed backup job and objects backed up in previous backups.

You can find a list of true image restore icons that appear at the following URL:

<http://entsupport.symantec.com/umi/V-269-12>

See “[Selecting backup sets that are enabled for true image restore](#)” on page 918.

See “[About true image restore](#)” on page 912.

See “[Requirements for true image restore](#)” on page 915.

See “[About true image catalogs](#)” on page 916.

See “[About duplicate backup set templates](#)” on page 548.

Selecting backup sets that are enabled for true image restore

You can view or select backup sets that are enabled for true image restore.

Selecting backup sets that are enabled for true image restore

- 1 On the navigation bar, click **Restore**.
- 2 Click **View by Resource**.
- 3 View or select the backup sets that are labeled as **True Image**.

If you select a true image backup set on the **View by Media** tab, the set is restored without true image functionality even though it is a true image backup set.

Troubleshooting tips for true image restore

Errors that can occur in regular restore jobs can also occur in true image restore jobs. If you cannot restore using true image backup sets, try restoring from individual backup sets in the **View by Media** tab.

See “[Selecting backup sets that are enabled for true image restore](#)” on page 918.

The following table lists troubleshooting tips:

Table B-4 Troubleshooting tips for true image restore operations

Problem	Explanation
The job fails with the error "Error retrieving catalog information."	Most likely, one or more prior catalogs are not present. If the catalog of the selected set is present but some prior catalogs are not present, check the job log to find this specific catalog error.
You cannot expand the View by Resource restore selections to display the true image restore selections.	If any prior catalogs are missing, the restore view cannot be expanded.

See ["About true image restore"](#) on page 912.

See ["Requirements for true image restore"](#) on page 915.

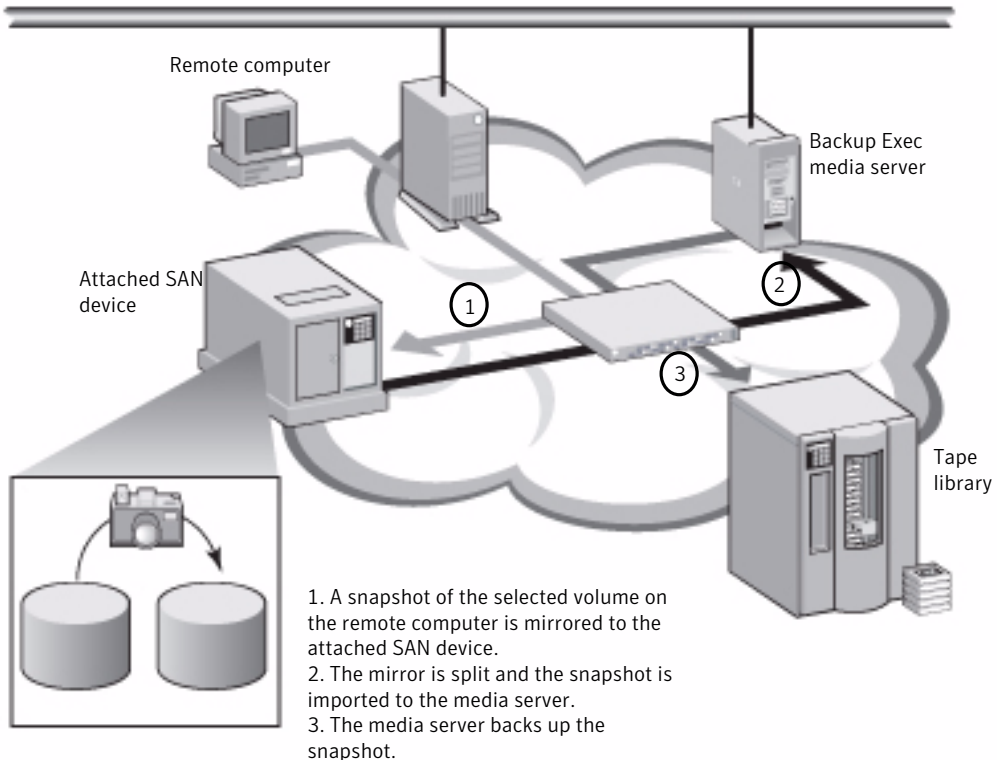
See ["About true image catalogs"](#) on page 916.

About offhost backup

Offhost backup enables Backup Exec to move backup processing from the host computer, which is the remote computer that contains the volumes selected for backup, to the Backup Exec media server. The offhost backup creates a snapshot of the volume or volumes that are selected for backup on the remote computer. The snapshots are then imported to the media server, where they are backed up.

The following illustrates the basic method for performing an offhost backup.

Figure B-4 Offhost Backup



After the backup, the snapshots are deported from the media server and mounted back on the remote computer and resynchronized with the source volume. This process requires solutions from hardware or software providers that can support transportable snapshots, that is, snapshots that can be imported to and deported from the media server. The Microsoft Volume Shadow Copy Services (VSS) provider that you select is used for each volume in the offhost backup. An offhost backup job is performed on one remote computer at a time.

Offhost backup supports the following:

- Microsoft Volume Shadow Copy Service (VSS).
- Veritas Storage Foundation for Windows (VSWF).

- Backups for NTFS volumes that use the full, incremental, and differential backup methods.
- SQL Agent backups for Microsoft SQL Server 2000 databases.
- Exchange Agent backups for Microsoft Exchange Server 2003 (Service Pack 1)/ 2007 instances that run on Windows Server 2003. Support for the option to use Backup Exec Granular Recovery Technology for Exchange Agent backups is included.

Advanced Disk-based Option offhost backup does not support the following:

- The option **Checkpoint Restart**.
- Volumes that run Windows BitLocker Drive Encryption.
- The option **Use the Microsoft Change Journal if available** for differential and incremental backups, unless you select the modified time method.
- Exchange Agent backup jobs that are configured to use Symantec Continuous Protection Server (CPS).

See [“Requirements for offhost backup when using the Veritas Storage Foundation for Windows Provider”](#) on page 922.

See [“Setting offhost backup options for a backup job”](#) on page 925.

See [“Troubleshooting the offhost backup”](#) on page 929.

See [“Browsing remote computers for installed snapshot providers”](#) on page 925.

See [“About restoring offhost backup data”](#) on page 929.

See [“Setting default backup and restore options for Exchange data”](#) on page 1123.

See [“How Granular Recovery Technology works with the Exchange Information Store”](#) on page 1106.

See [“How to use ADBO with the SQL Agent”](#) on page 1264.

Requirements for offhost backup

The following are requirements for offhost backup:

Table B-5 Offhost backup requirements

Item	Description
Media server	The following must be installed on the media server: <ul style="list-style-type: none"> ■ Backup Exec ■ Advanced Disk-based Backup Option

Table B-5 Offhost backup requirements (*continued*)

Item	Description
Remote computer	The Backup Exec Remote Agent for Windows Systems must be installed on the remote computer.
Media server and the remote computer	<p>The following must be installed on both the media server and on the remote computer:</p> <ul style="list-style-type: none"> ■ The same operating system, either Microsoft Windows Server 2003 with Service Pack 2 or Windows Server 2008. ■ The most recent Volume Shadow Copy Services (VSS) patches. ■ The Microsoft VSS hardware or software snapshot provider that you want to use. Otherwise, the snapshots of the volumes cannot be deported to the media server. ■ Ability to access the disks that are shared between the media server and the remote computer.
GRT-enabled offhost backup of Exchange Server resources	<p>The following requirements must be met:</p> <ul style="list-style-type: none"> ■ Microsoft Exchange Server 2003 (Service Pack 1) or Exchange Server 2007 instances that run on Windows Server 2003 must be installed on the Exchange Server. See “Requirements for using the Exchange Agent” on page 1095. ■ Device requirements for GRT-enabled jobs See “Recommended devices for backups that use Granular Recovery Technology” on page 324. See “About requirements for jobs that use Granular Recovery Technology” on page 325.

See [“Using checkpoint restart on Microsoft Cluster Server failover”](#) on page 822.

See [“About offhost backup”](#) on page 919.

See [“Best practices for offhost backup”](#) on page 923.

See [“Troubleshooting the offhost backup”](#) on page 929.

See [“Browsing remote computers for installed snapshot providers”](#) on page 925.

Requirements for offhost backup when using the Veritas Storage Foundation for Windows Provider

If you are using the Veritas Storage Foundation for Windows (VSW) FlashSnap option provider, read the following before running an offhost backup:

- VSW version 4.2 or later must be installed on the media server and on the computer that contains the volumes that you want to back up.

- The VSFW FlashSnap option must be installed on the computer that contains the volumes that you want to back up.
- Use the VSFW FlashSnap Snap Start command to mirror the volumes on the remote computer. The offhost backup option does not create mirrored volumes or resynchronize volumes that are already created and split.
See [“Using Snap Start on a Veritas Storage Foundation volume”](#) on page 944.
- Confirm that the mirrored volumes that you create with the VSFW FlashSnap option reside on disks that are shared between the remote computer (the computer containing the volumes to be backed up) and the media server.
- All volumes selected for offhost backup using the VSFW FlashSnap provider must belong to the same disk group. A maximum of seven volumes can be snapped at one time.
- Do not select dynamic volumes and basic volumes for the same offhost backup job because the VSFW FlashSnap option cannot perform snapshots of basic volumes. Symantec recommends that you use other backup methods for backing up basic volumes if the VSFW FlashSnap provider is selected.
- If the computer on which you want to perform an offhost backup is in an environment with the Central Admin Server Option and the Veritas Cluster Server installed, and if failover occurs to a Veritas Cluster Server node, you may need to manually clean up the snapshots before restarting the offhost backup on the failover node. Refer to the VSFW documentation for details.
See [“Troubleshooting the offhost backup”](#) on page 929.
See [“Browsing remote computers for installed snapshot providers”](#) on page 925.

Best practices for offhost backup

The following best practices are recommended:

- Keep source volumes and snapped volumes from sharing the same physical disks. If this is not maintained, then any attempt to split the snapshot volume from the original volume fails.
- Most hardware and software providers have some limitation about the types of volumes that are transportable. Therefore, Symantec recommends that you use offhost backup jobs only for backing up data for which all dependent volumes, or mounted volumes, can be imported and deported.
- Using offhost backup to back up Veritas Storage Foundation for Windows (VSFW) volumes requires that snapshot volumes in shared storage be transferred from host to host. Make sure that VSFW volumes that are backed up with offhost backup reside in VSFW disk groups that have either the "private protection" or "cluster disk group" disk group property. Private dynamic disk

group protection and cluster disk group property settings use hardware locking techniques to protect a dynamic disk group that is located on shared storage from being accessed by other hosts that are connected to the shared storage pool.

- The offhost backup will fail if any one volume that you select for backup is only supported by a Microsoft Volume Shadow Copy Services (VSS) provider and cannot be imported or deported, or if the required VSS hardware provider is not on a Symantec-approved compatibility list. You can choose to continue the backup if the offhost backup fails.

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

- The Hitachi Raid Manager log cannot be on a volume being snapped. Hitachi executes I/O to its Raid Manager log file during the snapshot commit process, and the VSS coordinator blocks I/O to any drive being snapped. Therefore, if the log directory for Raid Manager is on the volume that is being snapped, then log I/O is blocked and the snap process is deadlocked.
- If the Central Admin Server Option (CASO) is installed, for jobs that use offhost backup, you must manually select the destination device that will run the job rather than allowing the job to be delegated by the central administration server. Otherwise, the job could be delegated to a media server that does not have offhost capability.
See [“How to use media server pools in CASO ”](#) on page 1545.
- When performing an offhost backup using a VSS hardware provider in a Microsoft Cluster Server (MSCS) or Veritas Cluster Server environment, the media server and the remote computer must not be in the same cluster group. The cluster applications cannot support devices’ logical unit numbers (LUNs) that have duplicate signatures and partition layouts, therefore, the snapshots containing the LUNs must be transported to a host, or remote computer, that is outside the cluster.

See [“About Backup Exec and clusters”](#) on page 814.

See [“About offhost backup”](#) on page 919.

See [“Requirements for offhost backup”](#) on page 921.

See [“Requirements for offhost backup when using the Veritas Storage Foundation for Windows Provider”](#) on page 922.

See [“Setting offhost backup options for a backup job”](#) on page 925.

See [“Browsing remote computers for installed snapshot providers”](#) on page 925.

See [“Troubleshooting the offhost backup”](#) on page 929.

Browsing remote computers for installed snapshot providers

You can view the snapshot providers that are installed on a remote computer before you run an offhost backup for selected resources.

The Microsoft Volume Shadow Copy Services (VSS) hardware or software provider that you select when creating an offhost backup must also be installed on the remote computer that you want to back up. If the snapshot provider is not installed on the remote computer, the snapshots of the volumes cannot be imported to the media server.

To browse remote computers for installed snapshot providers

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Properties** pane, under **Source**, click **Selections**.
- 4 In the backup selections pane, do the following in the order listed:
 - Right-click the remote computer that contains the volumes you want to back up.
 - Click **List Snapshot Providers**.
- 5 View the list of the available snapshot providers on the remote computer.

See [“Requirements for offhost backup when using the Veritas Storage Foundation for Windows Provider”](#) on page 922.

See [“Requirements for offhost backup”](#) on page 921.

See [“Best practices for offhost backup”](#) on page 923.

See [“Setting offhost backup options for a backup job”](#) on page 925.

Setting offhost backup options for a backup job

You can set offhost backup options for each backup job.

For Exchange Server resources, you can create a GRT-enabled offhost backup job.

See [“Configuring a GRT-enabled offhost backup for Exchange resources”](#) on page 928.

Note: If the Central Admin Server Option (CASO) is installed, do not let the central administration server delegate the job. It can delegate the job to a media server that does not have offhost capability. You must manually select the destination device for the CASO jobs that use the offhost backup method.

See [“How to use media server pools in CASO ”](#) on page 1545.

To set offhost backup options for a backup job

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Properties** pane, under **Settings**, click **Advanced Disk-based Backup**.
- 4 Select the appropriate options, and then click **OK**.

See [“Backup options for the Advanced Disk-based Backup Option”](#) on page 926.

Backup options for the Advanced Disk-based Backup Option

Backup options for the Advanced Disk-based Backup Option provide information on the settings for offhost backup jobs.

See [“Setting offhost backup options for a backup job”](#) on page 925.

See [“Setting default options for offhost backup jobs ”](#) on page 928.

Table B-6 Backup options for the Advanced Disk-based Backup Option

Item	Description
Use offhost backup to move backup processing from remote computer to media server	Indicates if offhost backup is enabled. If this is enabled for a single job, or as a default for all backup jobs, then an offhost backup of all volumes will be performed if all requirements are met. See “About offhost backup” on page 919.

Table B-6 Backup options for the Advanced Disk-based Backup Option
(continued)

Item	Description
Snapshot provider	<p>Displays the following options:</p> <ul style="list-style-type: none"> ■ Automatic - Use hardware if available; otherwise use software Select this option to enable VSS to select the best provider for the selected volume. ■ Software - Use Veritas Storage Foundation for Windows ■ Hardware - Use technology provided by hardware manufacturer <p>If Software or Hardware are the snapshot providers, then the following information applies:</p> <ul style="list-style-type: none"> ■ The provider must support transportable snaps. ■ If multiple volumes are selected, then all volumes must be snappable by the same type of provider. ■ Software and hardware providers cannot both be used to snap different volumes in the same job. You must either create another job, or make sure that the option Process logical volumes for offhost backup one at a time is selected.
Continue the backup job (offhost backup is not used)	<p>Lets the backup job complete even if any of the volumes selected do not support offhost backup, or if an error occurs that is related to the snapshot or volume import. The backup will run according to all the other options that have been set for this job.</p>
Fail the backup job (further selections are not backed up after failure occurs)	<p>Terminates the offhost backup job if any of the selected volumes do not support offhost backup, or if an error occurs that is related to the snapshot or volume import.</p>

Table B-6 Backup options for the Advanced Disk-based Backup Option
(continued)

Item	Description
Process logical volumes for offhost backup one at a time	<p>Enables the backup of multiple volumes in one job, while creating a snapshot of only one logical volume at a time. To ensure database integrity, or if a volume contains mount points, multiple volumes may need to be snapped at one time.</p> <p>After the logical volume is snapped and backed up, the snapshot is deleted before the next logical volume is snapped. This option increases the ability to meet the minimum quiet time needed to complete a snapshot.</p> <p>A logical volume can comprise multiple physical volumes. A single logical volume can encompass all of the volumes on which databases reside.</p>

Setting default options for offhost backup jobs

You can set the defaults that are used for every backup job.

Note: If the Central Admin Server Option (CASO) is installed, do not let the central administration server delegate the job. It can delegate the job to a media server that does not have offhost capability. You must manually select the destination device for the CASO jobs that use the offhost backup method.

See [“How to use media server pools in CASO ”](#) on page 1545.

To set default options for offhost backup jobs

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Advanced Disk-based Backup**.
- 3 Select the appropriate options, and then click **OK**.

See [“Backup options for the Advanced Disk-based Backup Option”](#) on page 926.

Configuring a GRT-enabled offhost backup for Exchange resources

You can enable the Backup Exec Granular Recovery Technology (GRT) option for offhost backups of Exchange resources. When you select the GRT option for a

backup, Backup Exec collects additional information for the catalog. This information lets you restore individual mailboxes, mail messages, and public folders from Information Store backups.

Offhost backup does not support the Exchange Server backup jobs that are configured to use Symantec Continuous Protection Server (CPS).

You should perform consistency checks before you run an offhost backup.

To configure a GRT-enabled offhost backup for Exchange resources

- 1 Create an Exchange backup job.
See [“About backing up Exchange 2003/2007”](#) on page 1129.
- 2 Check **Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups.**
- 3 If you send the job to a backup-to-disk folder with file size limitations, ensure that Backup Exec can stage temporary metadata on the default path of C:\temp.
See [“Setting default backup options”](#) on page 391.
- 4 Set the offhost backup options.
See [“Setting offhost backup options for a backup job”](#) on page 925.
- 5 Select **Automatic - Use hardware if available; otherwise use software.**
The options for job disposition are not available.
- 6 If you include resources that are not supported for offhost backup, check the following check box: **Process logical volumes for offhost backup one at a time** to let the job complete with errors.

About restoring offhost backup data

Use a standard restore job to restore data that is backed up with the offhost backup method. Data is restored directly from the backup media to the original volumes on the remote computer.

See [“Restoring data by setting job properties”](#) on page 605.

Troubleshooting the offhost backup

Offhost backup requires that the VSS providers and the volumes that are to be transported are set up correctly. Not all arrays are supported with the Advanced Disk-based Option.

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

To troubleshoot offhost backup issues, Symantec recommends that you use the tools that are available from the VSS provider to verify the setup that is required for offhost backup.

The minimum set up requirements are as follows:

- Volumes that you want to back up are snappable.
- Volumes are shared between the remote computer and the media server.
For example, when using Veritas Storage Foundation for Windows (VSW) as the provider, you can use the Veritas Enterprise Administrator (VEA) to verify snapshots of the volumes, split the snapped volumes into a different Disk Group (DG), deport the DG from the remote computer and import it to the media server. All providers will have similar administration console or command-line tools that will allow you to take a snapshot of volumes, and deport and import volumes.
- An offhost backup job can only contain volumes that can be transported to the media server for backup.
See [“Requirements for offhost backup”](#) on page 921.
See [“Requirements for offhost backup when using the Veritas Storage Foundation for Windows Provider”](#) on page 922.

Other factors to consider are as follows:

- Microsoft Windows Server 2003 with Service Pack 1 or Windows Server 2008 must be installed on both the media server and the remote computer. Both computers must have the most recent Volume Shadow Copy Services (VSS) patches.
- Microsoft XML Core Services (MSXML4) must be installed and running on both the media server and the remote computer.

Troubleshooting offhost backup issues depends to an extent on the VSS provider used for the snapshots, but the following setup issues that are common to all providers may cause offhost backup to fail:

Table B-7 Common setup issues for offhost backup

Issue	Solution
The volumes are not shared.	For offhost backup to work, all volumes must reside on disks that are shared between the remote computer and the Backup Exec media server. It is the backup administrator's responsibility to confirm this. If the volumes are not shared, the import operation will fail, and you may need to clean up the snapshots and resynchronize the volumes manually.
The VSS provider is not installed on the media server and the remote computer.	The provider used for snapshot must be installed on both the media server as well as on the remote computer. If the provider is not installed on the media server, the import operation will fail, and you may need to clean up the snapshots and resynchronize the volumes manually.
All volumes are not transportable.	All volumes selected for backup must be transportable to the media server. If Microsoft SQL or Exchange, or other database applications are selected for backup, make sure that the databases and log files reside on transportable volumes.
The VSS provider cannot snap all of the selected volumes.	In addition to being transportable, all volumes selected for backup must be snappable by the same provider. It is the backup administrator's responsibility to ensure that all volumes in a backup job are supported by the same VSS provider.
The log path location is incorrect.	Log files created by the provider or by its supporting application during normal snapshot operation should not reside on any of the volumes being snapped. This prevents VSS from flushing the write buffers, and the snapshot will time-out. Change the log path to another volume.
The provider or VSS services are not started	Make sure that the provider service is running and make sure that the Microsoft Windows "Volume Shadow Copy" service has not been disabled.

Table B-7 Common setup issues for offhost backup (*continued*)

Issue	Solution
The credentials are incorrect.	Make sure that the machine-level credentials used for the job are the same on both the media server and the remote computer. Incorrect credentials will cause snapshots or the backup to fail.
The VSS provider is not installed on all media servers in a Central Admin Server Option (CASO) environment.	If a backup job is configured in a CASO environment, you must target the job to media servers on which the selected VSS provider is installed rather than allowing the job to be delegated by the central administration server. Otherwise, the job could be delegated to a media server that does not have offhost capability. See "How to use media server pools in CASO" on page 1545.

Table B-7 Common setup issues for offhost backup (*continued*)

Issue	Solution
The media server and the remote computer are in the same cluster group.	<p>When performing an offhost backup in a Microsoft Cluster Server (MSCS) or Veritas Cluster Server environment, the media server and the remote computer must not be in the same cluster group. The cluster applications cannot support devices' logical unit numbers (LUNs) that have duplicate signatures and partition layouts; therefore, the snapshots containing the LUNs must be transported to a media server that is outside the cluster in which the host cluster resides.</p> <p>See “How Backup Exec works in a Microsoft Cluster Server environment” on page 816.</p> <p>If you are using an Hitachi 9970 and attempt to protect a Microsoft Cluster Server (MSCS) resource using the Advanced Disk-based Backup Option (ADBO), you may receive the following error message:</p> <pre>The job failed with the following error: querying the Writer status.</pre> <p>To correct this problem, ensure that the RM Shadow Copy Provider for Volume Snapshot Service is present and running. If the service is not running, run RMVSSPRV.exe from c:\horcm\tool. If the service is still not running, contact Hitachi for support.</p>

Offhost backup failures when using VSFW as a provider

Following are the most common causes of snapshot failure and offhost backup failure when using the Veritas Storage Foundation for Windows (VSFW) software provider:

Table B-8 Common causes of snapshot failure and offhost backup failure

Issue	Description
The volume has not had a snap start.	<p>Backup Exec requires that you first snap start all volumes using the Veritas Enterprise Administrator administration console or command-line interface before attempting an offhost job.</p> <p>See “Using Snap Start on a Veritas Storage Foundation volume” on page 944.</p> <p>If you already performed the snap start for a previous snapshot operation, you must either snap start the volume again (to another physical disk) or snap back the previous snapshot volume. For details on how to snap back, refer to the documentation for the Veritas Storage Foundation for Windows software provider.</p>
The volumes that are selected for backup are basic volumes.	VSWF allows only dynamic volumes to be snapped.
The volume that is selected for backup resides on a disk with other volumes.	The disk group cannot be deported.
The snapshot volumes and the source volumes share the same physical disks.	Any attempt to split the snapshot volumes from the source volumes will fail if the snapshot volumes and the source volumes share the same physical disks. The administrator must keep the source volumes and the snap volumes from sharing the same physical disks.
The VSWF version is not supported.	Offhost backup requires that VSWF version 4.1 or later be installed on both the remote computer and on the media server. Previous releases are not supported. Compatible VSWF software versions must reside on both the media server and the remote computer. It is recommended that the same version of VSWF be installed on both computers.
Multiple disk groups are selected for the same off-host backup job.	All dynamic volumes designated for backup must be in the same Disk Group (DG). Multiple disk groups in the same offhost job are not supported.

Table B-8 Common causes of snapshot failure and offhost backup failure
(continued)

Issue	Description
More than seven volumes are selected for a snapshot operation.	Ensure that there are no more than seven volumes in a single snapshot operation. When more volumes are installed, they cannot all be snapped within the ten-second time-out that VSS imposes for snapshots to complete, and the snapshot will fail.
Basic volumes and dynamic volumes are selected for the off-host backup job.	Basic volumes and dynamic volumes cannot be mixed in a backup job that uses the offhost feature. A workaround is to make sure that the option Process logical volumes for offhost backup one at a time is selected. This restriction also applies to dynamic volumes that are mounted by mount points on basic volumes.
Dynamic boot volumes and system volumes are selected for the off-host backup job.	Dynamic boot and system volumes are restricted from VSFW FlashSnap functionality for compatibility reasons. Therefore, dynamic boot and system volumes are not supported for an offhost backup.

Note: Most VSS providers have some limitation about the types of volumes that are transportable. Therefore, it is recommended that you do not use offhost backup jobs for complete system protection. A best practice is to use offhost backup jobs to back up databases and logs when all of the dependent volumes on which data resides are transportable. Any volumes that are used to host mount points for data volumes must also be transportable because the offhost backup must snap both the data volumes and the volume with the mount point for backup.

Offhost backup issues when using a hardware provider

Hardware disk array vendors may support VSS snapshots and the transporting of volumes to the media server for backup in a SAN environment. Using hardware providers requires a sound understanding of how disk arrays are configured for shared access between the remote computer and the media server in a SAN.

Consult the documentation for your hardware disk array on how to set up such disk arrays for offhost backup. Specifically, note any limitations on using the disk

arrays in context with VSS snapshots, and note how to verify if the volumes are transportable. It is highly recommended that you make use of any tools provided by the vendor to help verify the setup and for troubleshooting issues.

Offhost backup issues that can occur when using Hitachi hardware include the following:

Table B-9 Offhost backup issues when using Hitachi hardware

Issue	Description
Hitachi supports only basic disks for offhost backup.	If a computer uses a combination of dynamic and basic disks, a complete system backup using the offhost backup feature is not possible when the Hitachi provider is used.
Veritas Cluster Server (VCS) is not supported with the Hitachi provider.	The Hitachi provider does not support dynamic disks for offhost backup.

See [“Requirements for offhost backup”](#) on page 921.

See [“Requirements for offhost backup when using the Veritas Storage Foundation for Windows Provider”](#) on page 922.

See [“Best practices for synthetic backup”](#) on page 902.

Symantec Backup Exec Advanced Open File Option

This appendix includes the following topics:

- [About the Advanced Open File Option](#)
- [How to install the Advanced Open File Option](#)
- [Setting default options for the Advanced Open File Option](#)
- [Configuring the Advanced Open File Option for backup jobs](#)
- [About the job log and the Advanced Open File Option](#)

About the Advanced Open File Option

The Symantec Backup Exec Advanced Open File Option (AOFO) uses advanced open file and image technologies designed to alleviate issues that are sometimes encountered during backup operations, such as protecting open files and managing shortened backup windows.

When a job is submitted for backup with the AOFO selected, a snapshot of each volume is created, providing a point-in-time record of the data. When creating a snapshot, Backup Exec uses snapshot technologies to momentarily suspend write activity to a volume so that a snapshot of the volume can be created.

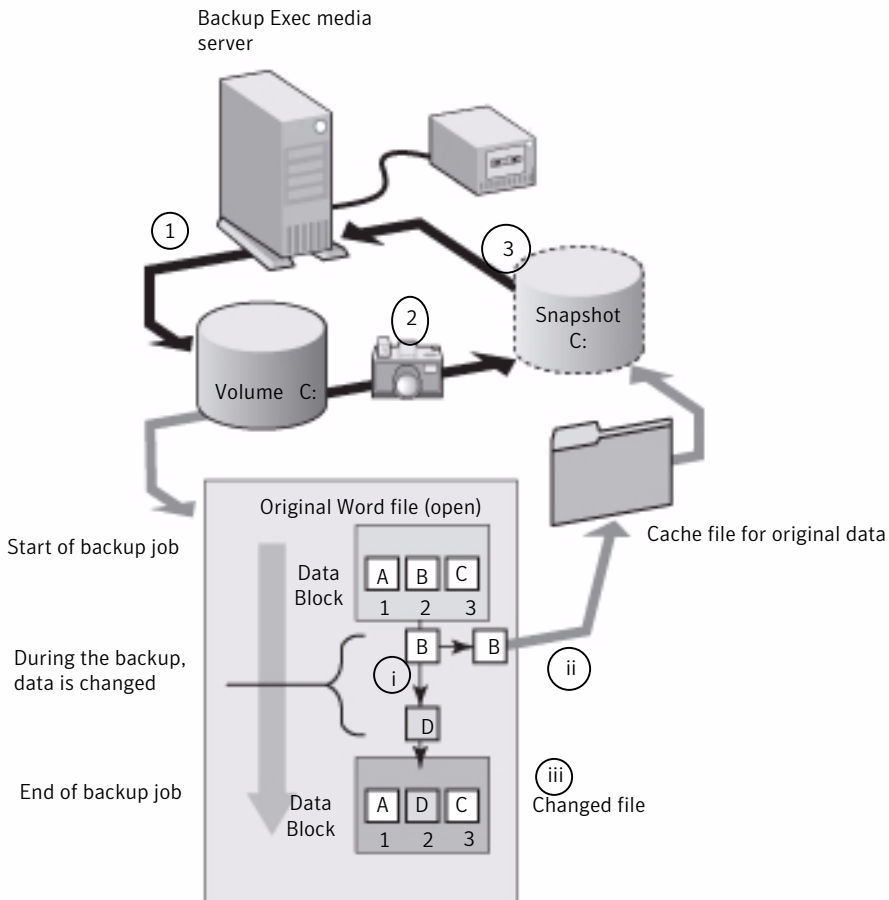
If the files selected for backup reside on more than one volume, by default Backup Exec creates a snapshot for each volume that contains data to be backed up. For example, if the data to be backed up resides on a single volume, a single snapshot is created. If data resides on four volumes, four snapshots are created. After creating a snapshot, the data is backed up from the snapshots, and then the snapshots are deleted.

During the backup, files can be open and data can be changed. Depending on the snapshot provider that you are using, open files are handled using different methods.

See “[Best practices for using the Symantec Volume Snapshot Provider](#)” on page 945.

The following graphic shows how AOFO works.

Figure C-1 Advanced Open File Option with Symantec Snapshot Provider



The illustration represents the following:

- 1 - A backup begins for volume C on a Windows server using AOFO.
- 2 - A snapshot is taken of volume C, that provides a point-in-time record of the data.
- 3 - After the snapshot is taken, the backup job starts and the data from volume C is written to tape.

During the backup job, files can be opened, and data can change. AOFO allows data to change by making a copy of the original data, named a cache file. The snapshot tracks the data changes, as illustrated in the cache file diagram.

For example, an open file such as a Microsoft Word document contains the data A, B, and C in blocks 1, 2, and 3.

i - During the backup job, "B" changes to "D" in block 2.

ii - The original data in block 2 is copied to a cache file. In this case, "B" is the original data.

iii - The changed file is now the most current file.

When the snapshot comes to a changed block, it replaces the changed block with the original data from the cache file. The snapshot sends the point-in-time data to Backup Exec. The data is then written to tape or disk.

When the backup is complete, the snapshot is deleted.

Symantec strongly recommends that Backup Exec database agents be used to back up databases. Backup Exec database agents provide selective restores of data and more integration with the database application while preventing backups of partial transactions. Agents also enable backups on a database that is spread across multiple disk volumes.

When you select AOFO for a volume-level backup of Microsoft SQL or Exchange servers but do not use the database agents, the SQL or Exchange databases are excluded from the backup.

When you select AOFO for backup of an Oracle server, databases are backed up automatically. To avoid duplicate backups of the database files, manually exclude the database files from the backup job.

You can use AOFO on the same volume as a database to provide open file support for other applications. AOFO provides generic protection for flat files when Backup Exec agents are not used, and provides protection for Microsoft Outlook PST files.

AOFO is not available for use with Extensible Firmware Interface (EFI) system partition backups.

See [“About the Backup Exec Exchange Agent”](#) on page 1094.

See [“Setting backup options for SQL”](#) on page 1273.

See [“About the Backup Exec Oracle Agent”](#) on page 1315.

See [“Setting default options for the Advanced Open File Option”](#) on page 943.

See [“Best practices for using the Symantec Volume Snapshot Provider”](#) on page 945.

About supported snapshot technologies

To use snapshot technologies with Backup Exec, you must install the Advanced Open File Option (AOFO). After making backup selections and selecting AOFO, you can configure Backup Exec to use the snapshot technologies installed on your computers.

Backup Exec supports the following snapshot technologies:

Table C-1 Supported snapshot technologies

Operating System	Snapshot Technology
On Windows 2000/XP 32-bit	Symantec Volume Snapshot Provider (VSP) Symantec Volume Snapshot Provider (VSP) is installed when you install AOFO.
On Windows 2000/2003	Veritas Storage Foundation™ by Symantec, formerly known as Veritas Volume Manager (VM), FlashSnap Option
On Windows Server 2003 and later	Microsoft Volume Shadow Copy Service (VSS) Third party software vendors also provide additional components that work in conjunction with the Microsoft Volume Shadow Copy Service. These components, called Writers, are used to flush application data or file data (if a file is open) residing in the computer’s memory before the Microsoft Volume Shadow Copy Service makes a snapshot of the volume to be backed up. See your software documentation for information about VSS Writers that may be provided by the application software vendor. For Windows Vista/Server 2008, VSS is always used by default. Note: If you turn off Active Directory, Microsoft Volume Shadow Copy Service (VSS) is not available. Jobs that require VSS fail.

See [“Setting default options for the Advanced Open File Option”](#) on page 943.

See [“Using Snap Start on a Veritas Storage Foundation volume”](#) on page 944.

See [“Best practices for using the Symantec Volume Snapshot Provider”](#) on page 945.

See [“About the Symantec Volume Snapshot Provider cache file location ”](#) on page 946.

Requirements for using Advanced Open File Option

The computer for which you want to use the Advanced Open File Option (AOFO) requires the following:

- AOFO must be installed.
- There must be enough free disk space on at least one volume to cache the data that changes during the backup job.
- The file system must be NTFS, FAT32, or FAT. To use Microsoft Volume Shadow Copy Service (VSS), at least one NTFS partition is required.
- To protect remote and local computers, the Backup Exec Remote Agent for Windows Systems must be installed. The Remote Agent is installed by default on the media server when Backup Exec is installed. When you install AOFO on remote computers, it automatically installs the Remote Agent.

Note: AOFO cannot be used on CD-ROM, floppy diskettes, or removable media. Additionally, the Checkpoint Restart option is not supported by AOFO.

See [“Enabling or disabling checkpoint restart ”](#) on page 824.

When backing up encrypted files with AOFO on Windows 2000 computers, a drive letter for the snapshot appears in Windows Explorer and on the Backup Exec Administration Console. Do not attempt to access or back up this drive letter. If no drive letters are available, encrypted files are backed up from the original volume, and the job is logged as Complete with Exceptions.

See [“About the Advanced Open File Option”](#) on page 937.

See [“Setting default options for the Advanced Open File Option”](#) on page 943.

See [“Using Snap Start on a Veritas Storage Foundation volume”](#) on page 944.

See [“Best practices for using the Symantec Volume Snapshot Provider”](#) on page 945.

See [“About the Symantec Volume Snapshot Provider cache file location ”](#) on page 946.

How to install the Advanced Open File Option

You select the Advanced Open File Option (AOFO) on the media server during installation. You must reboot the computer on which you are installing AOFO for Windows 2000/XP 32-bit after completing the installation.

You can install AOFO in the following ways:

- Install AOFO on a local media server.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

- Install AOFO on a remote media server.
See [“Push-installing the Remote Agent and Advanced Open File Option to remote computers”](#) on page 134.
- Use a command script to install the Remote Agent and AOFO.
See [“Using a command script to install the Remote Agent and AOFO ”](#) on page 148.
- Use the Windows command line to Install and uninstall AOFO on remote servers.
See [“Installing the Advanced Open File Option to remote Windows computers using the command line”](#) on page 942.

See [“About upgrading from previous versions of Backup Exec”](#) on page 180.

Installing the Advanced Open File Option to remote Windows computers using the command line

You can install the Advanced Open File Option (AOFO) using silent mode on a remote computer using the Windows command line. Silent mode runs the installation operation without the benefit of a user interface.

The AOFO files are installed on the remote computer in the following directory:

`\Program Files\Symantec\Backup Exec\RAWS`

The AOFO installation log file is created in the following directory:

`\Documents and Settings\All Users\Application Data\Symantec\Backup Exec\Logs\rawsinst.htm`

On Windows 7/Vista/Server 2008 R2/Server 2008 the AOFO installation log file is created in the following directory:

`\ProgramData\Symantec\Backup Exec\Logs\rawsinst.htm`

See [“About selecting data to back up ”](#) on page 280.

To install AOFO to remote computers using the command line

- 1 Move to a remote server.
- 2 Do one of the following:
 - Map a drive letter to the Backup Exec media server and change directories to the Advanced Open File Option install directory. By default, it is located at the following path:
`\Program Files\Symantec\Backup Exec\Agents`

- Copy the RAW32 and MSXML folders to a local directory.
- 3 Open a command prompt and enter the drive letter you mapped in step 2 and the following path:

\RAW32

- 4 Do one of the following:

To install AOFO without advertising enabled

At the command prompt, type the following:

```
setup.exe /AOFO: -s -boot
```

To install AOFO with advertising enabled

At the command prompt, type the following:

```
setup.exe /AOFO: -s /ADVRT:  
<media server name 1> <media server  
name 2>
```

The `-s` parameter is used to run the install operation in silent mode, without the benefit of a user interface.

The parameter `-boot` is used to automatically restart your computer. If you want to do this, add the parameter `-boot`; if not, you must restart the computer manually at your convenience in order to activate the Advanced Open File Option.

- 5 After the installation finishes, restart the computer at your convenience in order to activate the Advanced Open File Option.

Setting default options for the Advanced Open File Option

You can set the Advanced Open File Option (AOFO) default options for every backup job.

After a job is completed, check the section Backup Set Detail Information in the job log to make sure that AOFO was used during the backup.

See [“Configuring the Advanced Open File Option for backup jobs”](#) on page 948.

See [“Best practices for using the Symantec Volume Snapshot Provider”](#) on page 945.

To set default options for AOFO

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Advanced Open File**.

- 3 Select the appropriate options.
See [“Advanced Open File options”](#) on page 949.
- 4 Click **OK**.

About Snap Start on a Veritas Storage Foundation volume

Veritas Storage Foundation™ for Windows FlashSnap Option was formerly known as Volume Manager (VM). Before you can use the Veritas Storage Foundation for Windows FlashSnap option to back up volumes, you must use Veritas Enterprise Administrator (VEA) to snap-start the volumes. You must purchase VEA separately.

See [“Using Snap Start on a Veritas Storage Foundation volume”](#) on page 944.

If the computer on which you are running a backup using the Advanced Open File Option is in an environment with the Central Admin Server Option and the Veritas Cluster Server installed, and if failover occurs to a Veritas Cluster Service node, you must manually clean up the snapshots before restarting the backup on the failover node. Refer to the VSFW documentation for details.

When the Veritas Storage Foundation for Windows FlashSnap Option is used for AOFO backups, the SnapBack of the volume is done asynchronously since SnapBack can take a long time (depending on the size of the volume snapped and the changes that may have occurred during the backup). Rather than hold the job completion for the length of the time this operation may take, the job will instead Complete with Success (if no other error occurs). Use the VERITAS Enterprise Administrator to verify that the re-synchronization completed.

In rare cases, it is possible that the SnapBack fails and a broken mirror results. If this occurs, the next FlashSnap job submitted for the same volume may also fail, with the error Volume cannot be snapped or Volume has not been Snap-started or is not a dynamic volume. The job is logged as Completed with Exceptions. Use the VERITAS Enterprise Administrator to verify why the SnapBack did not complete and then correct the error.

Using Snap Start on a Veritas Storage Foundation volume

You can use Snap Start to start a volume. Snap Starting a volume only needs to be done once. The Snap Start procedure can take a considerable amount of time because it creates a mirror.

If you are backing up SQL or Exchange databases on the Snap started volume, you must make your selections using the Backup Exec SQL or Exchange database agents. Do not select a database or log at the volume level.

To use Snap Start on a Veritas Storage Foundation volume

- 1 Start Veritas Enterprise Administrator.
- 2 In the left pane, expand the Localhost object.
- 3 Under the Localhost object, expand the Volumes object, and then right-click the volume to Snap Start.
- 4 On the short cut menu, select **Snap**, and then select **Snap Start**.
- 5 On the Snap Start Volume screen, select either **Auto select disks** or **Manually select disks**.

Auto select disks enables Veritas Storage Foundation to make the disk selection for you, while Manually select disks enables you to make the selection.

- 6 Click **OK** to begin the snap start of the volume.

See [“About selecting data to back up”](#) on page 280.

See [“About backup strategies for SQL”](#) on page 1258.

See [“About the Backup Exec Exchange Agent”](#) on page 1094.

Best practices for using the Symantec Volume Snapshot Provider

Following are recommended best practices for using AOFO and the Symantec Volume Snapshot Provider (VSP):

- Ideally, allow AOFO exclusive use of a disk. This disk should not have any user data, should never be backed up, and should have the AOFO cache file location directed to it.
- Make sure there is sufficient space on the disk to encompass all of the changed data. Changed data can include user files, system files, and the NTFS Master File Table (MFT).
- Exclude the cache file from real-time virus scanning software. Do not run regular scans or disk utilities such as scan disk or defragmenters during backups with AOFO.
- To prevent the VSP cache file from excessively large growth on a disk during a backup operation, avoid the following:
 - Processes that write excessive data to the disk.
 - Operations that copy large amounts of data to the disk.
- Create backup-to-disk folders on a different physical disk than the disk you want to back up. For example, if AOFO is used to snap volumes during a backup,

and if the destination device is a backup-to-disk folder, the backup-to-disk folder should be on a separate volume that is not being snapped.

See [“About the Advanced Open File Option”](#) on page 937.

See [“Requirements for using Advanced Open File Option”](#) on page 941.

See [“Setting default options for the Advanced Open File Option”](#) on page 943.

See [“How to adjust the Symantec Volume Snapshot Provider cache file size”](#) on page 947.

See [“About the job log and the Advanced Open File Option”](#) on page 951.

About the Symantec Volume Snapshot Provider cache file location

When the Symantec Volume Snapshot Provider (VSP) and the Advanced Open File Option (AOFO) are used, Backup Exec creates a cache file on the disk to save any changes that are made to the files on the volume while it is being snapped. Backup Exec automatically calculates the sizes of the cache files needed for the backup, as well as the location of the cache files. The cache file is created in a hidden folder named Backup Exec AOFO Store in the root of the selected volume. The cache file extension is .vsp.

Backup Exec locates the Symantec Volume Snapshot Provider (VSP) cache files based on the following:

Table C-2 How Backup Exec locates VSP cache files

Item	Description	Notes
If you used the AOFO wizard to specify a location for the cache files	The location that you specified is used if it is not write-protected, and if it is not part of the resources that are being snapped.	If multiple source volumes (the volumes to be snapped) are being snapped, then multiple cache files (one for each source volume) are located on the volume you specified (if that volume is not being snapped).

Table C-2 How Backup Exec locates VSP cache files (*continued*)

Item	Description	Notes
If you have not specified a location for the cache file	<p>Backup Exec will attempt to locate the cache file on volumes other than the source volumes.</p> <p>If the criteria for the non-source volumes are not met, the cache file is created on the source volume.</p>	<p>These volumes must meet the following requirements:</p> <ul style="list-style-type: none"> ■ The volumes must be fixed drives. ■ The volumes must use a recognized file system (FAT, FAT32, or NTFS). ■ The volumes must be mounted locally. ■ The volumes must be valid cache file locations. <p>A valid cache file location must meet the following requirements:</p> <ul style="list-style-type: none"> ■ Cannot be a location for a snapshot by another job that is currently running. ■ Cannot contain another active cache file. ■ Cannot be write-protected or disabled.

See [“About the Advanced Open File Option”](#) on page 937.

See [“About supported snapshot technologies ”](#) on page 940.

See [“Setting default options for the Advanced Open File Option”](#) on page 943.

See [“About the job log and the Advanced Open File Option”](#) on page 951.

How to adjust the Symantec Volume Snapshot Provider cache file size

When you use the Advanced Open File Option with the option, **Automatically select open file technology**, Backup Exec chooses either the Symantec Volume Snapshot Provider or the Microsoft Volume Shadow Copy Service as the technology to use when open files are encountered. Both of these providers use "Copy-on-Write" technology to create snapshots. The snapshot itself consists of a virtual volume and a cache file. The cache file keeps track of the changes to the volume being snapped after the snapshot is taken. Using this technology, Backup Exec can back up data at a point-in-time, while also ensuring data consistency.

By default, Backup Exec determines the cache file location based on the amount of used disk space on the volume to be snapped and the availability of free disk space on other volumes.

The size of the cache file grows based on the amount of time the snapshot is active and by the rate of data change occurring on the volume during the time the snapshot is active. By default, Backup Exec sets a predetermined maximum cache file size, which in cases of heavy disk write activity during a backup job, can be exceeded. In such cases, the backup job will fail. If this occurs, you can increase the maximum size of the Advanced Open File Option cache file using the Advanced Open File Option wizard, or by running backup jobs during periods of low computer usage. In extreme cases, you may have to allocate a maximum cache file size equal to the used space on the volume being snapped.

Although rare, running out of disk space due to an insufficient cache size also can occur when running a backup operation while simultaneously running virus scan or disk defragmentation operations.

To change the cache file size, use the Advanced Open File Option wizard.

See [“Setting default options for the Advanced Open File Option”](#) on page 943.

See [“About the Symantec Volume Snapshot Provider cache file location ”](#) on page 946.

Configuring the Advanced Open File Option for backup jobs

You can set options for the Advanced Open File Option (AOFO) for each backup job.

To set options for an AOFO backup job

- 1 On the navigation bar click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 In the **Properties** pane, under **Source**, click **Selections**.
- 4 Select the resource you want to back up.
- 5 In the **Properties** pane, under **Settings**, click **Advanced Open File**.
- 6 Select the appropriate options.
See [“Advanced Open File options ”](#) on page 949.
- 7 Start the backup jobs or select other backup options from the Properties pane.

Advanced Open File options

You can set the following options for the Advanced Open File Option (AOFO).
See [“Setting default options for the Advanced Open File Option”](#) on page 943.
See [“Configuring the Advanced Open File Option for backup jobs”](#) on page 948.

Table C-3 Advanced Open File options

Item	Description
Use Advanced Open File Option	Indicates if AOFO is enabled for backup jobs. If you do not select this option, jobs saved before you installed the Advanced Open File Option will continue to use the previous settings for backing up open files.
Automatically select open file technology	Enables Backup Exec to select the best snapshot method to use for the type of data being backed up. However, even if selected, a snapshot method may not be used if the resources do not meet the requirements for using snapshot methods. If the volume you select for backup does not meet the requirements for using AOFO, then the open file backup options apply (Never; If closed within 30 seconds; With a lock; Without a lock). See “Advanced options for backup jobs” on page 348. If you select AOFO and the Microsoft VSS provider, then Backup Exec uses the first available hardware, software, or system provider to perform the snapshot.
Symantec Volume Snapshot Provider (Windows 2000 only)	Indicates if Symantec Volume Snapshot Provider (VSP) is enabled as a snapshot provider. You can only use VSP on computers that have Windows 2000/XP 32-bit. VSP is Symantec's snapshot technology for Backup Exec. You can use the AOFO Wizard to help you configure VSP. See “Best practices for using the Symantec Volume Snapshot Provider” on page 945.
AOFO Wizard	Starts a wizard that helps you configure the Symantec Volume Snapshot Provider for use with AOFO.
Veritas Storage Foundation™ for Windows FlashSnap Option (Windows 2000 and 2003 only)	Indicates if Veritas Storage Foundation™ is enabled as a snapshot provider. You can only use Veritas Storage Foundation™ on computers that have Windows 2000/2003. See “About Snap Start on a Veritas Storage Foundation volume” on page 944.

Table C-3 Advanced Open File options (*continued*)

Item	Description
<p>Microsoft Volume Shadow Copy Service (Windows 2003 and later)</p>	<p>Enables third-party hardware and software vendors to create snapshot add-ins for use with Microsoft’s technology.</p> <p>Microsoft, as well as other third party software vendors, often provide additional components that work in conjunction with VSS. These components, called Writers, are used to flush application data or file data (if a file is open) residing in the computer’s memory before the Microsoft Volume Shadow Copy Service makes a snapshot of the volume to be backed up.</p> <p>See your software documentation for information about VSS Writers that may be provided by the application software vendor.</p> <p>If you turn off Active Directory, Microsoft Volume Shadow Copy Service (VSS) is not available. Jobs that require VSS fail.</p>
<p>Snapshot provider</p>	<p>Indicates the snapshot provider to use for jobs.</p> <ul style="list-style-type: none"> ■ Automatic - Allow VSS to select the snapshot provider. Select this option to enable VSS to select the best provider for the selected volume. The order in which a snapshot provider is selected is hardware provider, software provider, and then the system provider. ■ System - Use Microsoft Software Shadow Copy Provider. ■ Software - Use Veritas Storage Foundation for Windows. ■ Hardware - Use technology provided by hardware manufacturer. <p>If you select Software or Hardware as the snapshot provider, then the following information applies:</p> <ul style="list-style-type: none"> ■ If multiple volumes are selected, then all volumes must be snappable by the same type of provider. ■ Software and hardware providers cannot both be used to snap different volumes in the same job. You must either create another job, or select the option Process logical volumes for backup one at a time.

Table C-3 Advanced Open File options (*continued*)

Item	Description
Process logical volumes for backup one at a time	<p>Enables the backup of multiple volumes in one job, while creating a snapshot of only one logical volume at a time. To ensure database integrity, or if a volume contains mount points, multiple volumes may need to be snapped at one time. A volume with mount points to other volumes is considered a logical volume for snapshot purposes. Therefore, that volume and the mount point volumes will be snapped together simultaneously.</p> <p>After the logical volume is snapped and backed up, the snapshot is deleted before the next logical volume is snapped. This option increases the ability to meet the minimum quiet time needed to complete a snapshot.</p> <p>A logical volume can comprise multiple physical volumes. A single logical volume can encompass all of the volumes on which databases reside.</p> <p>If this option is not selected, then a snapshot for all volumes in the backup job is created simultaneously. All volumes must meet the minimum quiet time.</p> <p>This option is only available for Symantec Volume Snapshot Provider (VSP) and Microsoft Volume Shadow Copy Service (VSS) jobs for logical volumes.</p> <p>The Shadow Copy Components snapshots are created using VSS. This is reported in the job log and job history.</p>

About the job log and the Advanced Open File Option

When a backup completes successfully while using the Advanced Open File Option (AOFO), information is displayed in the job log. Check the section Backup Set Detail Information in the job log to make sure that AOFO was used during the backup. If the backup included more than one volume, this information is repeated for each volume.

If AOFO fails on initialization, the backup still runs but the job is logged as Completed with Exceptions in the job log.

If AOFO fails during the backup of a device, that backup set is terminated and is reported as an error.

Note: If a job that uses the Symantec Volume Snapshot Provider fails, an active image may remain on the computer. The active image may cause subsequent jobs to fail with an unknown error. If this occurs, restart the computer to clear the active image.

See [“Setting default options for the Advanced Open File Option”](#) on page 943.

See [“Best practices for using the Symantec Volume Snapshot Provider”](#) on page 945.

See [“About the Symantec Volume Snapshot Provider cache file location ”](#) on page 946.

Symantec Backup Exec Agent for DB2 on Windows Servers

This appendix includes the following topics:

- [About the Backup Exec DB2 Agent](#)
- [Requirements for the DB2 Agent](#)
- [Configuring the DB2 Agent on Windows computers](#)
- [Backing up DB2 resources](#)
- [Restoring DB2 data](#)
- [About using DB2 to run DBA-initiated jobs](#)
- [Troubleshooting DB2](#)

About the Backup Exec DB2 Agent

The Symantec Backup Exec Agent for DB2 on Windows Servers (DB2 Agent) protects IBM DB2 databases on Microsoft Windows computers.

The following features are available with the DB2 Agent:

- The ability to initiate backup and restore operations:
 - From Backup Exec.
 - From the IBM DB2 Control Center or command-line processor as a Database Administrator (DBA). Operations that the DBA performs on the Control Center or command-line processor are referred to as DBA-initiated

operations. See your IBM DB2 documentation for information about the Control Center or command-line processor.

- Support for the DB2 log archiving methods that are known as user exit and vendor.
- Multiple data stream support for increased performance during backup and restore operations.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Note: Backup Exec does not support DB2 running as a 32-bit Windows-based application on a 64-bit Windows operating system.

Backup Exec does not support DB2 backup and restore jobs running on the IPv6 protocol.

See “[Requirements for the DB2 Agent](#)” on page 954.

See “[About using DB2 to run DBA-initiated jobs](#)” on page 972.

Requirements for the DB2 Agent

The DB2 Agent is installed as a separate, add-on component of Backup Exec 2010.

To protect local or remote DB2 instances, you must install the following Backup Exec options:

- Backup Exec DB2 Agent on the media server.
See “[Installing additional Backup Exec options to the local media server](#)” on page 123.
- Backup Exec Remote Agent for Windows Systems on remote Windows computers.
See “[About installing the Remote Agent for Windows Systems](#)” on page 139.

After you install the required components, you must configure them for the DB2 Agent before you can back up or restore any DB2 resources.

Do the following:

- On the computer on which the DB2 instances are installed, configure the DB2 Agent.
See “[Configuring the DB2 Agent on Windows computers](#)” on page 955.
- On the media server, configure database access for DB2 operations.

See [“Adding information about DB2 servers to the media server's authentication credentials list”](#) on page 955.

Configuring the DB2 Agent on Windows computers

Before you back up or restore DB2 databases, you must run the Remote Agent Utility to configure the DB2 Agent.

The information that you configure for an instance applies to all of the databases that are contained in that instance.

Whenever DB2 instance information changes, you must update the Remote Agent Utility. If credential information is not updated or is incorrect, the error "Unable to attach to a resource..." may appear when you run a backup job.

Table D-1 DB2 Agent configuration process

Step	Description
Step 1	Add the fully qualified domain name and the logon account name to the media server's list of DB2 servers and authentication credentials. See “Adding information about DB2 servers to the media server's authentication credentials list” on page 955.
Step 2	Set job options for DB2 operations. See “Creating a template for DBA-initiated jobs” on page 424.
Step 3	Configure database access for DB2 operations on Windows computers. See “Configuring database access for DB2 operations on Windows computers” on page 959.

Adding information about DB2 servers to the media server's authentication credentials list

You must add the DB2 server's fully qualified domain name and the logon account name to the media server's authentication credentials list. The media server has database access for operations on DB2 instances that are included in the authentication list. Before you run jobs, use the Remote Agent Utility to configure

instance information and database access on the computer that has the DB2 instances.

The logon account name must have administrative rights to the DB2 server. If the user name is incorrect or is not provided, or if it does not have administrative rights, then you cannot perform DB2 backup or restore operations to that computer.

See [“Creating a template for DBA-initiated jobs”](#) on page 424.

See [“Editing information about DB2 servers on the media server’s list of authentication credentials”](#) on page 956.

See [“Editing DBA-initiated jobs”](#) on page 434.

See [“Deleting information about DB2 servers from the media server’s list of authentication credentials”](#) on page 957.

To configure database access on the media server for DB2 operations

- 1 On the media server, on the **Tools** menu, click **Options**.
- 2 In the properties pane, under **Job Defaults**, click **DB2**.
- 3 Click **Modify list**.
- 4 Click **New**.
- 5 Enter the fully qualified domain name of the DB2 server on which the instance is installed.
- 6 To add the logon account name, do one of the following:

Click the arrow Select the logon account name that you want to add.

Click **New** On the Logon Account Selection dialog box, click **New**.

See [“Creating a Backup Exec logon account”](#) on page 187.

Use the same logon account format that you use when you enter the logon account name on the **Database Access** tab in the Remote Agent Utility. For example, if you entered Domainname\Username on the Remote Agent Utility, use that same format on the list of authentication credentials.

- 7 On the **Authentication Credentials for Oracle and DB2 Servers** dialog box, click **OK**.

Editing information about DB2 servers on the media server’s list of authentication credentials

You must update the media server’s DB2 servers and authentication credentials list if the fully qualified domain name or the logon account name changes. Make

the same changes on the DB2 server by using the Remote Agent Utility to configure instance information and database access.

The logon account name must have administrative rights to the DB2 server. If the user name is incorrect or is not provided, or if it does not have administrative rights, then you cannot perform DB2 backup or restore operations to that computer.

See [“Configuring the DB2 Agent on Windows computers”](#) on page 955.

See [“Adding information about DB2 servers to the media server's authentication credentials list”](#) on page 955.

See [“Deleting information about DB2 servers from the media server's list of authentication credentials”](#) on page 957.

To edit the fully qualified domain name or a logon account on the media server's list of authentication credentials

- 1 On the media server, on the **Tools** menu, click **Options**.
- 2 In the properties pane, under **Job Defaults**, click **DB2**.
- 3 Click **Modify list**.
- 4 Select the item that contains the server name or logon account that you want to edit.
- 5 Click **Edit**.
- 6 Change the fully qualified domain name or change the logon account name.
See [“Editing a Backup Exec logon account”](#) on page 190.
- 7 Click **OK**.

Deleting information about DB2 servers from the media server's list of authentication credentials

Delete the fully qualified domain name or logon account to remove the DB2 server from the media server's list of authentication credentials. If you want to back up the DB2 server again, add the information about the DB2 server to the media server's list of authentication credentials.

See [“Adding information about DB2 servers to the media server's authentication credentials list”](#) on page 955.

To delete the fully qualified domain name and a logon account from the media server's list of authentication credentials

- 1 On the media server, on the **Tools** menu, click **Options**.
- 2 In the properties pane, under **Job Defaults**, click **DB2**.

- 3 Click **Modify list**.
- 4 Select the item that contains the server name or logon account that you want to delete.
- 5 Click **Delete**.
See [“Deleting a Backup Exec logon account”](#) on page 192.
- 6 Click **OK**.

Editing default options for DB2

You can use the defaults that Backup Exec sets during installation for all DB2 backup jobs, or you can choose your own defaults.

See [“Backing up DB2 resources”](#) on page 964.

See [“Restoring DB2 data”](#) on page 967.

See [“Troubleshooting DB2”](#) on page 978.

To edit default options for DB2

- 1 On the **Tools** menu, click **Options**.
- 2 On the properties pane, under **Job Defaults**, click **DB2**.
- 3 Complete the appropriate options.
See [“DB2 default options”](#) on page 958.
- 4 Click **OK**.

DB2 default options

You can edit the defaults that Backup Exec sets during installation for all DB2 backup jobs.

See [“Editing default options for DB2”](#) on page 958.

The following table describes the DB2 default options.

Table D-2 DB2 default options

Item	Description
Backup method	Specifies which of the following backup methods is used for all backup jobs: <ul style="list-style-type: none"> ■ Full - Backup selections. Performs a full backup of DB2 selections. ■ Differential - Backup changes since last full. Backs up all database changes since the last full backup. ■ Incremental - Backup changes since last full or incremental. Backs up all database changes since the last full or incremental backup.
Perform backups offline	Takes the database offline before you start the backup job. Backup Exec brings the database online after the backup job is complete.
Quiesce the database before offline backup	Forces users off the database before it brings the database offline for the backup job. Users who are not actively running databases tasks are forced off the database. Users who are running database tasks can complete their current tasks before being forced off.
Modify list	Lets you add the DB2 computer name and the logon account name to the media server's list of authentication credentials for DB2 servers. See "Adding information about DB2 servers to the media server's authentication credentials list" on page 955.

Configuring database access for DB2 operations on Windows computers

Use the following steps to configure database access for DB2 operations on a Windows computer.

See ["About using the DB2 database archive logging methods"](#) on page 974.

See ["Backing up DB2 resources"](#) on page 964.

See ["Editing DBA-initiated jobs"](#) on page 434.

See ["Editing default options for DB2"](#) on page 958.

To configure database access for DB2 operations on Windows computers

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

When the Remote Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 On the **DB2** tab, enter the appropriate information.
See “[Backup Exec DB2 Agent configuration options](#)” on page 960.
- 3 On the **Database Access** tab, complete the appropriate options.
See “[Database access options for the Remote Agent Utility](#)” on page 1956.
- 4 Click **OK**.
- 5 On the media server, add the following to the media server’s list of authentication credentials:
 - Fully qualified domain name of the DB2 server
 - The user name that you entered on the **Database Access** tab

Backup Exec DB2 Agent configuration options

Use the following table to complete the following options when you use the Remote Agent Utility to configure the DB2 Agent on Windows computers.

See “[Configuring the DB2 Agent on Windows computers](#)” on page 955.

Table D-3 Backup Exec DB2 Agent configuration options

Item	Description
Local instance name	Specifies the name of a local DB2 instance. If you edit an instance, you cannot change the instance name.

Table D-3 Backup Exec DB2 Agent configuration options (*continued*)

Item	Description
User name	<p>Specifies the user name for the DB2 instance.</p> <p>The user name must have the following:</p> <ul style="list-style-type: none"> ■ A valid authorization ID, or the connect privilege to all the databases that reside in this DB2 instance. ■ The correct authority levels and privileges. <p>Correct authority levels may include SYSADM, SYSCTRL, SYSMAINT, and DBADM.</p> <p>If the credentials are incorrect, the error "Unable to attach to a resource..." may appear when you run a backup job.</p> <p>If the credentials for the DB2 instance change, you must update the credentials in this field.</p> <p>You must add this computer name and logon account to the list of authentication credentials for DB2 servers.</p> <p>See "Adding information about DB2 servers to the media server's authentication credentials list" on page 955.</p>
Change password	<p>Launches the Change password dialog box where you can change the password for the DB2 instance user name.</p> <p>See "Enter Password options" on page 339.</p>
Media server	<p>Specifies the name or IP address of the Backup Exec media server where you want to process the operations.</p> <p>You must use the same form of name resolution for all operations. For example, if you use the IP address of this computer for backup operations, you must also use the IP address for restore operations. If you use the full computer name for backup operations, you must also use the full computer name for restore operations.</p>

Table D-3 Backup Exec DB2 Agent configuration options (*continued*)

Item	Description
<p>Job template name</p>	<p>Specifies the name of the Backup Exec job template that you want the DBA-initiated job to use for backup and restore operations. You create the job template on the DBA-initiated Job Settings dialog box on the Backup Exec media server. If you do not specify a job template, the default job template is used.</p> <p>For the databases that have archive logging enabled, enter a separate archive logs template name in the Archive logs template name field.</p> <p>See “Editing DBA-initiated jobs” on page 434.</p>
<p>Archive logs template name</p>	<p>Specifies the name of the Backup Exec archive log template that uses the user exit or vendor method. These methods are specified for a database on the DB2 Control Center or command-line processor. If you use an incorrect job template name, then log files cannot be archived correctly.</p> <p>If you use the user exit or vendor method for a database, you must create a Backup Exec job template specifically for archiving logs. The template should specify the destination devices that are different than the devices that are specified in the job template for database backups.</p> <p>If the same device is used for both jobs, then the archive log backup must wait until the database backup completes. However, the database backup cannot complete until the archive log backup completes. If this situation occurs when the device is a backup-to-disk folder, increase the number of concurrent operations that are allowed on the backup-to-disk folder.</p> <p>See “Editing DBA-initiated jobs” on page 434.</p> <p>See “About using the DB2 database archive logging methods” on page 974.</p>

Adding a DB2 instance to the DB2 Agent on Windows computers that run the Remote Agent Utility

Use the following steps to add a DB2 instance to the DB2 Agent on Windows computers that run the Remote Agent Utility.

See [“About the Remote Agent Utility for Windows Systems”](#) on page 1948.

See [“Remote Agent Utility Command Line Applet switches”](#) on page 1959.

To add a DB2 instance to the the DB2 Agent on Windows computers that run the Remote Agent Utility

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

When the Remote Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 On the **DB2** tab, click **New**.
- 3 Complete the appropriate options.
See [“Backup Exec DB2 Agent configuration options”](#) on page 960.
- 4 Click **OK**.

Editing a DB2 instance by using the Remote Agent Utility

Use the following steps to edit a DB2 instance by using the Remote Agent Utility.

To edit a DB2 instance by using the Remote Agent Utility

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.
- 2 On the **DB2** tab, click **Edit**.
- 3 Edit the appropriate options.
See [“Backup Exec DB2 Agent configuration options”](#) on page 960.
- 4 Click **OK**.

Deleting a DB2 instance by using the Remote Agent Utility

Use the following steps to delete a DB2 instance by using the Remote Agent Utility.

To delete a DB2 instance by using the Remote Agent Utility

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.
- 2 On the **DB2** tab, click **Delete**.

Backing up DB2 resources

Before you back up DB2 resources, review the following:

- You must run the Remote Agent Utility on the DB2 server and add information about the instances before you can perform any backup or restore operations. When DB2 instance information changes, you must update the Remote Agent Utility. After you enter these changes, the Backup Exec media server discovers them.
See [“Requirements for the DB2 Agent”](#) on page 954.
- If you run database backups and archive logging, you must have at least two storage devices that are available for the jobs.
- If you use multiple data streams for the backup job, the number of backup devices that are available for the job must at least equal the number of data streams. If archive logging is enabled for the database, an additional backup device must be available.

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux Servers is installed.

See [“Creating a backup job by setting job properties”](#) on page 332.

See [“Editing DBA-initiated jobs”](#) on page 434.

See [“Editing default options for DB2”](#) on page 958.

To back up DB2 resources

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selection list, under **Favorite Resources**, expand **Windows Systems**.

- 4 Expand the DB2 server that you want to back up.
If the DB2 server is not listed under **Favorite Resources**, then you can add it.
See [“About the Favorite Resources node in the backup selections list”](#) on page 284.
- 5 Select the following objects that you want to back up:

Instance	Specifies a database instance to back up. All databases within the instance are backed up.
Database	Specifies a database to back up. All partitions within the database are also backed up. During a file system backup, online DB2 database files that are included in the selecton list are not automatically excluded. You must manually exclude the datafiles of an online DB2 database from the selection list.
Partition	Specifies a partition to back up. All partition tablespaces and log folders are backed up.
Tablespace	Specifies all of the tablespaces or individual tablespaces that you want to back up.
- 6 On the Properties pane, under **Settings**, click **DB2**.
- 7 Complete the appropriate options.
See [“DB2 backup options”](#) on page 966.
- 8 To configure multiple data streams for backup, under **Destination**, click **Device and Media**.
- 9 Complete the appropriate options as follows:

Maximum number of devices to use for resources that support multiple data streams	<p>Specifies the maximum number of devices that the backup job can use.</p> <p>When you run a DB2 database backup job, Symantec recommends that the number of backup devices that are available for the job is at least equal to the number of streams. If archive logging is enabled for the database, an additional backup device must be available.</p> <p>If you specify more than one device, you must choose one of the following items as a destination device for the backup job:</p> <ul style="list-style-type: none">■ A device pool.■ A backup-to-disk folder that has at least two concurrent operations enabled. <p>See “Creating a backup-to-disk folder by setting properties” on page 499.</p> <p>This feature is not available for DBA-initiated jobs.</p>
Minimum number of devices, terminate job if fewer devices are available	<p>Specifies the minimum number of devices that the job can use.</p> <p>If the job cannot acquire the minimum number of devices, the job fails.</p> <p>This feature is not available for DBA-initiated jobs.</p>

10 Complete the remaining backup job properties as necessary.

DB2 backup options

You can set specific backup options for DB2 resources when you create a backup job.

See [“Backing up DB2 resources”](#) on page 964.

The following table describes the DB2 backup options:

Table D-4 DB2 backup options

Item	Description
Backup method	<p>Specifies which of the following backup methods to use for the backup job:</p> <ul style="list-style-type: none"> ■ Full - Backup selections. Performs a full backup of DB2 selections. ■ Differential - Back up changes since last full. Backs up all database changes since the last full backup. ■ Incremental - Back up changes since last full or incremental. Backs up all database changes since the last full or incremental backup.
Perform the backup offline	<p>Takes the database offline before you start the backup job. Backup Exec brings the database online after the backup job is complete.</p> <p>If circular logging is enabled for the database, then you must select this option; otherwise the backup job fails.</p>
Quiesce the database before offline backup	<p>Forces all users off the database before it brings the database offline for the backup job. Users who are not actively running databases tasks are forced off the database. Users who are running database tasks can complete their current task before they are forced off.</p>

Restoring DB2 data

Before you restore DB2 resources, make sure that you have completed all of the preparations for installing and configuring the DB2 Agent.

See [“Requirements for the DB2 Agent”](#) on page 954.

Note: In a CASO environment, you can delegate a DB2 restore job to a managed media server. However, if the restore job uses encrypted DB2 backup sets from which to restore, the restore job may fail. An error message may appear that indicates the managed media server does not have the required encryption keys necessary to complete the job. You must create the encryption keys on the managed media server that runs the restore job.

See [“Creating an encryption key”](#) on page 420.

To restore DB2 data

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 On the properties pane, under **Source**, click **Selections**.

- 4 On the **View by Resource** tab, expand the **All Resources** icon.
- 5 Expand the system resource that contains the database instance that you want to restore.
- 6 Do one of the following:

To restore the entire database	Check Current Database . If the database is offline, the Current Database is empty and you cannot make selections.
To restore a backup set	Check a backup set, or a historical set.
To restore a table space	Expand either Current Database or a backup set, and then check any of the table spaces that you want to restore.
- 7 On the properties pane, under **Settings**, click **DB2**.
- 8 Complete the appropriate options.
See [“DB2 restore options”](#) on page 968.
- 9 Select other restore options from the Properties pane as appropriate, and then start the restore job.
See [“Restoring data by setting job properties”](#) on page 605.
- 10 Run a full backup of the restored database.

DB2 restore options

When you use the Agent for DB2 to create a restore job, you can select specific restore options.

See [“Restoring DB2 data”](#) on page 967.

The following table describes the restore options available for DB2.

Table D-5 DB2 restore options

Item	Description
Restore database from full and/or incremental backups	Enables restore options.
To the most recent available	Restores the DB2 database to the most recent full and incremental backups that are available.

Table D-5 DB2 restore options (*continued*)

Item	Description
To a point in time	<p>Restores data up to and including a point in time that is in the job log. After the point in time, recovery stops.</p> <p>In the Date box, select the part of the date that you want to change, and then enter a new date or click the arrow to display a calendar from which you can select a date.</p> <p>In the Time box, select the part of the time that you want to change, and then enter a new time or click the arrows to select a new time.</p>
Roll forward using logs	<p>Enables a restore of a database that has archive logging enabled. To restore a database completely, you must also check Restore from full and/or incremental backups.</p> <p>You can restore the database, and then perform the roll forward operation later.</p> <p>For databases that have circular logging enabled, you must uncheck Roll forward using logs or the job will fail.</p>
To the most recent available	<p>Rolls the DB2 database forward to the most recent logs that are available.</p>
To a point in time	<p>Restores logs up to and including a point in time. After the point in time, recovery stops.</p> <p>In the Date box, select the part of the date that you want to change, and then enter a new date or click the arrow to display a calendar from which you can select a date.</p> <p>In the Time box, select the part of the time that you want to change, and then enter a new time or click the arrows to select a new time.</p>
Override default log file path on DB2 server for this job	<p>Specifies an alternate log file path location to search for archive log files during the roll forward operation. You should specify an alternate location if archive log files were moved to a location other than the location that was specified by the logpath database configuration on the destination DB2 server. Type the full path for the archive logs location.</p> <p>Selecting this option does not change the archive log path configuration on the DB2 server.</p>
Bring database online when finished rolling forward	<p>Ensures that the database is brought online as soon as the recovery is finished.</p>

Redirecting a restore of DB2 data

The instance that you are redirecting DB2 data to must already exist. Backup Exec does not create a new instance.

You cannot redirect restores of the DMS containers or SMS containers in the following situations:

- You redirect a database restore to another database.
- You select a database that does not exist.

To redirect a restore of DB2 data

- 1 Create a restore job.
See [“Restoring DB2 data”](#) on page 967.
- 2 After you select options on the **Restore Job Properties** dialog box for DB2, on the properties pane, under **Destination**, click **DB2 Redirection**.
- 3 Select the appropriate options.
See [“DB2 redirection options”](#) on page 970.
- 4 Start the redirected restore job or select other restore options from the properties pane.

After the restore job is complete, Symantec recommends that you run a full backup of the restored data.

See [“Troubleshooting DB2”](#) on page 978.

DB2 redirection options

You can redirect DB2 data to another instance provided the instance already exists.

See [“Redirecting a restore of DB2 data”](#) on page 970.

The following table describes the Redirection options available for DB2:

Table D-6 DB2 redirection options

Item	Description
Redirect DB2 instance to server	Redirects the restore of the DB2 instance to a server other than the source server.
Server	Specifies the name of the server to which you want to redirect the restore job.

Table D-6 DB2 redirection options (*continued*)

Item	Description
Server logon account	Specifies a logon account that has rights to restore data to the server to which you want to redirect the restore job.
Redirect to a new instance	Redirects the restore of a database to another instance. The instance that you want to redirect to must already exist or the job fails.
Instance	Specifies the name of the instance to which you are redirecting the restore of the database.
Instance logon account	Specifies the instance logon account. If you want to change the logon account, type a logon account for the database that you want to restore. This logon account must have backup operator or administrator privileges.
Restore to a new database	Redirects the restore of the database to a new database other than the source server.
Database name	Indicates the name of the database to which you want to redirect the restore job.
Drive to restore to	Indicates a drive to which you want to redirect the DB2 database.
Restore log location	Specifies the full path of the location where you want the log files for the new database to reside.
Redirect containers	Specifies a different location for the DMS and/or SMS containers for the tablespaces or database that you want to restore.

Table D-6 DB2 redirection options (*continued*)

Item	Description
<p>SMS tablespace container relative path</p>	<p>Indicates the path to which you want to redirect the system managed space (SMS) table spaces containers.</p> <p>For example, you restore the tablespace TS1, which has the following SMS and DMS containers:</p> <p>C:\TS1Containers\SMS\SMSCONT001\ C:\TS1Containers\SMS\SMSCONT002\ C:\TS1Containers\DMS\DMSCONT001 C:\TS1Containers\DMS\DMSCONT002</p> <p>You can redirect the SMS and DMS containers to different locations by typing the path for the SMS container as D:\TS1SMS\. Type the path for the DMS container as D:\TS1DMS\. When the tablespace is restored, the containers are restored as follows:</p> <p>SMS containers D:\TS1SMS\SMSCONT001\ D:\TS1SMS\SMSCONT002\ DMS containers D:\TS1DMS\DMSCONT001 D:\TS1DMS\DMSCONT002</p>
<p>DMS tablespace container relative path</p>	<p>Indicates the path to which you want to redirect the database managed space (DMS) table space containers.</p> <p>See the example in the description of the SMS tablespace container relative path field.</p>

About using DB2 to run DBA-initiated jobs

Backup Exec supports DBA-initiated backup, restore, redirected restore, and recovery of DB2 databases. When you run a DBA-initiated DB2 job, you configure and start the job by using DB2 - not Backup Exec. After the DBA-initiated job

starts, you can monitor the job by using Backup Exec. All DBA-initiated jobs appear in the Backup Exec Job Monitor tab.

The following table describes the files that are installed on the Backup Exec media server during the DB2 Agent installation:

Table D-7 Files that are installed on the media server with the DB2 Agent

File	Description
The vendor dll file db2sqluv.dll and the user exit program db2uext2.exe	<p>Installed in the Windows System Directory. The Windows System Directory may have a pathname such as:</p> <p>C:\winnt\system32 or D:\windows\system32.</p> <p>You can use a vendor library or a user exit program from the DB2 Control Center as a method for archiving log files. However, doing so causes db2sqluv.dll and db2uext2.exe to be used by default.</p> <p>See “About using the DB2 database archive logging methods” on page 974.</p>
A configuration file named db2.conf	<p>Includes specifications for redirected restore jobs and roll forward operations. The Backup Exec vendor dll file and user exit program use the information in this file.</p> <p>The db2.conf file is installed to the following location on the media server:</p> <p>\Program Files\Symantec\Backup Exec\db2.conf</p> <p>The db2.conf file is installed to the following location on the remote DB2 server:</p> <p>\Program Files\Symantec\Backup Exec\RAWS\db2.conf</p> <p>See “About the db2.conf file” on page 975.</p>
Example scripts for backup and restore operations	<p>Available for you to run at the DB2 command-line processor.</p> <p>These scripts are installed to the following location:</p> <p>\Program Files\Symantec\Backup Exec\scripts\DB2</p>

Review the following notes before you run DBA-initiated jobs for DB2:

- Complete all of the preparations for installing and configuring the DB2 Agent. See [“Requirements for the DB2 Agent”](#) on page 954.
- If you use a domain administrator logon account to browse DB2 databases on a DB2 server, you may not be able to expand or select databases for operations from Backup Exec. If this situation occurs, add the domain administrator account to the DB2ADMNS group.

- On the media server, the logon account that you use to back up the DB2 resources should have backup operator or administrator rights.
- With DBA-initiated jobs in a CASO environment, the destination device that you select in the DBA-initiated job template must be locally attached to the central administration server. This includes DBA-initiated DB2 archive log jobs.
If the destination device includes a device pool, all devices in the pool must be locally attached to the central administration server.

See [“Troubleshooting DB2”](#) on page 978.

See [“About using the DB2 database archive logging methods”](#) on page 974.

About using the DB2 database archive logging methods

DB2 supports the user exit and vendor methods for archiving its log files. Backup Exec provides a user exit program and a vendor dll file to support these methods. When you use the user exit method, Backup Exec backs up the archive logs by using the user exit program named `db2uext2.exe`. When you use the vendor method, Backup Exec backs up the archive logs by using the Backup Exec vendor dll file named `db2sqluv.dll`.

Before you can use the user exit or vendor method, you must add information in the Remote Agent Utility about the DB2 instances that contain the following:

- The source database for the archive logging operations.
- The destination database for any roll forward operations.

You must also add the fully qualified domain name of the DB2 server that contains these instances to the media server’s list of DB2 servers and authentication credentials.

If you use archive logging for DB2 databases, you must create a Backup Exec DBA-initiated job settings template that is used exclusively by archive logging jobs. This job template must specify destination storage devices that are different than the devices that are specified in the job template that is used for database backups. You must add the name of the DBA-initiated job settings template for archive logs in the Remote Agent Utility.

Some errors for DBA-initiated jobs that use the vendor dll file `db2sqluv.dll` are described in the following table:

Table D-8 Errors that may occur while using db2sqluv.dll

Error	Description
514	Backup Exec cannot find the logon account information that is required for database access. Ensure that the information has been updated in the media server's list of authentication credentials.
SQL2062N	See the Application log in the Windows Event Viewer for details about the error.

See [“Configuring the DB2 Agent on Windows computers”](#) on page 955.

See [“Editing DBA-initiated jobs”](#) on page 434.

See [“Troubleshooting DB2”](#) on page 978.

See [“About the db2.conf file”](#) on page 975.

About the db2.conf file

The Backup Exec db2.conf file provides the settings for DBA-initiated redirected restore jobs and roll forward jobs. You must configure the required settings in the db2.conf file before you can run DBA-initiated redirected restore or roll forward jobs.

The Backup Exec DB2 configuration file, db2.conf, consists of a series of keywords and values that define how to back up the database and the archive logs. Use this file to define the source database and the source instance for the redirected restore operations and redirected roll forward operations.

Instructions and examples are in the db2.conf file.

See [“Editing DBA-initiated jobs”](#) on page 434.

See [“About using DB2 to run DBA-initiated jobs”](#) on page 972.

See [“Editing a db2.conf file”](#) on page 975.

Editing a db2.conf file

The db2.conf file contains two blocks of settings. The first block contains the settings that you can use to perform a redirected database restore using Backup Exec's vendor dll db2sqluv.dll. The second block contains the settings that you can use to perform a redirected roll forward of a database using the Backup Exec vendor dll db2sqluv.dll. or user exit program db2uext2.exe.

See [“Editing DBA-initiated jobs”](#) on page 434.

See [“About using DB2 to run DBA-initiated jobs”](#) on page 972.

See “[Example db2.conf file](#)” on page 976.

To edit a db2.conf file

- 1 On the computer on which the DB2 instances that you want to redirect are installed, open the db2.conf file for editing.

On the media server, the db2.conf file is located at the following path:

```
\Program Files\Symantec\Backup Exec\db2.conf
```

On remote DB2 servers, the db2.conf file is located at the following path:

```
\Program Files\Symantec\Backup Exec\RAWS\db2.conf
```

The db2.conf file consists of keyword lines that form object identifiers. The lines in each object identifier specify the database and other information.

- 2 Remove the pound sign (#) that precedes the lines and add the appropriate information.
- 3 Save and close the db2.conf file.
- 4 Repeat step 1 - step 3 on each DB2 server that you want to back up with the DB2 Agent.

After you complete a redirected restore job or a roll forward, you must remove the instructions for that database. If you do not remove the instructions, they apply to all subsequent restore operations.

Example db2.conf file

The sample db2.conf file is as follows:

```
#
# The following settings are used by Backup Exec to perform
# an alternate restore or a rollforward of a DB2 database during a
# a DBA-initiated operation using Backup Exec's vendor dll
# db2sqluv.dll or user exit program db2uext2.exe.
# Reminders:
#
# Uncomment the following lines by removing the # preceding every line
# and add appropriate data to perform an alternate restore/rollforward
# operation.
# You can add more blocks for any additional alternate restore/rollforward
# operations.
# -----
# Settings for alternate database restore using
```

```
# Backup Exec's vendor dll db2sqluv.dll
# -----

#OBJECTTYPE ALTERNATE      # Specifies an alternate restore
#SRCINST srcinstname      # Names the source instance that was backed up
#SRCALIAS srcaliasname    # Names the source database alias that was backed up
#DESTINST destinstname    # Names the destination instance name
#DESTALIAS destaliasname  # Names the destination database alias name
#ENDOPER                  # Ends the object identifier

OBJECTTYPE ALTERNATE      # Specifies an alternate restore
SRCINST myinst1          # Names the source instance that was backed up
SRCALIAS mydb1          # Names the source database alias that was backed up
DESTINST myinst2        # Names the destination instance name
DESTALIAS mydb2        # Names the destination database alias name
ENDOPER                  # Ends the object identifier

OBJECTTYPE ALTERNATE      # Specifies an alternate restore
SRCINST myinst3          # Names the source instance that was backed up
SRCALIAS mydb3          # Names the source database alias that was backed up
DESTINST myinst4        # Names the destination instance name
DESTALIAS mydb4        # Names the destination database alias name
ENDOPER                  # Ends the object identifier

# -----
# The following are settings for an alternate database rollforward operation
# using Backup Exec's vendor dll db2sqluv.dll or user exit program
# db2uext2.exe. Use this block to indicate the source database
# if the log files were archived from a different source database.
# -----
# If DB2 log file archiving is enabled (DB2 USEREXIT ON), DB2 will invoke
# the Backup Exec user exit program to back up and restore DB2 archive
# log files. If DB2 log file archiving is enabled for vendor dll, DB2 will
# Backup Exec vendor dll db2sqluv.dll to back up and restore DB2 archive log files.
# invoke the DESTALIAS parameter indicates the destination database alias
# for the user exit. DESTINST parameter indicates the destination instance alias
# for the user exit. SRCALIAS parameter indicates the source database alias
# from which log files were archived and should now be used for the
# rollforward operation. SRCINST parameter indicates the source instance from
# which log files were archived and should now be used for the rollforward
# operation.
#
```

```
#OBJECTTYPE ARCHIVE # Specifies that this block is for
#alternate rollforward.
#ARCFUNC SAVE
#DESTALIAS destaliasname # Names the destination database alias name
for which this setting applies.
#DESTINST destinstname # Names the destination instance name name
for which this setting applies.
#SRCALIAS srcaliasname # Names the source database alias from which
log files were archived.
#SRCINST srcinstname # Names the source instance from which log
files were archived.
#ENDOPER # Ends the object identifier

OBJECTTYPE ARCHIVE # Specifies that this block is for alternate
rollforward.
ARCFUNC SAVE
DESTALIAS mydb1 # Names the destination database alias name for
which this setting applies.
DESTINST myinst1 # Names the destination instance name for
which this setting applies.
SRCALIAS mydb2 # Names the source database alias from which
log files were archived.
SRCINST myinst1 # Names the source instance from which
log files were archived.
ENDOPER # Ends the object identifier
```

Troubleshooting DB2

What should I do if roll forward operations fail when I redirect a restore of DB2?

For complete recovery of databases for which archive logging is used, you must restore both the database and the archived logs. This operation is known as roll forward. If you used the vendor method, then db2sqluv.dll is located in the Windows system directory. This directory can be located in different paths, such as:

C:\winnt\system32 or D:\windows\system32

Information about the archival method and the location of the db2sqluv.dll is in each backup of the database. If you restore the backup to another computer, the information about the location of the db2sqluv.dll points to the same path that is on the source computer. However, the computer that you redirected the restore to may have a different path for the Windows system directory. For example, on

the source computer, the db2sqluv.dll may be located in the Windows system directory, in the path:

C:\winnt\system32

On the computer that you redirect the restore to, the Windows system directory may use another path, such as:

D:\windows\system32

When you run the roll forward on the restored database, DB2 attempts to run db2sqluv.dll from the Windows system directory path on the source computer. The roll forward fails because the db2sqluv.dll is not located in that path. To run a successful roll forward, db2sqluv.dll must exist in the same path on source and destination computers.

Symantec Backup Exec Agent for Enterprise Vault

This appendix includes the following topics:

- [About the Enterprise Vault Agent](#)
- [Requirements for the Enterprise Vault Agent](#)
- [About installing the Enterprise Vault Agent](#)
- [About backup methods for Enterprise Vault backup jobs](#)
- [About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)
- [Backing up an Enterprise Vault open partition](#)
- [Backing up Enterprise Vault closed partitions](#)
- [Backing up Enterprise Vault 8.x ready partitions](#)
- [Backing up the Enterprise Vault Directory database](#)
- [Backing up the Enterprise Vault Monitoring database](#)
- [Backing up an Enterprise Vault vault store database](#)
- [Backing up the Enterprise Vault 8.x Audit database](#)
- [Backing up the Enterprise Vault 8.x FSA Reporting database](#)
- [Backing up the Enterprise Vault 8.x Fingerprint database](#)
- [Backing up the Enterprise Vault 8.x Compliance Accelerator Configuration database and Compliance Accelerator customer databases](#)

- [Backing up the Enterprise Vault 8.x Discovery Accelerator Configuration database and Discovery Accelerator customer databases](#)
- [Backing up the Discovery Accelerator Custodian database](#)
- [Backing up an Enterprise Vault vault store](#)
- [About backing up an Enterprise Vault 7.x server and an Enterprise 8.x site](#)
- [About restoring Enterprise Vault](#)
- [Best practices for the Enterprise Vault Agent](#)
- [About the Backup Exec Migrator for Enterprise Vault](#)

About the Enterprise Vault Agent

The Symantec Backup Exec Agent for Enterprise Vault (Enterprise Vault Agent) is installed as a separate, add-on component of Backup Exec. The Enterprise Vault Agent provides data protection for both Enterprise Vault 7.x and 8.x components:

Enterprise Vault 7.x components include the following:

- Sites
- Servers
- Databases
- Indexes
- Vault partitions

Enterprise Vault 8.x components include the following:

- Sites
- Vault Store Groups
- Databases
- Indexes
- Vault partitions

The Enterprise Vault Agent can help provide a disaster recovery solution for data that is archived with Enterprise Vault. Recovery of the archived data is not dependent on the archive source, such as Exchange Server or a specific file system.

The Enterprise Vault Agent lets you do the following:

- Back up and restore Enterprise Vault archives from open or closed vault store partitions.

- Back up and restore individual Enterprise Vault 7.x servers and Enterprise Vault 8.x vault store groups from an Enterprise Vault site.
- Back up and restore Enterprise Vault sites, databases, and index locations.

When you back up Enterprise Vault servers the following Enterprise Vault components can be backed up along with the vault partitions:

- Enterprise Vault Directory and Monitoring databases
- Enterprise Vault Audit, FSA Reporting, and Fingerprint databases (Enterprise Vault 8.x only)
- Enterprise Vault vault store databases
- Enterprise Vault indexing files

If you install the Enterprise Vault Compliance and Discovery Accelerator products, the following components can be backed up:

- Enterprise Vault Compliance Accelerator and Discovery Accelerator Configuration databases
- Enterprise Vault Compliance and Discovery Accelerator Customer databases
- Enterprise Vault Discovery Accelerator Custodian database

The Enterprise Vault Agent backs up Enterprise Vault 7.x by placing the Enterprise Vault 7.x services in read-only mode. Read-only mode ensures that database, database index, and vault store synchronization remains while the backup job runs. As a result, Enterprise Vault does not archive new data while in read-only mode.

After the backup job successfully completes, the Enterprise Vault Agent takes the Enterprise Vault services out of read-only mode so that archival operations can continue.

With later versions of Enterprise Vault (Enterprise Vault 8.x), the Enterprise Vault Agent uses Enterprise Vault Backup mode rather than read-only mode to back up Enterprise Vault components. By using the Backup mode, the Enterprise Vault Agent can back up Enterprise Vault components without having to suspend Enterprise Vault archiving operations.

For example, when you select a vault store group or site for backup, the individual vault store or indexes are placed in Backup mode. Backup mode lets Enterprise Vault continue archiving operations in other vault store groups or sites. After the backup job successfully completes, the Enterprise Vault Agent takes the Enterprise Vault components out of Backup mode so that those components can continue archival operations.

The Enterprise Vault Agent backs up the Compliance Accelerator and Discovery Accelerator application databases while they are online. It does not place the databases in Read-only mode or Backup mode before it backs them up.

The Enterprise Vault Agent runs a physical check on each Enterprise Vault database before it backs them up. The Enterprise Vault Agent also runs a physical check on each database before you restore them.

Note: The Enterprise Vault Agent uses physical database consistency checks because physical checks consume less system resources than other types of Database Consistency Checking options.

Backing up and restoring Enterprise Vault 8.x databases and related components required specific user account credentials for each Enterprise Vault component you protect.

When you back up specific Enterprise Vault components, other Enterprise Vault components are automatically backed up at the same time. Backup Exec includes these components to hasten an Enterprise Vault recovery.

Table E-1 Supported user accounts that are required to back up and restore Enterprise Vault 8.x components

Enterprise Vault 8.x components	User credentials
Enterprise Vault 8.x databases and components (vault store, indexes, partitions, vault store database, Directory, Monitoring, Fingerprint, FSA Reporting, and Audit databases)	<p>The following credentials are required:</p> <ul style="list-style-type: none"> ■ Vault Service account ■ Domain Admin account with Role Based Admin privileges <p>You can also use any domain user account that meets the following requirements:</p> <ul style="list-style-type: none"> ■ The user account must be included in the Administrators group on all Enterprise Vault servers. ■ The user account must be included in the Backup Operators group on all computers where Enterprise Vault databases reside. With Enterprise Vault partitions, the user account must be included in the Administrators group. ■ The user account must have backup-related Role Based Admin privileges for the vault store and the index location. <p>Backup-related Role Based Admin privileges include:</p> <ul style="list-style-type: none"> ■ EVT Manage Vault Store Backup Mode ■ EVT Manage Index Location Backup Mode <p>To configure Role Based Admin privileges for a Windows Domain Admin account, see your Enterprise Vault documentation.</p> <p>Note: To back up a computer that has a partition, or a partition and a database, you must be a member of the computer's Administrator's group. To back up a computer that has only an Enterprise Vault database on it, you need only to be a member of the Backup Operators group.</p>

Table E-1 Supported user accounts that are required to back up and restore Enterprise Vault 8.x components (*continued*)

Enterprise Vault 8.x components	User credentials
Compliance Accelerator and Discovery Accelerator	<p>The following credentials are required:</p> <ul style="list-style-type: none"> ■ Vault Service account ■ Domain Admin account <p>A user account that is a member of the following groups:</p> <ul style="list-style-type: none"> ■ Administrator's group on those computers where the Compliance and Discovery Accelerator applications reside. ■ Administrator's group and Backup Operator's group on those computers where the Accelerator databases reside.

When you back up specific Enterprise Vault components, the other Enterprise Vault components are automatically backed up at the same time. Backup Exec includes these components to hasten an Enterprise Vault recovery.

Table E-2 Enterprise Vault components that are automatically backed up

When you back up this	Backup Exec automatically backs up this	Description
Enterprise Vault site	Directory database	Backup Exec automatically backs up the Directory database that is associated with the Enterprise Vault site.
Enterprise Vault 7.x server	Directory database	After Backup Exec sequentially backs up all of the selected servers, it finishes the backup job by automatically backing up the Directory database. Backup Exec places the Directory database in a backup set that remains separate from the sets that contain the server components.

Table E-2 Enterprise Vault components that are automatically backed up
(continued)

When you back up this	Backup Exec automatically backs up this	Description
Open partition	Vault store database	Backup Exec automatically backs up the vault store database that is associated with the open partition.

Over time the amount of data that Enterprise Vault stores continues to grow. At some point, you may observe that as the data moves through its usage lifecycle, you no longer access it as frequently. You can use the Backup Exec Migrator for Enterprise Vault to automatically migrate the older Enterprise Vault data to the storage devices that Backup Exec manages.

See [“About the Backup Exec Migrator for Enterprise Vault”](#) on page 1041.

Requirements for the Enterprise Vault Agent

Review the following requirements before you use the Enterprise Vault Agent.

- You must create at least one partition on an Enterprise Vault server before the Enterprise Vault server can publish itself to Backup Exec.
- You must install the Backup Exec Remote Agent for Windows Systems and license the Enterprise Vault Agent on any computer that hosts an Enterprise Vault component.

Note: The Enterprise Vault Agent uses the Remote Agent to back up all NTFS shares on a remote computer that contains Enterprise Vault data. However, if the Remote Agent is not installed, the Enterprise Vault Agent uses Microsoft’s Common Internet File System (CIFS) to back up the data.

For a device or a filer that does not support the Remote Agent, the Enterprise Vault Agent uses CIFS to back up the data. Symantec recommends that you create separate backup jobs when you want to do NDMP backups of Enterprise Vault data. You may see a significant performance improvement of NDMP backups with the Symantec Backup Exec NDMP Option.

About installing the Enterprise Vault Agent

The Enterprise Vault Agent is installed locally as a separate, add-on component of Backup Exec. To back up all Enterprise Vault servers, the Enterprise Vault Agent must be installed on each Enterprise Vault server in your environment. In addition, the Enterprise Vault Agent must also be installed on any remote computer where Enterprise Vault components are installed. If the Compliance and Discovery Accelerators are installed on remote computers, the Enterprise Vault Agent must be installed on those computers too.

Note: You cannot back up Enterprise Vault databases to devices that are attached to a computer on which the Remote Media Agent for Linux Servers is installed.

You can install the Enterprise Vault Agent in the following ways:

- Install it automatically from the Backup Exec media server as part of a Remote Agent installation to the local Enterprise Vault server. After you finish the installation, you may need to configure the Enterprise Vault Agent to publish itself to a media server of your choice.

See [“About publishing the Remote Agent for Windows Systems to media servers”](#) on page 1951.

- Install the required Enterprise Vault Agent license keys on the media server. After you install the license keys, you can push-install the Backup Exec Remote Agent to all Enterprise Vault servers and the computers where other Enterprise Vault components are installed.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

See [“Push-installing the Remote Agent and Advanced Open File Option to remote computers”](#) on page 134.

About backup methods for Enterprise Vault backup jobs

You can select a backup method that depends on the Enterprise Vault object that you want to backup.

The following table describes the type of Enterprise Vault backup jobs you can run. The table also describes the backup methods that are available for each type of backup job.

Table E-3 Backup methods to use with Enterprise Vault backup jobs

To back up:	Select:	Description
Directory and Monitoring databases Audit database and FSA Reporting database (Enterprise Vault 8.x only)	Full, differential, or incremental backup method	Directory, Monitoring, Audit, and FSA Reporting database backups can use the full and incremental backup methods. These databases cannot be backed up using the differential backup method. If you select the differerential backup method, Backup Exec does a full backup instead . Note: Selecting an incremental backup method backs up the database transaction logs and them truncates them.
Vault database and Fingerprint database	Full, differential, or incremental backup method	Vault database and Fingerprint database backups can use all three backup methods: Full, differential, and incremental. Note: Selecting an incremental backup method backs up the database transaction logs and them truncates them.
Vault partitions and index locations	Full, differential, or incremental backup methods.	You can use all of the backup methods that are available for standard file system backup jobs.

When you combine Enterprise Vault components in a backup job, each component may use a backup method that differs from what you selected for the overall job. For example, you create a job that uses the differential backup method to back up both a Directory database and a partition. However, because a Directory database cannot be backed up using the differential method, Backup Exec uses the full backup method to back up the Directory database. This results in fast and easy restores. After the Directory database is backed up, Backup Exec uses the differential backup method to back up the partition.

Use the following table as a guide.

Table E-4 Actual backup methods that are used for Enterprise Vault components

Enterprise Vault component	Full (F)	Differential (D)	Incremental (I)
Directory and Monitoring databases	F	F	I Always truncates the transaction logs
Vault store database	F	D	I Always truncates the transaction logs
Audit database (Enterprise Vault 8.x only)	F	F	I Always truncates the transaction logs
FSAReporting database (Enterprise Vault 8.x only)	F	F	I Always truncates the transaction logs
Fingerprint database (Enterprise Vault 8.x only)	F	D	I Always truncates the transaction logs
Partition	F	D	I Always truncates the transaction logs
Index root path	F	D	I Always truncates the transaction logs

Table E-4 Actual backup methods that are used for Enterprise Vault components (*continued*)

Enterprise Vault component	Full (F)	Differential (D)	Incremental (I)
Compliance Accelerator/Discovery Accelerator Configuration database (Enterprise Vault 8.x only) Note: Also includes the Compliance Accelerator and Discovery Accelerator databases that are installed with runtime versions of Enterprise Vault.	F	F	I Always truncates the transaction logs
Compliance Accelerator/Discovery Accelerator Customer database (Enterprise Vault 8.x only) Note: Also includes the Compliance Accelerator and Discovery Accelerator databases that are installed with runtime versions of Enterprise Vault.	F	D	I Always truncates the transaction logs

Table E-4 Actual backup methods that are used for Enterprise Vault components (*continued*)

Enterprise Vault component	Full (F)	Differential (D)	Incremental (I)
Discovery Accelerator Custodian database (Enterprise Vault 8.x only) Note: Also includes the Discovery Accelerator Custodian databases that are installed with runtime versions of Enterprise Vault.	F	D	I Always truncates the transaction logs

See [“About backup methods”](#) on page 274.

See [“Backing up an Enterprise Vault open partition”](#) on page 994.

See [“Backing up Enterprise Vault closed partitions”](#) on page 995.

See [“Backing up the Enterprise Vault Directory database”](#) on page 998.

See [“Backing up the Enterprise Vault Monitoring database”](#) on page 999.

See [“Backing up an Enterprise Vault vault store database”](#) on page 1000.

See [“Backing up an Enterprise Vault vault store”](#) on page 1007.

See [“Backing up an Enterprise Vault 7.x server”](#) on page 1009.

See [“Backing up an Enterprise Vault site”](#) on page 1010.

See [“Backing up Enterprise Vault index locations”](#) on page 1011.

Enterprise Vault backup options

You can select a backup method that is based on the type of Enterprise Vault database that you want to back up.

See [“About backup methods for Enterprise Vault backup jobs”](#) on page 988.

Setting a default backup method for Enterprise Vault backup jobs

You can set a default backup method that you can use for all Enterprise Vault backup jobs.

In some cases, Backup Exec may override the default backup method when you run a backup job.

See [“About backup methods for Enterprise Vault backup jobs”](#) on page 988.

To set a default backup method for the Enterprise Vault backup jobs

- 1 On the **Tools** menu, click **Options**.
- 2 In the task pane, under **Job Defaults**, click **Enterprise Vault**.
- 3 Select an appropriate backup method. Choices include Full, Differential, or Incremental.

See [“About default Enterprise Vault backup options”](#) on page 993.

- 4 Click **OK**.

About default Enterprise Vault backup options

You can select a default backup method that is based on the type of Enterprise Vault database that you want to back up.

Note: For Directory and Monitoring databases, and the Enterprise Vault 8.x Audit and FSA Reporting databases, the Full backup method is substituted for the differential backup method.

When you use the incremental backup method for the Enterprise Vault databases, the transaction logs are backed up and then truncated.

See [“About backup methods for Enterprise Vault backup jobs”](#) on page 988.

About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases

Backup Exec automatically checks the physical consistency of an Enterprise Vault database before a backup job and after a restore job. It also checks the consistency of the Compliance and Discovery databases before a backup job and after a restore job. Backup Exec uses Microsoft SQL Server's Physical Check Only utility for consistency checks of the databases. In the event a consistency check fails, Backup

Exec continues with the job and reports the consistency check failures in the Backup Exec job log.

If consistency checks fail during a restore operation, Backup Exec continues the job and reports the consistency check failures in the Backup Exec job log.

For more information about the Physical Check Only utility, see the Microsoft SQL Server documentation.

Backing up an Enterprise Vault open partition

When you back up an open partition, Backup Exec automatically backs up the partition's associated vault store database in the same backup job. Backup Exec includes vault store database to help maintain synchronization between the vault store database and the open partition if a recovery operation is required.

See [“Scheduling jobs”](#) on page 356.

See [“About restoring Enterprise Vault”](#) on page 1012.

To back up an open partition

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Backup Selections** pane, expand **Enterprise Vault**.
- 4 Expand a Directory on *<Computer_name>* that contains the partition you want to back up.
- 5 Do the following:

To back up Enterprise Vault 7.x partitions Do the following in the order listed:

- Expand an Enterprise Vault site that contains the server where the vault store partition that you want to back up resides.
- Expand the Enterprise Vault server that contains the open partition that you want to back up.

To back up Enterprise Vault 8.x partitions Do the following in the order listed:

- Expand an Enterprise Vault site that contains the vault store group where the vault store partition that you want to back up resides.
- Expand the vault store group.
- Expand the vault store that contains the open partition that you want to back up.

6 Expand a vault store that contains the partition that you want to back up.

7 Expand **All Partitions**.

8 Select **Open Partitions**.

You must select the open partition that you want to back up from the backup selections view. You cannot select the open partition from the results pane.

Backup Exec automatically includes the open partition's associated vault store database in the backup job when you select an open partition for backup.

9 In the task pane, under **Settings**, click **Enterprise Vault**.

10 Select a backup method.

See "[About backup methods for Enterprise Vault backup jobs](#)" on page 988.

11 In the task pane, select other backup options as appropriate.

12 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See "[Scheduling jobs](#)" on page 356.
- Click **Submit**.

Backing up Enterprise Vault closed partitions

Use the following steps to back up vault store closed partitions.

To back up vault store closed partitions

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Backup Selections** pane, expand **Enterprise Vault**.
- 4 Expand a Directory on <Computer name> where the partition that you want to back up resides.
- 5 Do the following:

To back up an Enterprise Vault 7.x closed partition Do the following in the order listed:

- Expand an Enterprise Vault site that contains the server where the vault store partition that you want to back up resides.
- Expand the Enterprise Vault server that contains the closed partition that you want to back up.

To back up an Enterprise Vault 8.x closed partition Do the following in the order listed:

- Expand an Enterprise Vault site that contains the vault store group where the vault store partition that you want to back up resides.
- Expand the vault store group.
- Expand the vault store that contains the closed partition that you want to back up.

- 6 Expand **All Partitions**.
- 7 Double-click **Closed Partitions**.
- 8 In the results pane, select the partitions that you want to back up.
- 9 In the task pane, under **Settings**, click **Enterprise Vault**.
- 10 Select a backup method.
See [“About backup methods for Enterprise Vault backup jobs”](#) on page 988.
- 11 In the task pane, select other backup options as appropriate.
- 12 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See "[Scheduling jobs](#)" on page 356.
- Click **Submit**.

Backing up Enterprise Vault 8.x ready partitions

Use the following steps to back up Enterprise Vault 8.x **ready** partitions.

To back up Enterprise Vault 8.x ready partitions

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Backup Selections** pane, expand **Enterprise Vault**.
- 4 Expand a Directory on <Computer name> where the **ready** partition that you want to back up resides.
- 5 Expand an Enterprise Vault site that contains the vault store group where the **ready** partition that you want to back up resides.
- 6 Expand the vault store group.
- 7 Expand the vault store that contains the **ready** partition that you want to back up.
- 8 Expand **All Partitions**.
- 9 In the results pane, select the **ready** partition.
- 10 In the task pane, under **Settings**, click **Enterprise Vault**.
- 11 Select a backup method.
See "[About backup methods for Enterprise Vault backup jobs](#)" on page 988.
- 12 In the task pane, select other backup options as appropriate.
- 13 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[About scheduling jobs](#)” on page 356.
- Click **Submit**.

Backing up the Enterprise Vault Directory database

Use the following steps to back up the Directory database.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Directory database before it backs it up.

See “[About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)” on page 993.

To back up the Directory database

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **backup selections view**, expand **Enterprise Vault**.
- 4 Expand a Directory on <Computer name> that contains the Directory database that you want to back up.
- 5 Select **Directory DB (<SQLServer/instance>/EnterpriseVaultDirectory)**.
- 6 In the task pane, under **Settings**, click **Enterprise Vault**.
- 7 Select a backup method.
See “[About backup methods for Enterprise Vault backup jobs](#)” on page 988.
- 8 In the task pane, select other backup options as appropriate.
- 9 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

Backing up the Enterprise Vault Monitoring database

Use the following steps to back up the Monitoring database.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Monitoring database before it backs it up.

See [“About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases”](#) on page 993.

To back up the Monitoring database

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Expand a Directory on *<Computer name>* that contains the Monitoring database that you want to back up.
- 5 Select **MonitoringDB(<SQLServer/instance>/EnterpriseVaultMonitoring)**.
- 6 In the task pane, under **Settings**, click **Enterprise Vault**.
- 7 Select a backup method.
See [“About backup methods for Enterprise Vault backup jobs”](#) on page 988.
- 8 In the task pane, select other backup options as appropriate.
- 9 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

Backing up an Enterprise Vault vault store database

Use the following steps to back up a vault store database.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the vault store database before it backs it up.

See [“About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases”](#) on page 993.

To back up a vault store database

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Do one of the following:

To back up the Enterprise Vault 7.x vault store database

To the following in the order listed:

- Expand a Directory on *<Computer_name>* that contains the vault store database that you want to back up.
- Expand the Enterprise Vault site that contains the vault store database that you want to back up.
- Expand the Enterprise Vault server that contains the vault store that you want to back up.
- Expand the vault store that contains the vault store database that you want to back up..

To back up the Enterprise Vault 8.x vault store database Do the following in the order listed:

- Expand a Directory on <Computer_name> that contains the vault store database that you want to back up.
- Expand the Enterprise Vault site that contains the vault store database that you want to back up.
- Expand a vault store group.
- Expand a vault store that contains the vault store database that you want to back up.

- 5 Select **Vault Store DB** (<VaultStore_SQL_Server_name/instance>/<vault_storedatabase_name>)
- 6 In the **task** pane, under **Settings**, click **Enterprise Vault**.
- 7 Select a backup method
See “[About backup methods for Enterprise Vault backup jobs](#)” on page 988.
- 8 In the task pane, select other backup options as appropriate.
- 9 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

Backing up the Enterprise Vault 8.x Audit database

Use the following steps to backup up the Audit database.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Audit database before it backs it up.

See “[About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)” on page 993.

To backup the Enterprise Vault 8.x Audit database

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Expand a Directory on <Computer_name> that contains the audit database that you want to back up.
- 5 Select **Audit DB (<SQLServer/instance>/EnterpriseVaultAudit)**.
- 6 In the **task** pane, under **Settings**, click **Enterprise Vault**.
- 7 Select a backup method.
See “[About backup methods for Enterprise Vault backup jobs](#)” on page 988.
- 8 In the task pane, select other backup options as appropriate.
- 9 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

Backing up the Enterprise Vault 8.x FSA Reporting database

Use the following steps to backup up the FSA Reporting database.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the FSA Reporting database before it backs it up.

See “[About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)” on page 993.

To back up the Enterprise Vault 8.x FSA Reporting database

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.

- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Expand a Directory on <Computer_name> that contains the FSA Reporting database that you want to back up.
- 5 Select **FSAREporting DB (<SQLServer/instance>/EnterpriseVaultFSAREporting)**.
- 6 In the **task** pane, under **Settings**, click **Enterprise Vault**.
- 7 Select a backup method.
See “[About backup methods for Enterprise Vault backup jobs](#)” on page 988.
- 8 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

Backing up the Enterprise Vault 8.x Fingerprint database

Use the following steps to back up the Fingerprint database

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Fingerprint database before it backs it up.

See “[About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)” on page 993.

To back up the Enterprise Vault 8.x Fingerprint database

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Expand a Directory on <Computer_name> that contains the Fingerprint database that you want to back up.

- 5 Expand the vault site.
- 6 Expand the vault store group.
- 7 Expand **Fingerprint Databases**.
- 8 Select a Fingerprint database.

For example, **Fingerprint DB**
(**<SQLServer/instance>/EnterpriseVaultFingerprint**)

Fingerprint database names are based on a naming convention that you determine.

- 9 In the **task** pane, under **Settings**, click **Enterprise Vault**.
- 10 Select a backup method.

See [“About backup methods for Enterprise Vault backup jobs”](#) on page 988.

- 11 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

Backing up the Enterprise Vault 8.x Compliance Accelerator Configuration database and Compliance Accelerator customer databases

Use the following steps to back up the Compliance Accelerator Configuration database. You can also use these steps to back up the Compliance Accelerator customer databases.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Configuration database before it backs it up.

See [“About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases”](#) on page 993.

To back up the Enterprise Vault 8.x Compliance Accelerator Configuration database

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Expand **Accelerators**.
- 5 Expand **Compliance on <server_name>**.
- 6 Select **Configuration DB (<SQLServer/instance>/EVConfiguration)**.
- 7 Select Compliance Accelerator customer databases, if desired.

For example, **mycompanyABC_cpml Customer DB (<SQLServer/instance>/mycompanyABC_cpml)**

- 8 In the **task** pane, under **Settings**, click **Enterprise Vault**.
- 9 Select a backup method.

See “[About backup methods for Enterprise Vault backup jobs](#)” on page 988.

- 10 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

Backing up the Enterprise Vault 8.x Discovery Accelerator Configuration database and Discovery Accelerator customer databases

Use the following steps to back up the Discovery Accelerator Configuration database. You can also use these steps to back up the Discovery Accelerator customer databases.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Configuration database before it backs it up.

See “[About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)” on page 993.

To back up the Enterprise Vault 8.x Discovery Accelerator Configuration database and Discovery Accelerator customer databases

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Expand **Accelerators**.
- 5 Expand **Discovery <server_name>**.
- 6 Select **Discovery DB (<SQLServer/instance>/EVDISCOVERY)**.
- 7 Select Discovery Accelerator customer databases, if desired.
- 8 In the **task** pane, under **Settings**, click **Enterprise Vault**.
- 9 Select a backup method.

See “[About backup methods for Enterprise Vault backup jobs](#)” on page 988.

- 10 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

Backing up the Discovery Accelerator Custodian database

Use the following steps to back up the Discovery Accelerator Custodian database.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Custodian database before it backs it up.

See [“About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases”](#) on page 993.

See [“SQL backup options”](#) on page 1273.

To back up the Discovery Accelerator Custodian database

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Expand **Accelerators**.
- 5 Expand **Discovery <server_name>**.
- 6 Select **<database_name> Custodian DB (<SQLServer/instance>/<database_name>)**.
- 7 In the **task** pane, under **Settings**, click **Enterprise Vault**.
- 8 Select a backup method.

See [“About backup methods for Enterprise Vault backup jobs”](#) on page 988.

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

Backing up an Enterprise Vault vault store

When you back up a vault store, all closed partitions, open partitions, the vault store database, and Ready partitions are backed up.

See [“Backing up an Enterprise Vault open partition”](#) on page 994.

See [“Backing up Enterprise Vault closed partitions”](#) on page 995.

To back up a vault store

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Do one of the following:

To back up an Enterprise Vault 7.x vault store

To the following in the order listed:

- Expand a Directory on *<Computer_name>* that contains the vault store that you want to back up.
- Expand the Enterprise Vault site that contains the vault store that you want to back up.
- Expand the Enterprise Vault server that contains the vault store that you want to back up.

To back up an Enterprise Vault 8.x vault store

Do the following in the order listed:

- Expand a Directory on *<Computer_name>* that contains the vault store that you want to back up.
- Expand the Enterprise Vault site that contains the vault store that you want to back up.
- Expand a vault store group.

- 5 Select a vault store.
- 6 In the task pane, under **Settings**, click **Enterprise Vault**.
- 7 Select a backup method.
See "[About backup methods for Enterprise Vault backup jobs](#)" on page 988.
- 8 In the task pane, select other backup options as appropriate.
- 9 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

About backing up an Enterprise Vault 7.x server and an Enterprise 8.x site

When you back up an Enterprise Vault 7.x server, the following items are backed up:

- Index locations
- Vault stores
- Vault store closed partitions, if present
- Vault store open partitions
- Vault store databases

When you back up an Enterprise Vault 8.x site, all of the same items listed above are included. In addition, the following Enterprise Vault 8.x components are also backed up:

- Audit, Fingerprint, and FSA Reporting databases
- Vault store groups
- Vault store ready partitions, if present

Backup Exec also does an automatic backup of the Directory database when you back up a Enterprise Vault 7.x server or an Enterprise Vault 8.x site.

See [“Backing up an Enterprise Vault 7.x server”](#) on page 1009.

See [“Backing up an Enterprise Vault site”](#) on page 1010.

Backing up an Enterprise Vault 7.x server

Use the following steps to back up an Enterprise Vault 7.x server.

See [“About backing up an Enterprise Vault 7.x server and an Enterprise 8.x site”](#) on page 1009.

To back up an Enterprise Vault 7.x server

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Expand an Enterprise Vault directory where the server that you want to back up resides.
- 5 Expand an Enterprise Vault site.
- 6 Select an Enterprise Vault server.
- 7 In the task pane, under **Settings**, click **Enterprise Vault**.
- 8 Select a backup method.
See “[About backup methods for Enterprise Vault backup jobs](#)” on page 988.
- 9 In the task pane, select other backup options as appropriate.
- 10 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the task pane, under Frequency, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

Backing up an Enterprise Vault site

When you back up an Enterprise Vault site, Backup Exec also does an automatic backup of the Directory database.

To back up an Enterprise Vault site

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Expand an Enterprise Vault directory that contains the site that you want to back up.
- 5 Select the Enterprise Vault site.
- 6 In the task pane, under **Settings**, click **Enterprise Vault**.

- 7 Select a backup method.
See “[About backup methods for Enterprise Vault backup jobs](#)” on page 988.
- 8 In the task pane, select other backup options as appropriate.
- 9 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

Backing up Enterprise Vault index locations

Use the following steps to back up Enterprise Vault index locations.

To back up index locations

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the backup selections view, expand **Enterprise Vault**.
- 4 Expand a Directory on *<Computer name>* that contains the index locations that you want to back up.
- 5 Do the following:

To back up an Enterprise Vault 7.x index location To the following in the order listed:

- Expand a Directory on *<Computer_name>* that contains the index location that you want to back up.
- Expand the Enterprise Vault site that contains the index location that you want to back up.
- Expand the Enterprise Vault server that contains the index location that you want to back up.

To back up an Enterprise Vault 8.x index location Do the following in the order listed:

- Expand a Directory on <Computer_name> that contains the index location that you want to back up.
- Expand the Enterprise Vault site that contains the index location that you want to back up.

6 Do one of the following:

To back up all index locations Check **Index Locations**.

To back up individual index locations Do the following in the order listed:

- Click the **Index Locations** icon.
- In the results pane, select the individual index locations that you want to back up.

7 In the task pane, under **Settings**, click **Enterprise Vault**.

8 Select a backup method.

See “[About backup methods for Enterprise Vault backup jobs](#)” on page 988.

9 In the task pane, select other backup options as appropriate.

10 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the task pane, under Frequency, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

About restoring Enterprise Vault

Review the following before you begin an Enterprise Vault restore operation.

- When you restore an Enterprise Vault installation, you should restore the Directory database in a separate restore job. After you successfully restore the

Directory database, you can restore other Enterprise Vault components and partitions.

See [“Restoring the Enterprise Vault Directory database”](#) on page 1015.

- When you restore Enterprise Vault databases, you can select the options that either leave databases in a ready-to-use state or in a non-operational state. The non-operational state options that you select apply to all Enterprise Vault databases except the vault store database. When you restore an Enterprise Vault 8.x vault store database, the Enterprise Vault Agent places the vault store database in Enterprise Vault 8.x Backup mode. If the vault store database remains in a non-operational state after the restore job completes, the Enterprise Vault Agent cannot remove it from Backup mode.

If you select the option that leaves the databases ready to use, the following applies:

- The Enterprise Vault Agent restores the vault store database in a ready-to-use, operational state. The vault store database's operational status is maintained even when you select additional backup sets for restore in the same vault store database restore job. Additional backup sets can include Full, Differential, and Incremental backup methods.

When you choose the option that leaves the databases in a nonoperational state, the following applies:

- The Enterprise Vault Agent prompts you to stop the **Enterprise Vault Storage Service** before you start the vault store database restore operation. You can restart the vault store restore operation again after the Enterprise Vault Storage Service stops.

As a best practice, Symantec recommends that you restore the vault store database in a ready-to-use state. When you restore the vault store database in a nonoperational state, Enterprise Vault cannot remove it from Backup mode after the restore operation finishes.

See [“Enterprise Vault restore options”](#) on page 1034.

- You can individually restore Enterprise Vault components. Before you begin the restore, the databases and other components may or may not exist on the destination Enterprise Vault server. If the databases do not exist, you can restore them using the Enterprise Vault Agent. After the restore job completes, you must configure Enterprise Vault to use the restored databases. To configure Enterprise Vault to use the restored databases, see your Enterprise Vault documentation.

These items include the:

- Enterprise Vault 7.x and 8.x Directory, Monitoring, Audit, FSAReporting, and Fingerprint databases.

- Vault store databases, indexes, and partitions.
- Compliance and Discovery Accelerator Configuration and Customer databases.
- Discovery Accelerator Custodian database
- Symantec recommends that you use the Enterprise Vault service account or an account with rights to access the restore selections as the default logon account. Otherwise, you may have to enter proper credentials for each Enterprise Vault resource that you select for restore.
- After you restore Enterprise Vault, a message appears that says you need to run Enterprise Vault recovery tools. The recovery tools are used to re-synchronize Enterprise Vault with the newly restored databases after you complete the restore.
For information on running the Enterprise Vault recovery tools, see your Enterprise Vault documentation.

Before you restore Enterprise Vault sites, servers or other components, you should have the following items installed on the destination computer:

- Enterprise Vault
- The Backup Exec Remote Agent for Windows Systems

Note: You must install the Remote Agent on remote Enterprise Vault computers where you want to restore Enterprise Vault components.

See [“Restoring the Enterprise Vault Directory database”](#) on page 1015.

See [“Restoring Enterprise Vault partitions”](#) on page 1017.

See [“Restoring an Enterprise Vault 7.x server to its original location”](#) on page 1033.

See [“Redirecting the restore of Enterprise Vault databases to a different Microsoft SQL Server computer”](#) on page 1039.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

About automatic redirection of Enterprise Vault components under an Enterprise Vault server

You can change the location of the vault store databases, Enterprise Vault 8.x Fingerprint databases, or partitions to a location that differs from where they were backed up. During restores of the vault store database, Enterprise Vault 8.x Fingerprint databases, or partitions, the Enterprise Vault Agent detects the location change. It then automatically redirects the component restores to the new location.

Note: Automatic redirected restores of the vault databases, partitions, or Enterprise Vault 8.x Fingerprint databases occur when you change only the location of these Enterprise Vault components. The names of the partitions, vault stores, and vault store groups must not change from the time the partition was originally backed up.

See [“Restoring the Enterprise Vault Directory database”](#) on page 1015.

See [“Restoring Enterprise Vault partitions”](#) on page 1017.

See [“Restoring an Enterprise Vault 7.x server to its original location”](#) on page 1033.

See [“Redirecting the restore of Enterprise Vault databases to a different Microsoft SQL Server computer”](#) on page 1039.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

Restoring the Enterprise Vault Directory database

Use the following steps to restore the Enterprise Vault Directory database. You also can redirect the restore of the Directory database to a different Microsoft SQL Server computer.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Directory database after the database is restored.

See [“About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases”](#) on page 993.

To restore the Enterprise Vault Directory database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** pane, expand an Enterprise Vault installation that contains the Directory database that you want to restore.
For example, expand Directory on *<Computer name>*.
- 4 Expand **Directory DB (<SQL Server name>/<instance>/EnterpriseVaultDirectory)**.
- 5 Select the backup set that you want to restore.
- 6 In the task pane, under **Settings**, click **Enterprise Vault**.

- 7 Click **Terminate the database connections automatically when restoring selected databases. (Do not terminate database connections for Vault Store database.)**

If you do not select this option, you must stop the Enterprise Vault Admin and Directory services on the Enterprise Vault computer where you want to restore the Directory database. If other Enterprise Vault servers connect to the Directory database, stop the Admin and Directory services on those computers too.

- 8 Select other restore options that you want to use.
See [“Enterprise Vault restore options”](#) on page 1034.
- 9 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

- 10 After the restore successfully completes, restart all Directory and Admin services.

Restoring the Enterprise Vault Monitoring database

Use the following steps to restore the Monitoring database to its original location.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Monitoring database after the database is restored.

See [“About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases”](#) on page 993.

To restore the Monitoring database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** pane, expand **All Resources**.

- 4 Expand an Enterprise Vault installation that contains the Directory database that you want to restore.

For example, expand Directory on *<Computer name>*.

- 5 Expand **Monitoring DB (<SQL Server name>/<instance>/EnterpriseVaultMonitoring)**.
- 6 Select the backup set that you want to restore.
- 7 In the task pane, under **Settings**, click **Enterprise Vault**.
- 8 Click **Terminate the database connections automatically when restoring selected databases. (Do not terminate database connections for Vault Store database.)**

If you do not select this option, you must stop the Enterprise Vault Admin and Directory services on the Enterprise Vault computer where you want to restore the Monitoring database. If other Enterprise Vault servers connect to the Monitoring database, stop the Admin and Directory services on those computers too.

- 9 Select other restore options that you want to use.
See [“Enterprise Vault restore options”](#) on page 1034.
- 10 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

- 11 After the restore successfully completes, restart all Directory and Admin services.

Restoring Enterprise Vault partitions

When you restore an open partition, Backup Exec automatically restores the partition’s associated vault store database in the same restore job. By including the vault store database, Backup Exec maintains synchronization between the two components.

Note: Restoring an open partition means that the partition that you select for restore is currently open on the destination Enterprise Vault server.

During an open partition restore job, the Enterprise Vault Agent restores the vault database that was backed up at the time the partition was backed up. If the vault database backup does not exist, an existing backup of the vault database is restored instead. The Enterprise Vault Agent selects a backup of the vault database that was backed up closest to the time of the partition backup.

For example, if you restore an open partition that was backed up at 10:00 A.M., the Enterprise Vault Agent restores the 10:00 A.M. backup of the vault database. If you do not have 10:00 A.M. backup of the vault database, but you have a 9:45 A.M. backup, the Enterprise Vault Agent automatically restores the 9:45 A.M. backup.

After you restore open, closed, or Enterprise Vault 8.x **Ready** partitions, you must run the Enterprise Vault recovery tool. The recovery tool maintains synchronization between the vault store database and its associated partitions.

See your Enterprise Vault documentation.

See [“Restoring an Enterprise Vault 7.x server to its original location”](#) on page 1033.

See [“Redirecting the restore of Enterprise Vault databases to a different Microsoft SQL Server computer”](#) on page 1039.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

To restore an Enterprise Vault partition

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **Restore Job Properties** pane, expand **All Resources**.
- 4 Expand a Directory on <Computer_name> that contains the partition that you want to restore.
- 5 Do the following:

To restore an Enterprise Vault 7.x partition

Do the following in the order listed:

- Expand an Enterprise Vault site that contains the server where the partition that you want to restore resides.
- Expand the Enterprise Vault server that contains the partition that you want to restore.

To restore an Enterprise Vault 8.x partition

Do the following in the order listed:

- Expand an Enterprise Vault site that contains the partition that you want to restore.
- Expand the vault store group that contains the partition that you want to restore.

6 Expand **Partitions**.

7 Expand a partition that contains the partition that you want to restore.

8 Select the backup set that you want to restore.

9 In the task pane, under **Settings**, select **Enterprise Vault**.

10 Select the restore options you want to use.

See “[Enterprise Vault restore options](#)” on page 1034.

11 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

12 After the restore successfully completes, run the Enterprise Vault recovery tool.

Restoring an Enterprise Vault vault store database

Use the following steps to restore a vault store database.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the vault store database after the database is restored.

See “[About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)” on page 993.

To restore a vault store database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** pane, expand a Directory on *<Computer name>* that contains the vault store database that you want to restore.
- 4 Do the following:

To restore an Enterprise Vault 7.x vault store database

Do the following in the order listed:

- Expand an Enterprise Vault site that contains the vault store database that you want to restore.
- Expand an Enterprise Vault server that contains the vault store database that you want to restore.

To back up an Enterprise Vault 8.x vault store database

Do the following in the order listed:

- Expand an Enterprise Vault site that contains the vault store database that you want to restore.
- Expand the vault store group that contains the vault store database that you want to restore.

- 5 Expand the vault store.
- 6 Expand **Vault Store DB** (**<SQL_Server_name>/<instance>/EV<vault_store_database_name>**).
- 7 Select the backup set that you want to restore.
- 8 In the task pane, under **Settings**, select **Enterprise Vault**.
- 9 Select the restore options you want to use.
See [“Enterprise Vault restore options”](#) on page 1034.
- 10 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

Restoring an Enterprise Vault 8.x Audit database

Use the following steps to restore an Audit database to its original location. You can also redirect the restore location of the Audit database.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Audit database after the database is restored.

See [“About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases”](#) on page 993.

To restore an Enterprise Vault 8.x Audit database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** pane, expand an Enterprise Vault 8.x Directory on *<Computer name>* that contains the audit database that you want to restore.
- 4 Expand **AuditDB(<SQL_Server_name>/<instance>/EnterpriseVaultAudit)**.
Audit database names are based on naming conventions that you determine.
- 5 Select the backup set that you want to restore.
- 6 In the task pane, under **Settings**, select **Enterprise Vault**.
- 7 Select the restore options you want to use.
- 8 See [“Enterprise Vault restore options”](#) on page 1034.
- 9 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

Restoring the Enterprise Vault 8.x FSA Reporting database

Use the following steps to restore an FSAREporting database to its original location. You can also redirect the restore location of the FSA Reporting database..

See “[Redirecting an Enterprise Vault restore job](#)” on page 1036.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the FSA Reporting database after the database is restored.

See “[About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)” on page 993.

To restore the Enterprise Vault 8.x FSA Reporting database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** pane, expand an Enterprise Vault 8.x Directory on *<Computer name>* that contains the FSAREporting database that you want to restore.
- 4 Expand **<FSAREporting_database_name> DB (<SQL_Server_name>/<instance>/EnterpriseVaultFSAREporting)**
FSAREporting database names are based on naming conventions that you determine.
- 5 Select the backup set that you want to restore.
- 6 In the task pane, under **Settings**, select **Enterprise Vault**.
- 7 Check **Terminate the database connections automatically when restoring selected databases (Do not terminate database connections for Vault Store database).**
- 8 Select the other restore options you want to use.

9 See [“Enterprise Vault restore options”](#) on page 1034.

10 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

Restoring the Enterprise Vault 8.x Fingerprint database

Use the following steps to restore an Fingerprint database to its original location. You can also redirect the restore location of the Fingerprint database.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Fingerprint database after the database is restored.

See [“About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases”](#) on page 993.

To restore the Enterprise Vault 8.x Fingerprint database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** pane, expand an Enterprise Vault 8.x Directory on *<Computer name>* that contains the Fingerprint database that you want to restore.
- 4 Expand an Enterprise Vault site that contains the Fingerprint database that you want to restore.
- 5 Expand a vault store group that contains the Fingerprint database that you want to restore.
- 6 Expand **Fingerprint Databases**.
- 7 Expand Fingerprint DB
(*<SQL_Server_name>/<instance>/<SQL_Server_name/vault_store_group_name>*).
- 8 Select the backup set that you want to restore.

- 9 In the **task** pane, under **Settings**, select **Enterprise Vault**.
- 10 Check **Terminate the database connections automatically when restoring selected databases. (Do not terminate database connections for the Vault Store database)**.
- 11 Select other restore options as appropriate.
See [“Enterprise Vault restore options”](#) on page 1034.
- 12 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

Restoring the Compliance Accelerator Configuration database

Use the following steps to restore an Compliance Accelerator Configuration database to its original location. You can also redirect the restore location of the Configuration database.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Compliance Accelerator Configuration database after the database is restored.

See [“About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases”](#) on page 993.

To restore the Compliance Accelerator Configuration database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 In the **View by Resource** pane, expand **Accelerators**.
- 3 Expand **Compliance on <computer_name>**.
- 4 Expand **Configuration DB <SQL_Server_name>/<instance>/EVConfiguration>**.
- 5 Select the backup set that you want to restore.

- 6 In the task pane, under **Settings**, select **Enterprise Vault**.
- 7 Check **Terminate the database connections automatically when restoring selected databases. (Do not terminate database connections for the Vault Store database).**

If you do not use this option, you must stop the Accelerator Manager service on the computer where you restore the Compliance Accelerator Configuration database.

- 8 Select other restore options that you want to use.
See [“Enterprise Vault restore options”](#) on page 1034.
- 9 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

- 10 After Backup Exec successfully restores the database, restart the Accelerator Manager service on the Compliance Accelerator server.

Restoring the Compliance Accelerator Customer database

Use the following steps to restore one or more Compliance Accelerator Customer databases to its original location. You can also redirect the restore location of the Customer database.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

To restore the Compliance Accelerator Customer database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 In the **View by Resource** pane, expand **Accelerators**.
- 3 Expand **Compliance on <computer_name>**.
- 4 Expand **<database_name> Customer DB <SQL_Server_name>/<instance>/CA/<database_name>**.
- 5 Select the backup set that you want to restore.
- 6 To restore multiple Customer databases, repeat steps 6 and 7.

- 7 In the task pane, under **Settings**, select **Enterprise Vault**.
- 8 Check **Terminate the database connections automatically when restoring selected databases. (Do not terminate database connections for the Vault Store database)**.

If you do not use this option, you must stop the Accelerator Manager service on the computer where you restore the Compliance Accelerator Customer database.

- 9 Select other restore options that you want to use.
See “[Enterprise Vault restore options](#)” on page 1034.
- 10 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

- 11 After Backup Exec successfully restores the database or databases, restart the Accelerator Manager service on the Compliance Accelerator server.

Restoring the Discovery Accelerator Configuration database

Use the following steps to restore the Discovery Accelerator Configuration database to its original location. You can also redirect the restore location of the Configuration database.

See “[Redirecting an Enterprise Vault restore job](#)” on page 1036.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Discovery Accelerator Configuration database after the database is restored.

See “[About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)” on page 993.

To restore the Discovery Accelerator Configuration database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 In the **View by Resource** pane, expand **Accelerators**.

- 3 Expand **Discovery on <computer_name>**.
- 4 Expand **Configuration DB <SQL_Server_name>/<instance>/DA**.
- 5 Select the backup set that you want to restore.
- 6 In the task pane, under **Settings**, select **Enterprise Vault**.
- 7 Check **Terminate the database connections automatically when restoring selected databases. (Do not terminate database connections for the Vault Store database)**.

If you do not use this option, you must stop the Accelerator Manager service on the computer where you restore the Discovery Accelerator Configuration database.

- 8 Select other restore options that you want to use.
See [“Enterprise Vault restore options”](#) on page 1034.
- 9 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

- 10 After Backup Exec successfully restores the database or databases, restart the Accelerator Manager service on the Discovery Accelerator server.

Restoring the Discovery Accelerator Custodian database

Use the following steps to restore the Discovery Accelerator Custodian database to the original location. You can also redirect the restore location of the Custodian database.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

Note: The Enterprise Vault Agent automatically runs a physical consistency check of the Discovery Accelerator Custodian database after the database is restored.

See [“About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases”](#) on page 993.

To restore the Discovery Accelerator Custodian database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 In the **View by Resource** pane, expand **Accelerators**.
- 3 Expand **Discovery on <computer_name>**.
- 4 Expand **<database_name> Custodian DB <SQL_Server_name>/<instance>/<database_name>**.
- 5 Select the backup set that you want to restore.
- 6 Check **Terminate the database connections automatically when restoring selected databases. (Do not terminate database connections for the Vault Store database)**.

If you do not use this option, you must stop the Accelerator Manager service on the computer where you restore the Discovery Accelerator Custodian database.

- 7 Select other restore options that you want to use.
See “[Enterprise Vault restore options](#)” on page 1034.
- 8 Do the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

- 9 After Backup Exec successfully restores the database or databases, restart the Accelerator Manager service on the Discovery Accelerator server.

Restoring the Discovery Accelerator Customer database

Use the following steps to restore one or more Discovery Accelerator Customer databases to the original location. You can also redirect the restore location of the Customer database.

See “[Redirecting an Enterprise Vault restore job](#)” on page 1036.

To restore the Discovery Accelerator Customer database

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 In the **View by Resource** pane, expand **Accelerators**.
- 3 Expand **Discovery on <computer_name>**.
- 4 Expand **<database_name> Customer DB <SQL_Server_name>/<instance>/<database_name>**.
- 5 Select the backup set that you want to restore.
- 6 Check **Terminate the database connections automatically when restoring selected databases. (Do not terminate database connections for the Vault Store database)**.

If you do not use this option, you must stop the Accelerator Manager service on the computer where you restore the Discovery Accelerator Customer database.

- 7 Select other restore options as appropriate.
See “[Enterprise Vault restore options](#)” on page 1034.
- 8 Do the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

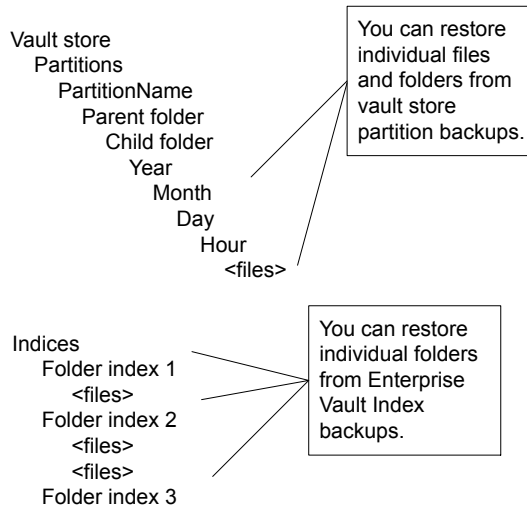
- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

- 9 After Backup Exec successfully restores the database or databases, restart the Accelerator Manager service on the Discovery Accelerator server.

About restoring individual files and folders with the Enterprise Vault Agent

The Enterprise Vault Agent supports individual file and folder restores from vault store partition backups. You can also restore complete index locations or individual folders from Enterprise Vault index backups.

Figure E-1 Restoring individual files from vault store partitions and complete folders from an Enterprise Vault index



See [“Restoring individual files from partitions by using the Enterprise Vault Agent”](#) on page 1030.

See [“Restoring individual folders from an Enterprise Vault index backup”](#) on page 1032.

Restoring individual files from partitions by using the Enterprise Vault Agent

Use the following steps to restore individual files from open and closed partitions.

See [“About restoring individual files and folders with the Enterprise Vault Agent”](#) on page 1029.

To restore individual files from partitions by using the Enterprise Vault Agent

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **Restore Job Properties** pane, expand **All Resources**.
- 4 Expand a directory on *<Computer_name>* that contains the partition that you want to restore.
- 5 Do the following:

To restore individual files from an Enterprise Vault 7.x partition

Do the following in the order listed:

- Expand an Enterprise Vault site that contains the server where the partition information that you want to restore resides.
- Expand the Enterprise Vault server where the partition information that you want to restore resides.

To restore individual files from an Enterprise Vault 8.x partition

Do the following in the order listed:

- Expand an Enterprise Vault site that contains the vault store where the partition information that you want to restore resides.
- Expand the vault store group that contains the vault store where the partition information that you want to restore resides.

- 6 Expand a vault store that contains the partitions and the files that you want to restore.
- 7 Expand **Partitions**.
- 8 Expand a partition that contains the files that you want to restore.
- 9 Select the backup set that you want to restore.
- 10 In the Results pane, select one or more files and/or folders that you want to restore.
- 11 In the **Restore Job Properties** pane, select other restore options, if appropriate.
- 12 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

- 13 After the restore successfully completes, run the Enterprise Vault recovery tool. See your Enterprise Vault documentation for more information.

Restoring individual folders from an Enterprise Vault index backup

Use the following steps to restore complete folders from an Enterprise Vault index.

See “[About restoring individual files and folders with the Enterprise Vault Agent](#)” on page 1029.

To restore folders from an Enterprise Vault index backup

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** pane, expand **All Resources**.
- 4 Expand **Enterprise Vault**.
- 5 Expand a directory on *<Computer_Name>* that contains the index locations that you want to restore.
- 6 Do the following:

To restore Enterprise Vault 7.x folders from an index backup

Do the following in the order listed:

- Expand an Enterprise Vault site that contains the index folders that you want to restore.
- Expand the Enterprise Vault server that contains the index folders that you want to restore.

To restore Enterprise Vault 8.x folders from an index backup

Do the following:

- Expand an Enterprise Vault site that contains the index folders that you want to restore.

- 7 Expand **Index Locations**.

- 8 Expand the path that contains the folders that you want to restore.
- 9 Expand the backup set that contains the folder that you want to restore.
- 10 Select the index folder to restore.
- 11 In the **Restore Job Properties** pane, select other restore options as needed.
See [“Restoring data by setting job properties”](#) on page 605.
- 12 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

Restoring an Enterprise Vault 7.x server to its original location

Use the following steps to restore an Enterprise Vault server to its original location. You also can redirect the restore of the server to a different computer.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

See [“Redirecting the restore of Enterprise Vault databases to a different Microsoft SQL Server computer”](#) on page 1039.

To restore an Enterprise Vault server to its original location

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the Restore Job Properties pane, expand **All Resources**.
- 4 Expand a Directory on <Computer_name> that contains the server that you want to restore
- 5 Expand an Enterprise Vault site that contains the server that you want to restore.
- 6 Expand an Enterprise Vault server.
- 7 Expand the vault store.
- 8 Expand **Partitions**.
- 9 Expand each partition.

10 Select the backup sets for each partition.

11 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

12 After the restore successfully completes, restart all Directory and Admin services.

See your Enterprise Vault documentation.

Enterprise Vault restore options

Use the following table to select the restore option you want to use when you restore the Enterprise Vault databases.

Table E-5 Enterprise Vault restore options

Item	Description
Terminate the database connections automatically when restoring selected databases. (Do not terminate database connections for Vault Store database.)	<p>Takes the shared Enterprise Vault Directory, Monitoring, Auditing, FSA Reporting, and Fingerprint databases offline so Backup Exec can replace them during a restore job.</p> <p>Note: If you don't use this option, you must stop the Directory and Admin services on all Enterprise Vault servers before you restore the previously mentioned databases. In addition, you must also stop the Accelerator Manager server on all of the Compliance Accelerator servers and the Discovery servers. Only after you stop the Accelerator Manager can you restore the Customer, Configuration, and Custodian databases.</p> <p>This option causes the Enterprise Vault Admin and Directory services on all related Enterprise Vault servers to terminate the connection to the Directory database that you restore.</p> <p>It also terminates connections to the following:</p> <ul style="list-style-type: none">■ Monitoring database■ Audit, Fingerprint, and FSA Reporting databases (Enterprise Vault 8.x only)■ Configuration, Customer, and Custodian databases <p>When the restore job completes, you must manually restart the Enterprise Vault Admin and Directory services on your Enterprise Vault server. After you restart the services, the services reconnect to the restored databases and Enterprise Vault begins archival operations again.</p> <p>Note: This option causes the Enterprise Vault Admin and Directory services on all Enterprise Vault servers to terminate their connections to the Directory database that you restore. It also terminates the connections to the Enterprise Vault Accelerator Manager database.</p>

Table E-5 Enterprise Vault restore options (*continued*)

Item	Description
Leave the database ready to use; additional transaction logs or differential backups cannot be restored	<p>Rolls back all uncompleted transactions when you restore the last database, differential, or log backup. After the recovery operation, the database is ready for use. If you do not select this option, the database is left in an intermediate state and is not usable.</p> <p>If you select this option, you cannot continue to restore backups. You must restart the restore operation from the beginning.</p>
Leave the database nonoperational; additional transaction logs or differential backups can be restored	<p>Creates and maintains a standby database.</p> <p>By using this option, you can continue restoring other backups sets for non-operational databases.</p> <p>See your SQL documentation for information on standby databases.</p>

Note: Symantec recommends that you select all required backup sets when you run a single restore job for a vault store database. All required backup sets can include Full, Differential, and Incremental backup sets. The vault store database should also be restored in a ready-to-use state after the restore job completes.

See [“Restoring the Enterprise Vault Directory database”](#) on page 1015.

See [“Restoring the Enterprise Vault Monitoring database”](#) on page 1016.

See [“Restoring Enterprise Vault partitions”](#) on page 1017.

Redirecting an Enterprise Vault restore job

Use the following steps to redirect an Enterprise Vault restore job.

You also can restore the Directory database to a different Microsoft SQL Server computer.

See [“Redirecting the restore of Enterprise Vault databases to a different Microsoft SQL Server computer”](#) on page 1039.

To redirect an Enterprise Vault restore job

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **Restore Job Properties** pane, expand **All Resources**.
- 4 Navigate to and select the Enterprise Vault components that you want to redirect during restore.
- 5 In the **Restore Job Properties** pane, under **Destination**, click **Enterprise Vault Redirection**.
- 6 Select the type of redirected restore that you want to do.
See [“Redirection options for Enterprise Vault”](#) on page 1037.
- 7 Do one of the following:
Use the default logon account as indicated.
Click **Change** to select a different one.
- 8 In the **Restore Job Properties** pane, select other restore options as needed.
See [“Restoring data by setting job properties”](#) on page 605.
- 9 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

Redirection options for Enterprise Vault

You can redirect a restore job for Enterprise Vault components.

See [“Redirecting an Enterprise Vault restore job”](#) on page 1036.

Table E-6 Redirection options for Enterprise Vault

Item	Description
Redirect Enterprise Vault to server (Enterprise Vault 7.x only)	Redirects the restore of Enterprise Vault 7.x backup to a different server.

Table E-6 Redirection options for Enterprise Vault (*continued*)

Item	Description
Redirect Enterprise Vault to server	Indicates the destination server where you want to redirect the Enterprise Vault 7.x restore job.
Redirect to a new Microsoft SQL server	Redirects the restore jobs of Enterprise Vault databases and Accelerator databases to a different SQL Server. Note: Vault store databases are restored for Enterprise Vault 8.0 only.
Server	Displays the name of the server to which you want to redirect the restore job for a vault store.
Instance	Displays the name of the instance of the SQL Server to which you want to redirect the restore job for a vault store.
Restore index root to a new location	Redirects the restore job for the index root to a new location. If you redirect the restore of the Enterprise Vault server, you can specify an alternate path on the destination server. You can also redirect the index root location to an alternate path on the original server.
Path	Displays the path name to which you want to redirect the restore job for an index root.
Restore partition to a new location	Redirects the restore job for a vault store partition to a new location. Partitions are restored for Enterprise Vault 8.0 only.
Path	Displays the path name to which you want to redirect the restore job for a vault store partition.
Enterprise Vault logon account	Specifies the logon account to use.

Redirecting the restore of Enterprise Vault databases to a different Microsoft SQL Server computer

Use the following steps to redirect the restore of the Enterprise Vault databases to a different Microsoft SQL Server computer.

See [“About restoring Enterprise Vault”](#) on page 1012.

To redirect the restore of the Directory database to a different Microsoft SQL server computer

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Navigate to and select the Enterprise Vault Directory database that you want to restore.
- 4 In the task pane, under **Destination**, click **Enterprise Vault Redirection**.
- 5 Check **Redirect to a new Microsoft SQL server**.
- 6 In the server field, type the name of the SQL Server name to which you want to restore.

Use the following format: \\servername.

- 7 Check **Instance** to redirect the restore to a named SQL instance and then type the instance name. If you want to restore to the default instance, leave the field empty.
- 8 In the **Restore Job Properties** pane, select other restore options as needed.
See [“Restoring data by setting job properties”](#) on page 605.
- 9 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See [“Scheduling jobs”](#) on page 356.
- Click **Submit**.

- 10 After the restore job completes, configure Enterprise Vault to use the new name of the SQL database server.

See [“Configuring Enterprise Vault to use the name of the new SQL Server that holds the Directory database”](#) on page 1040.

Configuring Enterprise Vault to use the name of the new SQL Server that holds the Directory database

Use the following steps to configure Enterprise Vault to use the name of the new SQL Server that holds the Directory database.

See [“Redirecting the restore of Enterprise Vault databases to a different Microsoft SQL Server computer”](#) on page 1039.

To configure Enterprise Vault to use the name of the new SQL Server that holds the Directory database

- 1 At each Enterprise Vault server, use Enterprise Vault to change the name of the previous SQL Server computer. Change the name to the name of the SQL Server computer that now holds the Directory database.

See your Enterprise Vault documentation.

- 2 Restart the Enterprise Vault Admin service on all Enterprise Vault servers that use the Directory database.

Two directory names appear in the backup selections view after you restart the Enterprise Vault Admin service on the Enterprise Vault server.

For example, **Directory on <OldSQL_computer_name>** and **Directory on <NewSQL_computer_name>**.

- 3 In Backup Exec, on the navigation bar, click the arrow next to Backup.
- 4 Click **New Backup Job**.
- 5 Expand **Enterprise Vault**.
- 6 Expand **Directory on <SQL server computer where you moved the Directory database>**.
- 7 Expand all items under **Directory on <SQL server computer where you moved the Directory database>**.

The **Directory** and **Monitoring** databases, Enterprise Vault 8.x **FSA Reporting** and **Audit** databases, and the Enterprise Vault sites should appear. In addition, the Directory database should display the new SQL Server name and instance where it was redirected.

When you configure a new Directory database backup job, you must select the Directory database from the current Directory server. Backup Exec automatically removes the previous Directory server name 13 days after you complete the Directory database move.

- 8 To manually remove the previous server name, right-click **Directory on <OldSQL_computer_name>**.
- 9 Click **Delete**.

Best practices for the Enterprise Vault Agent

Symantec recommends the following best practices when you use the Enterprise Vault Agent.

- Back up the Enterprise Vault Directory database after you make any configuration changes in Enterprise Vault.
- Restore the Enterprise Vault Directory database in a separate Backup Exec restore job.
- Restore all Full, Differential, and Incremental backup sets of the vault store database in a single restore job.
- Do not allow the backup window and archive window to overlap.
- Do not allow the backup window and the migration window to overlap.
- Make sure Enterprise Vault 8.x components are not in Backup mode before you back up the Enterprise Vault 8.x Directory database.
- If you install both the Symantec Backup Exec NDMP Option and the Enterprise Vault Agent, pick only one product to protect an Enterprise Vault partition that resides on NDMP filers.
- Do not change the recovery model of any database that is created by Enterprise Vault. Enterprise Vault configures each database in full recovery mode when it creates them.

About the Backup Exec Migrator for Enterprise Vault

The Backup Exec Migrator for Enterprise Vault (Backup Exec Migrator) lets you automatically migrate archived Enterprise Vault data to the storage devices that Backup Exec manages. By migrating the archived Enterprise Vault data from a partition, you can reclaim disk space on the Enterprise Vault server without incurring the cost of additional hardware.

By migrating Enterprise Vault archive data to the Backup Exec media server storage devices, you also ensure an added level of storage redundancy using an off-host environment.

See [“How the Backup Exec Migrator works”](#) on page 1042.

See [“Configuring the Backup Exec Migrator”](#) on page 1049.

Backup Exec Migrator for Enterprise Vault requirements

Before you configure the Backup Exec Migrator, ensure that your Enterprise Vault server meets the following requirements:

- Backup Exec Agent for Enterprise Vault must be installed on the Enterprise Vault server.
- Enterprise Vault migration and collections must be enabled for the Enterprise Vault partition from which you want to migrate data.
- Enterprise Vault 8.0 SP3 or higher must be installed on the Enterprise Vault server.

How the Backup Exec Migrator works

Enterprise Vault automatically initiates all data migration operations from the Enterprise Vault server after you configure the Backup Exec Migrator. Enterprise Vault makes decisions on what should be migrated based on the archival policies and the data retention policies that you configure in the Enterprise Vault Administration Console. The Backup Exec Migrator then migrates the archived data to a Backup Exec media server after Enterprise Vault collects the eligible data from the vault store partitions. When you configure migration options for a partition, you can set the migration period. All migration options are configured at the Enterprise Vault server.

Table E-7 Enterprise Vault data migration process

Action	Notes
Enterprise Vault archives eligible partition data that is based on the file size or the file creation date.	All data that is eligible for archive is determined in the partition where you want to migrate data. See your Enterprise Vault documentation.

Table E-7 Enterprise Vault data migration process (*continued*)

Action	Notes
After Enterprise Vault completes the archival process, an Enterprise Vault collection process collects the archived data.	<p>The collection process places the archived data into Windows .cab files. The .cab files are stored in the partition where the migration occurs.</p> <p>Eligible data can include Enterprise Vault files with the following extensions:</p> <ul style="list-style-type: none">■ .dvf■ .dvssp■ .dvsc■ .dvs <p>Note: Some eligible data cannot be compressed into .cab files due to file size restrictions. However, the Backup Exec Migrator still migrates the data during the migration operation.</p> <p>See your Enterprise Vault documentation.</p>

Table E-7 Enterprise Vault data migration process (*continued*)

Action	Notes
The Backup Exec Migrator initiates the migration of the archived data files to a Backup Exec media server.	

Table E-7 Enterprise Vault data migration process (*continued*)

Action	Notes
	<p>Migration period schedules are determined when you configure migration for a partition and when you configure a collection schedule for the partition.</p> <p>See “Configuring Enterprise Vault collections” on page 1050.</p> <p>See “Configuring the Backup Exec Migrator to communicate with Enterprise Vault” on page 1054.</p> <p>If you follow the Symantec configuration recommendations for the Backup Exec Migrator and Enterprise Vault partitions, one migration job for each partition runs during a migration period. However, the Backup Exec Migrator may create separate migration jobs for each partition folder if you do not follow the configuration recommendations. If separate jobs are created, the resulting overhead that is required to run the jobs results in degraded migration and retrieval performance.</p> <p>Note: If you schedule a file retrieval request from the Enterprise Vault server between migration periods, separate jobs are created even though you followed the configuration recommendations. In this case, the Backup Exec Migrator automatically creates separate jobs to facilitate retrieval of the requested file. During a migration operation, the restore job can be scheduled to run between migration jobs.</p> <p>If you do not follow the configuration recommendations, file retrieval performance can be affected.</p> <p>To ensure the most efficient migration and retrieval performance possible, follow the Symantec recommendations when you configure the Backup Exec Migrator and the Enterprise Vault partitions.</p> <p>See “Configuring the Backup Exec Migrator”</p>

Table E-7 Enterprise Vault data migration process (*continued*)

Action	Notes
	on page 1049.
Backup Exec completes the migration process by moving all of the migrated files to storage devices.	<p>Symantec recommends configuring two storage devices for staged migration operations.</p> <p>See “About using staged migrations with Backup Exec and the Backup Exec Migrator” on page 1046.</p> <p>See “Configuring the Backup Exec Migrator” on page 1049.</p>

After Backup Exec migrates the .cab files to the storage devices, you can review the migration details in the **Job History** pane on the Backup Exec **Job Monitor** tab.

About using staged migrations with Backup Exec and the Backup Exec Migrator

When you configure Backup Exec to work with the Backup Exec Migrator, Symantec recommends that you configure two storage devices for staged migration operations. When you consider the devices to use, consider selecting a high performance backup-to-disk folder and a slower performance tape device. By using two devices, archived data can be migrated in two stages.

During the first stage, Backup Exec migrates the data it receives from the Backup Exec Migrator to a backup-to-disk folder on a high performance hard drive. By using a backup-to-disk folder, you can minimize the amount of time it takes to perform the initial migration. During the second migration stage, Backup Exec creates a duplicate job to migrate the archived data from the backup-to-disk folder to a tape device. You can schedule the duplicate job to move the archived data to a tape device at times when media server activity is low.

See [“Configuring the Backup Exec Migrator to work with a Backup Exec media server”](#) on page 1051.

See [“Configuring the Backup Exec Migrator to communicate with Enterprise Vault”](#) on page 1054.

About Backup Exec Migrator events

The Backup Exec Migrator generates events that specify the status of the tasks that it runs. The events also provide useful information for troubleshooting

purposes. You can view the events on the computer where you installed the Enterprise Vault Storage Service by viewing the Windows Event Viewer. From the Event Viewer, you can see the events under **Enterprise Vault**. You can also view the events in the Enterprise Vault Dtrace Utility.

For more information on the Enterprise Vault Dtrace Utility, see your Enterprise Vault documentation.

See “[About Backup Exec Migrator logs](#)” on page 1047.

About Backup Exec Migrator logs

The Backup Exec Migrator can create log files that log all migration activity. The log files reside on both the Enterprise Vault server and the Backup Exec media server. Backup Exec Migrator log files can help you troubleshoot migration issues.

Before you can view the log files, you must enable Backup Exec Migrator logging on the Enterprise Vault server and on the Backup Exec media server. To enable Backup Exec Migrator logs on the Enterprise Vault server, edit the Windows registry.

For information on enabling Backup Exec Migrator logging on the Enterprise Vault server, see the following:

<http://entsupport.symantec.com/umi/V-269-27>

To enable Backup Exec Migrator logging on the media server See “[Using the Backup Exec Debug Monitor for troubleshooting](#)” on page 810.

Note: Partition Recovery Utility log files are enabled by default.

After you enable logging on the Enterprise Vault server and on the Backup Exec media server, the following types of log files are created:

- VxBSA log files
For example, <computer_name>-vxbsa<00>.log
- Partition Recovery Utility log files
For example, partitionrecovery<00>.log
- Backup Exec media server log files
For example, <computer_name>-bengine<00>.log

Each time the Backup Exec Migrator is started, separate VxBSA log files are created. As a result, each new log file's sequential number increments by one.

For example, <computer_name>vxbsa00.log, <computer_name>vxbsa01.log.

Similarly, a new log file is created each time the Partition Recovery Utility is started. As a result, each new Partition Recovery Utility log file's sequential number increments by one.

For example, `partitionrecovery00.log`, `partitionrecovery01.log`

Backup Exec media server log file numbers also increment by one as multiple log files are created.

For example, `<computer_name>-bengine00.log`, `<computer_name>-bengine01.log`

You can find the log files in the following locations.

Table E-8 Backup Exec Migrator and Partition Recovery Utility log file locations

Log file	Computer	Directory location
VxBSA log files Partition Recovery Utility log files	Enterprise Vault server	C:\Program Files\Symantec\BACKUP EXEC\RAWS\logs
Backup Exec media server log files	Backup Exec media server	C:\Program Files\Symantec\Backup Exec\Logs

See [“About Backup Exec Migrator events”](#) on page 1046.

About deleting files migrated by Backup Exec Migrator

Enterprise Vault automatically deletes archived items when the item's Enterprise Vault retention periods expire. An Enterprise Vault retention period indicates how long Enterprise Vault retains archived items before it deletes them.

The Backup Exec Migrator maintains existing Enterprise Vault retention periods for archived items when it migrates the archived items to tape. As a result, when an item's data retention period expires, Enterprise Vault issues a command to delete the item from the storage tape that Backup Exec manages. To delete the expired archive item, the `.cab` file it resides in must be deleted from tape.

Note: Although the Backup Exec Migrator maintains existing Enterprise Vault retention periods, it does not initiate the deletion of expired archived items or archived partitions from tape. Only Enterprise Vault can initiate the deletion of expired items and partitions.

For more information on deleting expired items, see your Enterprise Vault documentation.

Because the .cab files may contain archived items with different retention periods, an expired item may be marked as deleted in the Backup Exec catalogs. However, it may not be immediately deleted from tape. All archived items in a .cab file must have expired retention periods before Enterprise Vault issues a command to delete the .cab file from tape.

Enterprise Vault can also delete entire archived vault store partitions from tape. After you delete an active Enterprise Vault vault store partition by using the Enterprise Vault Administration Console, Enterprise Vault deletes the associated archived partition from tape.

Backup Exec automatically recycles the tapes when all of the items on the tape are marked as deleted in the catalogs. Backup Exec checks for expired Enterprise Vault Migrator media once every 24 hours. If Backup Exec detects such media, it logically moves the media to the **Scratch media** node and then generates an information alert informing you of the move.

Note: Expired Enterprise Vault Migrator media is defined as media that contains only migrated Enterprise Vault data that is marked as deleted in the Backup Exec catalogs.

See [“About media in Backup Exec”](#) on page 219.

Note: You should ensure that migrated Enterprise Vault data remains accessible on the tapes that are used for migration purposes until the Enterprise Vault data retention periods expire. Therefore, Symantec recommends that you configure an indefinite retention period for all tapes that are used for migration purposes.

See [“About media overwrite protection”](#) on page 223.

Configuring the Backup Exec Migrator

All of the program files that are required to run the Backup Exec Migrator are installed when you install the Enterprise Vault Agent on the Enterprise Vault server. However, before you can use the Backup Exec Migrator, you must configure it to work with both a destination Backup Exec media server and the Enterprise Vault server.

Table E-9 Enterprise Vault configuration process

Step	Description
Step 1	Configure Enterprise Vault collections. See “Vault store partition properties - Collections” on page 1051.
Step 2	Configure the Backup Exec Migrator to work with a Backup Exec media server. See “Configuring the Backup Exec Migrator to work with a Backup Exec media server” on page 1051.
Step 3	Configure the Backup Exec Migrator to work with Enterprise Vault. See “Configuring the Backup Exec Migrator to communicate with Enterprise Vault” on page 1054.

Use the following configuration recommendations for both the Backup Exec Migrator and the Enterprise Vault partitions:

- Configure the Enterprise Vault partitions to save migrated data locally.
Do not configure Enterprise Vault partitions to delete files immediately after a migration operation finishes.
See your Enterprise Vault documentation for details on configuring a partition for migration.
- Configure the Backup Exec media server template to run staged migrations.
See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1046.

Failure to follow the configuration recommendations results in degraded migration and retrieval performance.

Configuring Enterprise Vault collections

Before you can use the Backup Exec Migrator to migrate Enterprise Vault archived data from a partition, Enterprise Vault first needs to collect the data.

To configure Enterprise Vault collections

- 1 From the Enterprise Vault Console, navigate to a vault store partition from which you want to migrate data.
- 2 Right-click the partiition, and then click **Properties**.

- 3 On the **Collections** tab, check **Use collection files**.
- 4 Set collection options as appropriate.
See [“Vault store partition properties - Collections”](#) on page 1051.
- 5 Click **OK**.

Vault store partition properties - Collections

Before you can use the Backup Exec Migrator to migrate Enterprise Vault archived data from a partition, Enterprise Vault needs to collect the data to be migrated.

See [“Configuring Enterprise Vault collections”](#) on page 1050.

Table E-10 Vault store partition properties - Collection options

Item	Description
Use collection details	Lets you set Enterprise Vault as the collector.
Start at	Indicates the local time at which you want collection to start.
End at	Indicates the local time at which you want collection to finish. Enterprise Vault stops collecting at this time or when it has no more files to collect, whichever comes first.
Limit collection files to <number> megabytes	Indicates the maximum size for collection files. The default size is 10 MB, although you can specify a file size range from 1 MB to 99 MB. You may want to change this value to optimize the use of your backup media
Collect files older than	Indicates the amount of time that must elapse since items were archived before they are eligible for collection.

Configuring the Backup Exec Migrator to work with a Backup Exec media server

Use the following steps to configure the Backup Exec Migrator to work with a destination Backup Exec media server.

Note: Symantec recommends that you configure two media server storages devices when you configure the Backup Exec Migrator to work with Backup Exec. Configuring two storage devices lets you create a staged migration for your archived Enterprise Vault data.

See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1046.

See [“Configuring the Backup Exec Migrator to communicate with Enterprise Vault”](#) on page 1054.

To configure the Backup Exec Migrator to work with a Backup Exec media server

- 1 At the Backup Exec media server, start Backup Exec.
- 2 Create a logon account that uses the Enterprise Vault server Vault Service account credentials.

Vault Service account credentials are used so that Backup Exec and the Backup Exec Migrator can complete the migration operation

See [“Creating a Backup Exec logon account”](#) on page 187.
- 3 On the navigation bar, click **Tools** and then click **Options**.
- 4 Under **Job Defaults**, click **DBA-initiated Job Settings**.
- 5 Select the **DEFAULT** template, and then click **Edit**.

You can also use an existing template, or you can create a new template specifically for Enterprise Vault migrations.
- 6 Under **Backup Job Template**, click **Device and Media**.
- 7 Select a backup-to-disk folder as the primary storage location for migrated data, and then set the options you want to use with the device.
- 8 Under **Migrator for Enterprise Vault**, click the down arrow next the field for **Vault Service account credentials**.
- 9 Select the logon account that you created in step 2.

See [“Migrator for Enterprise Vault options”](#) on page 1053.
- 10 Under **Backup Job Template**, set other options as appropriate.

See [“Editing DBA-initiated jobs”](#) on page 434.
- 11 Do one of the following:

- If you want to configure staged migrations Do the following in the order listed.
- See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1046.
- Under **Duplicate Job Template**, click **Settings**.
 - Check **Enable settings to duplicate backup sets for this job**.
 - In the **Device** list, select a tape device.
 - Set other options as appropriate. See [“Duplicate job template settings for DBA-initiated jobs”](#) on page 430.
 - Click **OK**.
- If you do not want to configure staged migrations Continue with step 12.

12 Click **OK**.

13 Configure Backup Exec Migrator to work with Enterprise Vault.

See [“Configuring the Backup Exec Migrator to communicate with Enterprise Vault”](#) on page 1054.

Migrator for Enterprise Vault options

The Backup Exec Migrator uses the Enterprise Vault server Vault Service account during the Backup Exec Migrator to Backup Exec media server authentication process.

Table E-11 Migrator for Enterprise Vault options

Item	Description
<p>Vault Service account credentials</p>	<p>Specifies the Enterprise Vault server Vault Service account credentials to use so that Backup Exec and the Backup Exec Migrator can complete the migration operation.</p> <p>The Vault Service account must be included in either the Administrators group or the Backup Operators group at the Backup Exec media server.</p> <p>Note: If the Enterprise Vault server and the Backup Exec media server are in different domains, a trust relationship must be established between the domains. The Vault Service account user must be a trusted user at the Backup Exec media server. Trust relationships are required so that the Microsoft Security Support Provider Interface (SSPI) can authenticate the Vault Service account user.</p> <p>For more information on domain trust relationships, see your Microsoft documentation.</p>
<p>New</p>	<p>Lets you create a new logon account or edit an existing account.</p> <p>See “Creating a Backup Exec logon account” on page 187.</p>

Configuring the Backup Exec Migrator to communicate with Enterprise Vault

Use the following steps to configure the Backup Exec Migrator to communicate with Enterprise Vault.

See [“Configuring the Backup Exec Migrator”](#) on page 1049.

To configure the Backup Exec Migrator to communicate with Enterprise Vault

- 1 At the Enterprise Vault server, navigate to a vault store partition from which to migrate data.
- 2 Right-click the vault store partition, and then click **Properties**.
- 3 On the **Migration** tab, check **Migrate files**.

- 4 In **Remove collection files from primary storage**, set the time period for this option to something longer than zero days.

Do not set it to zero days. Setting the time period to zero days causes Enterprise Vault to immediately delete the migrated data from the partition. More importantly, it causes the Backup Exec Migrator to create separate migration jobs for each partition folder being migrated during a migration period. If separate jobs are created, the resulting overhead that is required to run the jobs results in degraded migration and retrieval performance.

See [“Configuring the Backup Exec Migrator”](#) on page 1049.
- 5 Set other migration options as appropriate.

See [“Vault store partition properties - Migration options”](#) on page 1056.
- 6 On the **Advanced** tab, ensure that **Symantec Backup Exec** appears in the **List setting from** field.
- 7 In the window below the **List setting from** field, select **Backup Exec media server**.
- 8 Click **Modify**.
- 9 Type the name or the IP address of the destination Backup Exec server.
- 10 Click **OK**.
- 11 Select **Backup Exec DBA-initiated template**.
- 12 Click **Modify**.
- 13 Enter the name of an existing template that uses the Enterprise Vault server Vault Service account credentials.

The template that you select must be configured to use the Enterprise Vault server Vault Service account. The template you use must also match the template name that you used when you configured the Backup Exec Migrator to work with a media server.

See [“Configuring the Backup Exec Migrator to work with a Backup Exec media server”](#) on page 1051.
- 14 Click **OK**.
- 15 Ensure that the name of the template that contains the Enterprise Vault server Vault Service account credentials appears in the **Setting** pane.

See [“Configuring the Backup Exec Migrator to work with a Backup Exec media server”](#) on page 1051.
- 16 To test the communications between the Enterprise Vault server and the Backup Exec media server, click **Test Configuration**.

- 17 If the test fails, ensure that you used the correct credentials for the Vault Service account , and then click **Test Configuration** again.
- 18 Click **OK** after the test successfully completes.
- 19 Click **OK**.

Vault store partition properties - Migration options

Select the Enterprise Vault migration property options that you want to use.

Table E-12 Vault store partition properties - Migration options

Item	Description
Migrate files	Lets you migrate archived Enterprise Vault data to a Backup Exec storage device. Migration can help reduce storage costs by moving collection files to tertiary storage devices. However, retrieval times can increase. See your Enterprise Vault documentation.
Migrator	Indicates the name of the migration application. Symantec Backup Exec must appear in this field.
Migrate files older than	Indicates the amount of time that must elapse since files were last modified before they are eligible for migration. See your Enterprise Vault documentation.

Table E-12 Vault store partition properties - Migration options (*continued*)

Item	Description
Remove collection files from primary storage	<p>Indicates the age at which migrated collection files are removed from the primary storage location.</p> <p>Files that have been migrated to Backup Exec storage media can remain in their primary location for the period of time you specify.</p> <p>Note: Symantec recommends that you set the time period for this option to something longer than zero days, with a longer time period being best. Do not set it to zero days. Setting the time period to zero days causes the Backup Exec Migrator to create separate migration jobs in a migration period for each partition being migrated. If separate jobs are created, the resulting overhead that is required to run the jobs results in degraded migration and retrieval performance.</p> <p>See “Configuring the Backup Exec Migrator” on page 1049.</p>

About the Restore view of migrated Enterprise Vault data

The Backup Exec restore view lets you visually verify the success of an archived Enterprise Vault data migration operation. Backup sets containing the migrated .cab files appear under a partition name that reflects the Enterprise Vault partition from where the data was migrated. Because the restore view displays the archived data in a read-only mode, you cannot select the data for restore. However, you can retrieve the data in the application where the data resides.

Note: You can completely retrieve of all archived items that appear in the restore view by using the Partition Recovery Utility.

See [“About the Partition Recovery Utility”](#) on page 1059.

See [“About retrieving migrated Enterprise Vault data”](#) on page 1058.

About retrieving migrated Enterprise Vault data

All file retrieve operations start from the Enterprise Vault server console. You cannot restore archived Enterprise Vault data from Backup Exec.

When files are migrated from a partition, Enterprise Vault creates a shortcut in the partition that replaces the migrated file. The shortcut also links to the storage location of the migrated file. You retrieve files by double-clicking their shortcuts in the Enterprise Vault partition itself. If a partition retains a local copy of the migrated files, Enterprise Vault retrieves the files from the local copies. If Enterprise Vault deletes the migrated files because the partition's file retention period passes, the requested files must be retrieved from Backup Exec storage media.

Table E-13 How migrated data is retrieved

Action	Notes
Enterprise Vault works with the Backup Exec Migrator to begin the process.	The Backup Exec Migrator identifies the Backup Exec media server where the files are stored.
The Backup Exec Migrator schedules a Backup Exec restore job at the media server.	Backup Exec restores the requested files.
The Backup Exec Migrator migrates the restored files to the Enterprise Vault server partition from the Backup Exec media server.	The Backup Exec Migrator moves the restored files to a location specified by Enterprise Vault, using the name provided by Enterprise Vault.

The retrieval process is automatic after you start the operation at the Enterprise Vault server. It requires no user intervention other than perhaps placing a tape in the tape device if you removed the storage media.

See [“Retrieving migrated Enterprise Vault data”](#) on page 1058.

Retrieving migrated Enterprise Vault data

Use the following steps to restore migrated Enterprise Vault files.

Note: To successfully retrieve the files you want, you may need to place a tape in a tape drive at the Backup Exec media server.

To retrieve migrated Enterprise Vault data

- 1 At the Enterprise Vault server, navigate to the partition where you want to retrieve the data.
- 2 Double-click the file that you want to retrieve.

About the Partition Recovery Utility

The Partition Recovery Utility is a command-line application that is automatically installed when you install the Backup Exec Remote Agent for Windows Systems. The utility lets you restore all of a partition's archived files from the Backup Exec storage media in a single operation. You can also use it to recover the archived partition data for each of the Enterprise Vault partitions in a disaster recovery situation.

After you use the Partition Recovery Utility, you can review the restore details in the **Job History** pane on the Backup Exec **Job Monitor** tab.

See [“Partition Recovery Utility requirements”](#) on page 1059.

See [“Finding an archive ID”](#) on page 1060.

See [“Starting the Partition Recovery Utility”](#) on page 1060.

Partition Recovery Utility requirements

You must know the following when you use the Partition Recovery Utility:

- The vault store partition name for the data that you want to recover.
- The Archive ID of the partition data that you want to recover.
- An Enterprise Vault server user account with Vault Service Account privileges.

Note: If you run the Partition Recovery Utility on a Windows Server 2008/2008 R2 computer, administrator privileges are required.

In addition, the Partition Recovery Utility must run at the Enterprise Vault server that originally migrated the data you want to restore.

See [“Finding an archive ID”](#) on page 1060.

See [“Starting the Partition Recovery Utility”](#) on page 1060.

Finding an archive ID

You use the archive ID of the data you want to restore along with the vault store partition name when you run the Partition Recovery Utility. The archive ID is an alpha-numeric number of considerable length.

For example, 1D69957C6D917714FB12FEA54C9A8299A1110000ev8archive.EVMBE

You can find the Archive ID listed among the properties of an archived file set.

To find an archive ID

- 1 In the left view of the the Enterprise Vault Administration Console, expand **Archives**.
- 2 Navigate the folder structure and select the folder of the type for data you want to restore.
- 3 In the right view, right-click an archive, and then select **Properties**.
- 4 On the **Advanced** tab, note the archive ID at the bottom.

Starting the Partition Recovery Utility

Use the following steps to start the Partition Recovery Utility.

To start the Partition Recovery Utility

- 1 From the Enterprise Vault server, open a Windows command prompt.
- 2 Navigate to the Enterprise Vault Agent installation directory.
For example, C:\Program Files\Symantec\Backup Exec\RAWS
- 3 Do the following:

If you start the Partition Recovery Utility on a Windows Server 2008/2008 R2 computer	Type the following command: <pre>runas /user:<domain\administrator> partitionrecovery.exe -vs <vault_store_name> -ap <archive_ID></pre>
---	--

If you start the Partition Recovery Utility on all other supported Windows operating system versions	Type the following command: <pre>partitionrecovery.exe -vs <vault_store_name> -ap <archive_ID></pre>
--	---

- 4 Press **Enter**.

Best practices for using the Backup Exec Migrator

Consider the following best practices when you use the Backup Exec Migrator:

- Symantec recommends that you regularly back up the Backup Exec catalogs. In the event the catalogs become corrupt, you can restore them from backups. After you restore the catalogs, you must re-catalog the storage media on which Backup Exec Migrator data is stored. Re-cataloging the storage media ensures that the latest catalog entries are available.
- For best performance, configure the Backup Exec Migrator to migrate data to a backup-to-disk folder and then to a tape device by using a duplicate job. See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1046. See [“About duplicating backed up data”](#) on page 369.
- In the Enterprise Vault **Migration** options tab, set the time period for **Remove collection files from primary storage** to something longer than zero days. Setting the time period to zero days causes Enterprise Vault to immediately delete the migrated data from the partition.

If you set the time period to zero days, Symantec recommends the following:

- Increase the number of concurrent jobs that are allowed for the backup-to-disk folder you use for migration purposes. Increase the number of concurrent jobs based on the following formula:
<number of recommended concurrent jobs> = <number of installed tape drives plus two>

For example, if you have two installed tape drives, you should configure the backup-to-disk folder to allow four concurrent jobs.

Concurrent jobs let the Backup Exec Migrator continue to migrate data to disk storage while tape drives process duplicate jobs in a staged migration environment.

Note: You can increase the number of concurrent jobs that run by increasing the total concurrency level of the backup-to-disk devices.

- Symantec recommends that you first collect all of the archived files in one collection and migration operation and then migrate them in the next collection and migration operation. This process helps ensure that the Backup Exec Migrator creates a single job for each migration operation, which improves the migration performance.

Troubleshooting Backup Exec Migrator and Partition Recovery Utility issues

Review the following error messages for possible solutions to errors that you may encounter:

- The Backup Exec Migrator logs migration activity in the Windows Event Viewer and in the Enterprise Vault Dtrace Utility on the Enterprise Vault server. It also logs migration activity on the Backup Exec media server. The details that are provided in the log files can help you troubleshoot issues with the Backup Exec Migrator.
See [“About Backup Exec Migrator events”](#) on page 1046.
See [“About Backup Exec Migrator logs”](#) on page 1047.
- The Partition Recovery Utility cannot find any files to be recalled. There are no file to be recalled from the vault store database using the Archive ID that you provided.
- The Partition Recovery Utility operation will be terminated due to a user request.
You may have stopped the Partition Recovery Utility operation by pressing **Ctrl + C** or **Ctrl + Break**.
- The migrated file name <file_name> with ID <migrated_file_id> was not found in the Backup Exec backup sets. The recall is skipped for this file. The Partition Recovery Utility skips collection files if they already exist in the vault store database. To restore the files, delete them from the vault store database, and then run the Partition Recovery Utility again.
- The Partition Recovery Utility cannot find any partitions. Ensure that the name of the vault store is valid, and that there are partitions in the vault store. The vault store name that you provided may be invalid.

Symantec Backup Exec Agent for Lotus Domino Server

This appendix includes the following topics:

- [About the Agent for Lotus Domino Server](#)
- [Lotus Domino Agent requirements](#)
- [About installing the Lotus Domino Agent on the media server](#)
- [About the Lotus Domino Agent and the Domino Attachment and Object Service \(DAOS\)](#)
- [Viewing Lotus Domino databases that are created while Backup Exec is running](#)
- [Viewing Lotus Domino databases that are on the local server](#)
- [Viewing Lotus Domino databases that are on remote computers](#)
- [Configuring default Lotus Domino options](#)
- [About backing up Lotus Domino databases](#)
- [About selecting Lotus Domino databases for backup](#)
- [Selecting backup options for Lotus Domino databases](#)
- [Restoring Lotus Domino databases](#)
- [About selecting Lotus Domino databases for restore](#)
- [Selecting restore options for Lotus Domino databases](#)

- [Redirecting restore jobs for Lotus Domino databases](#)
- [Redirecting the restore of DAOS NLO files](#)
- [How to prepare for disaster recovery on a Lotus Domino server](#)

About the Agent for Lotus Domino Server

The Symantec Backup Exec Agent for Lotus Domino Server (Lotus Domino Agent) is installed as a separate, add-on component of Backup Exec.

You can use the Lotus Domino Agent to back up and restore Lotus Domino on local media servers and on remote computers. The Lotus Domino Agent backs up Lotus Domino databases, Domino Attachment and Object Service (DAOS) - related NLO files, and transaction logs. You can integrate Lotus Domino database backups with regular server backups without separately administering them or using dedicated hardware.

The Lotus Domino Agent provides support for the following:

- Full, incremental, and differential online backups of Lotus Domino databases, DAOS-related NLO files, and transaction logs using Lotus Domino APIs.
- Restores of Lotus Domino databases, .nlo files, archived transaction logs, and point-in-time restores.
- Recycling of archived Lotus Domino transaction logs after a successful backup.
- Flexible scheduling capabilities.
- Backup and restore of partitioned and clustered Lotus Domino servers.
- Lotus Domino databases in a Microsoft Cluster Server cluster in both Active-Active and Active-Passive configurations.

See [“About installing the Lotus Domino Agent on the media server”](#) on page 1066.

Lotus Domino Agent requirements

The Lotus Domino Agent supports the backup and restore of Lotus Domino versions 7.x and 8.x.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Following are the requirements for backing up Lotus Domino database files residing on the media server, or for remote Windows computers and workstations.

Note: Backup Exec does not support two versions of Lotus Domino on the same computer.

If the Lotus Domino files you want to back up are on the local media server, the server must have the following:

- Backup Exec
- An Intel-compatible processor
- The Lotus Domino data directory on the Lotus Domino server

If the Lotus Domino files you want to back up are on a remote computer, the remote computer must have the following:

- Windows operating system
- Backup Exec Remote Agent for Windows Systems
- An Intel-compatible processor
- Corresponding Windows Administrative Share for each volume that contains Lotus Domino databases
- The Lotus Domino data directory on the Lotus Domino server

The following are required to back up Lotus Domino transaction logs:

- Archive-style transaction logging must be enabled to perform differential and incremental backups and to perform point in time recovery.
- The Lotus Domino logging style must be set to archive if you want to back up the transaction logs.

The following are required to back up Lotus Domino DAOS-related NLO files:

- The DAOS state must be in read-only mode or enabled.
- The DAOS catalog should be synchronized.

If the Lotus Domino databases are running in a Microsoft Cluster Server cluster, you must have the following:

- Lotus Domino Server must be running on a Microsoft Cluster Server cluster. For more information, see your Lotus Domino documentation for instructions for setting up Lotus Domino in a Microsoft Cluster Server cluster.
- The Backup Exec Lotus Domino Agent must be installed on all nodes in the Microsoft Cluster Server cluster.

See [“Viewing Lotus Domino databases that are created while Backup Exec is running”](#) on page 1068.

About installing the Lotus Domino Agent on the media server

The Symantec Backup Exec Agent for Lotus Domino is installed locally as a separate, add-on component of Backup Exec. It can protect local or remote Lotus Domino databases.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

See [“Push-installing the Remote Agent and Advanced Open File Option to remote computers”](#) on page 134.

See [“Using a command prompt to install the Remote Agent on a remote computer”](#) on page 145.

Note: If you install Lotus Domino on the same server on which Backup Exec is already installed, you must restart Backup Exec services to display Lotus Domino Database selections.

See [“Starting and stopping Backup Exec services”](#) on page 168.

See [“Configuring default Lotus Domino options”](#) on page 1069.

See [“About backing up Lotus Domino databases”](#) on page 1071.

See [“Restoring Lotus Domino databases”](#) on page 1078.

About the Lotus Domino Agent and the Domino Attachment and Object Service (DAOS)

Lotus Domino 8.5 incorporates the Domino Attachment and Object Service (DAOS). DAOS-enabled databases (DAOS databases) save significant hard drive space by sharing data between applications on a server. DAOS databases do not save separate copies of every document attachment. Instead, DAOS databases save a single copy of a file attachment to an internal repository. The databases then create and save reference pointers to the stored file attachments.

File attachments are saved to the internal repository with the .nlo file extension. During a full backup of the entire Lotus Domino server, Backup Exec backs up all .nlo files, along with the Domino <server>.id file.

Backup Exec adds one container per partition under **Lotus Domino Databases** called **Domino Attachment and Object Service** in the restore selections view. All backed up DAOS NLO files reside in backup sets under **Domino Attachment and**

Object Service. In addition, all backed up <server>.id files reside in the **Databases** container under **Lotus Domino Databases**.

Note: Domino uses <server>.id for NLO encryption purposes. If you enable NLO file encryption at the Domino server, the <server>.id file must be backed up.

When you select individual DAOS-enabled databases for backup, the referenced .nlo files for each database are included in the backup job. However, the <server>.id file is excluded.

With incremental backups, only the databases and the .nlo files that are created since the last Full backup of the server are backed up.

Note: In cases where incremental backup jobs fully back up DAOS-enabled databases, all of the .nlo files that each database references are backed up. This scenario occurs when DAOS-enabled databases use circular logging, or when DAOS-enabled databases are in archive log mode and their DBIID changes.

During a full DAOS-enabled Domino database restore, all database data, .nlo files, and the <server>.id file are restored. When you restore individual DAOS-enabled databases, Backup Exec restores all database data, including the .nlo files. However, Backup Exec does not restore any .nlo files that match .nlo files in the internal repository. After the DAOS-enabled databases are restored, Backup Exec resynchronizes the Domino DAOS catalog.

During a point-in-time restore of a DAOS-enabled database, some required .nlo files may not be generated when the archive transaction logs are replayed. When this condition occurs, Backup Exec reports the names of the missing .nlo files. You can individually restore the missing .nlo files and then start a Domino DAOS catalog resynchronization operation at the Domino server.

For information on Domino DAOS catalog resynchronization operations, see your Lotus Domino documentation.

Best practices for restoring the missing .nlo files

If you decide to individually restore the missing .nlo files, Symantec recommends the following best practices:

- Always restore the .nlo files to the internal repository of the current DAOS-enabled Domino server.
- Instead of making random .nlo files selections, select all of them when you make your selections in the restore selections view. Then use the restore option,

Skip if file exists. By using the **Skip if file exists** option, Backup Exec restores only the missing .nlo files.

See [“Redirecting the restore of DAOS NLO files”](#) on page 1084.

Viewing Lotus Domino databases that are created while Backup Exec is running

Use the following steps to view Domino databases that are created while Backup Exec is running.

See [“About backing up Lotus Domino databases”](#) on page 1071.

See [“Selecting Lotus Domino databases for backup”](#) on page 1076.

Note: Use the virtual computer name or the virtual IP address of the Domino server to browse or submit jobs in a Microsoft Cluster Server environment.

To view the databases that are created while Backup Exec is running

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the selections pane, expand **Lotus Domino Databases**.
- 4 Press **F5**.

Viewing Lotus Domino databases that are on the local server

Use the following steps to view Domino databases that are on the local server.

Lotus Domino transaction logs do not appear under **Lotus Domino Databases**; however, when you select the database for backup, the transaction logs are automatically included.

The same process applies to DAOS NLO files. They do not appear under **Lotus Domino Databases**; however, when you select the database for backup, the .nlo files are automatically included.

See [“About backing up Lotus Domino databases”](#) on page 1071.

See [“Selecting Lotus Domino databases for backup”](#) on page 1076.

To view Lotus Domino databases on the local server

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the selections pane, expand **Lotus Domino Databases**.

Viewing Lotus Domino databases that are on remote computers

Use the following steps to view Domino databases that are on remote computers.

See [“About backing up Lotus Domino databases”](#) on page 1071.

See [“Selecting Lotus Domino databases for backup”](#) on page 1076.

See [“How to prepare for disaster recovery on a Lotus Domino server”](#) on page 1086.

To view Lotus Domino databases on remote computers

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 Click **Remote Selections**, and then click **Microsoft Windows Network**.
- 5 If necessary, click the domain that contains the Lotus Domino installations, and then click the computer in which the Lotus Domino database is located.

A list of shared network directories appears, along with an icon that represents the Lotus Domino Databases.

Configuring default Lotus Domino options

You can configure default settings for Lotus Domino databases for all new jobs you create. When you create a job, you can use the default settings or modify the Domino properties for the job.

See [“About backing up Lotus Domino databases”](#) on page 1071.

To configure default Domino options for all new jobs

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Lotus Domino**.

- 3 Select the appropriate options.
 See [“Lotus Domino default options”](#) on page 1070.
- 4 Click **OK** to save the options or select other options from the Properties pane.

Lotus Domino default options

You can use the default options that were set when Backup Exec is installed or you can change the options for all Lotus Domino jobs.

See [“Configuring default Lotus Domino options”](#) on page 1069.

Table F-1 Lotus Domino default options

Item	Description
Backup method	<p>Specifies one of the following backup methods to use:</p> <ul style="list-style-type: none"> ■ Full - Back up Database and Logs - Reset Archive Bit. Select this to back up all the selected databases. To properly back up your Lotus Domino data, you should perform regular full backups of the database. This backup method should also be used when the DBIID for the database has changed since prior transactions cannot be applied to the new database. ■ Differential - Changed Database and Logs. Select this to back up files that were modified since the last Full backup. This backup method is smaller and faster than a Full backup because only archived transaction logs, unlogged databases, and logged databases with DBIIDs that have changed will be backed up. Note: Only the .nlo files that change are backed up when DAOS-enabled databases are in archive log mode and their DBIIDs do not change. ■ Incremental - Changed Database and Logs - Reset Archive Bit. Select this to back up files that were modified since the last Full or Incremental backup. This backup method is smaller and faster than a Full backup because only archived transaction logs, unlogged databases, and logged databases with DBIIDs that have changed will be backed up. Note: Only the .nlo files that change are backed up when DAOS-enabled databases are in archive log mode and their DBIIDs do not change.

Table F-1 Lotus Domino default options (*continued*)

Item	Description
Mark archive logs for recycling	<p>Reuses the transaction log after it has been backed up.</p> <p>Backup Exec will not delete the transaction log. Selecting this option only indicates that the transaction log is ready to be reused after it has been backed up successfully; the Lotus Domino server actually deletes the transaction logs.</p> <p>This option is selected automatically when you select the full backup method. You cannot clear this option when you are using the full backup method.</p> <p>If the option is selected when you perform a differential or incremental backup job, transaction logs that are needed to maintain the differential backups will be reused. However, it should be selected regularly to create space for new transaction logs.</p>
Seconds to wait for the database to go offline	<p>Specifies the number of seconds for the restore process to wait for a database that is in use. When a Lotus database is restored it must first be taken offline. This will ensure that the database is not being accessed, closed, or deleted while the restore operation is being processed. If the database is still in use and cannot be taken offline after the specified wait time, the restore will fail.</p>
Retain original IDs	<p>Restores the original database IDs.</p>
Assign new database ID	<p>Assigns new IDs to the database.</p>
Assign new database ID and replica ID	<p>Assigns new IDs to the database. A replica ID is used to synchronize two or more databases that are being replicated in the Lotus Domino environment. You can assign a new replica ID during a restore to prevent other databases under replication from overwriting the restored database files.</p>

About backing up Lotus Domino databases

When a Lotus Domino backup job is submitted, Backup Exec uses Lotus Domino APIs to obtain the backup of the database. When you back up a DAOS-enabled Domino database, the DAOS NLO files are automatically included. In addition, the transaction logs associated with the Lotus Domino databases are included in the backup only if archive logging is turned at the server. If they are backed up, the archive logs are stored in a separate backup set that is stored within the Lotus Domino database backup set.

The Lotus Domino Agent supports the backup of the following types of files:

- .ntf - Lotus Notes Template Files
- .nsf - Lotus Notes Database Files
- .box - Lotus Mailbox Files
- .dsk - Cache Files
- .txn - Transaction log files
- .nlo - DAOS attachment files

Note: Transaction log files and DAOS attachment files do not appear in the Backup Exec backup selections view; however, they do appear in restore selections view.

You must back up .nsf, .ntf, and .box files to properly recover Lotus Domino databases. If you want to back up .njf, .ncf, .id, .dic, or notes.ini files, you must select them for backup from the volume in which the Lotus Domino Program directory is located.

Although DAOS and non-DAOS Domino servers use additional Domino-related databases and support files, Backup Exec does not back them up. Domino automatically recreates the items after you restart the Domino servers.

Backup Exec excludes the following support files from backup jobs:

- daos.cfg
- daoscat.nsf
- dbdirman.nsf

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux Servers is installed.

See [“About Lotus Domino transaction logs”](#) on page 1074.

See [“Selecting backup options for Lotus Domino databases”](#) on page 1076.

See [“Selecting restore options for Lotus Domino databases”](#) on page 1082.

About automatic exclusion of Lotus Domino files during volume-level backups

If you select a volume that contains Lotus Domino data for backup, the Lotus Domino Agent determines which Domino data should not be included in a volume level backup. For example, .ntf and .nsf files, nlo files, <server>.id files, as well as any active log files, should not be part of the backup because they are opened for exclusive use by the Lotus Domino system. These files will be automatically excluded for backup by a feature called Active File Exclusion. If this exclusion did not happen during a non-snapshot backup, these files would appear as in use - skipped. If this exclusion did not happen during a snapshot backup, the files would be backed up in a possible inconsistent state, which could create restore issues.

About supported Lotus Domino database configurations

You can back up the following types of Lotus Domino database configurations using the Lotus Domino Agent:

- **Domino Server Databases.**

Domino Server databases can be Logged or Unlogged, with their DAOS states being not-enabled, read-only, or enabled. DAOS cannot be enabled on Domino databases that do not use logging. Domino databases are located in a folder in the Domino data directory, typically Lotus\Domino\Data, but may also be linked to the Domino data directory using Lotus Linked Databases.

The following types of Lotus Domino databases are supported:

- **Logged Domino Server Databases.**

A logged Domino Server database logs transactions for one or more Lotus databases. If transaction logging is enabled on the server, all database transactions go into a single transaction log.

- **Unlogged Domino Server Databases.**

An unlogged Domino Server database does not have transaction logging enabled, or the transaction logging has been disabled for specific server databases. Unlogged Domino Server databases will be backed up in their entirety when a full, differential, or incremental backup is performed, but the database can only be restored to the point of the latest database backup.

- **Local Databases.**

Lotus databases are considered Local when they cannot be found in the Domino data directory, cannot be shared, and cannot be logged. This type of database requires a backup of the database itself when using any of the Lotus Domino backup methods. The database can be restored only to the point of the latest database backup.

About Lotus Domino transaction logs

Lotus Domino has the ability to log transactions for one or more Lotus Domino databases. Lotus Domino databases are logged by default when transaction logging is enabled on the Lotus Domino server and the database is in the Domino data directory.

When transaction logging is enabled on the server, each Lotus Domino database is assigned a database instance ID (DBIID). Each transaction recorded in the log includes the DBIID, which is used to match transactions to the database during a restore.

A new DBIID may be assigned to the database when some Lotus Domino operations are performed. When the new DBIID is assigned, all new transactions recorded in the log use the new DBIID; however, previous transactions have the old DBIID and will not match the new DBIID for the database. To prevent data loss, it is recommended that a full backup be performed when a database receives a new DBIID since transactions with the old DBIID cannot be restored to the database. A full backup includes all current transactions on the database and ensures that only the transactions with the new DBIID are needed to restore the database.

You can select only one logging style when transaction logging is enabled on the server.

Following are the two styles of logging for Lotus Domino databases:

- **Archive logging.**

This logging style produces a transaction log that is limited only by the capacity of your mass storage. Archive logging is the recommended logging style to be used with the Lotus Domino Agent since all the transaction logs can be backed up and marked for recycling. When the transaction logs are recycled the Lotus Domino server reuses the existing transaction logs after they are backed up to create space for new transaction logs.

- **Circular logging.**

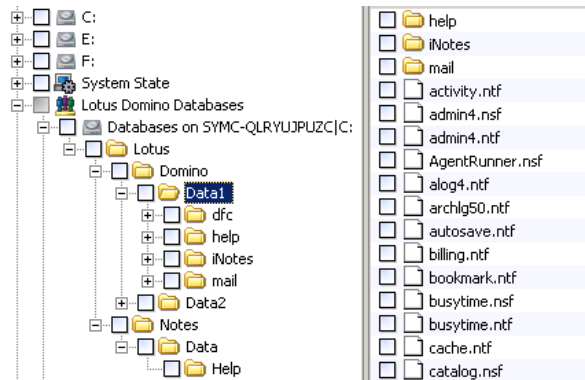
This logging style reuses the log file after a specific log file size is reached. By reusing the log file you are saving resources; however, you are also limiting your recovery options because the database can only be recovered to the point of the last full backup. If the incremental or differential backup method is selected for a backup job, a full backup of the changed databases is performed since transaction logs cannot be backed up.

Caution: When circular logging is enabled, the circular transaction log cannot be backed up, which could result in the loss of changes made to the database since the last backup was performed.

About selecting Lotus Domino databases for backup

After installing the Lotus Domino Agent, you can select existing Lotus Domino databases in the selections pane.

Figure F-1 Domino server file types that appear in the Backup Exec selections view



The following file types appear in the view for the Lotus Domino server:

- filename.nsf - Lotus Domino database files
- filename.ntf - Lotus Domino template files
- filename.box - shared mail database
- filename.dsk - cache files

You must back up all of these files in order to properly recover Lotus Domino Databases.

Only database files and the <server>.id file appear under the Lotus Domino Databases view. Domino Program files and other files such as .id and notes.ini appear in the volume in which the Lotus Domino Program directory is located. They must be backed up separately as part of a system backup.

Note: A full backup of the Domino server includes the <server>.id file. As a result, the Active File Exclusion feature automatically excludes the <server>.id file.

See [“Selecting Lotus Domino databases for backup”](#) on page 1076.

Selecting Lotus Domino databases for backup

After installing the Lotus Domino Agent, you can select existing Lotus Domino databases in the selections pane.

See [“About selecting Lotus Domino databases for backup”](#) on page 1075.

To select Lotus Domino databases

- ◆ Select the check box next to the volume to choose all of the databases in a volume, or expand the volume and select specific folders and databases. When selecting databases to back up, the databases must be local to the Lotus Domino server.

Selecting backup options for Lotus Domino databases

This procedure details how to select backup job properties for Lotus Domino databases. You should back up Lotus Domino databases during off-peak hours and disable Lotus Domino or third-party Lotus Domino agents before you run the backup. The archived transaction logs are included automatically.

See [“Creating a backup job by using the Backup Wizard”](#) on page 331.

Caution: All Lotus Domino databases and transaction logs that reside on single or multiple volumes must be backed up by the same media server. In addition, you should not back up a Lotus Domino server simultaneously from multiple media servers.

To select backup job properties for Lotus Domino databases

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Settings**, click **Lotus Domino**.
- 4 Select the appropriate options.
See [“Lotus Domino backup job options”](#) on page 1076.
- 5 Start the backup job or select other backup options from the **Properties** pane.
See [“Creating a backup job by setting job properties”](#) on page 332.

Lotus Domino backup job options

You can set Lotus Domino-specific options when you create a backup job.

See [“Selecting backup options for Lotus Domino databases”](#) on page 1076.

The following table describes the Lotus Domino options that can be set when you create a backup job.

Table F-2 Lotus Domino Backup Job Properties

Item	Description
Backup method	<p>Specifies one of the following backup methods:</p> <ul style="list-style-type: none"> ■ Full - Back up Database and Logs - Reset Archive Bit. Backs up all the selected databases. To properly back up your Lotus Domino data, you should perform regular full backups of the database. This backup method should also be used when the DBIID for the database has changed since prior transactions cannot be applied to the new database. ■ Differential - Changed Database and Logs. Backs up files that were modified since the last Full backup. This backup method is smaller and faster than a Full backup because only archived transaction logs, unlogged databases, and logged databases with DBIIDs that have changed will be backed up. Note: Only the .nlo files that change are backed up when DAOS-enabled databases are in archive log mode and their DBIIDs do not change. ■ Incremental - Changed Database and Logs - Reset Archive Bit. Backs up files that were modified since the last Full or Incremental backup. This backup method is smaller and faster than a Full backup because only archived transaction logs, unlogged databases, and logged databases with DBIIDs that have changed will be backed up. Note: Only the .nlo files that change are backed up when DAOS-enabled databases are in archive log mode and their DBIIDs do not change.
Mark archive logs for recycling	<p>Reuses the transaction log after it has been backed up.</p> <p>Backup Exec will not delete the transaction log. Selecting this option only indicates that the transaction log is ready to be reused after it has been backed up successfully; the Lotus Domino server actually deletes the transaction logs.</p> <p>This option is selected automatically when you select the full backup method. You cannot clear this option when you are using the full backup method.</p> <p>If the option is selected when you perform a differential or incremental backup job, transaction logs that are needed to maintain the differential backups will be reused. However, it should be selected regularly to create space for new transaction logs.</p>

Restoring Lotus Domino databases

Restoring a Lotus Domino database is a three-part process.

Table F-3 Restoring a Lotus Domino database

Step	Description
Step 1	<p>Restore database files to the Domino server.</p> <p>During a restore of the Lotus Domino database, the existing database is taken offline and deleted, the database is restored, and changed records contained in the backup job are applied to the database.</p> <p>Note: Domino servers include databases with names such as <code>admin4.nsf</code>, <code>names.nsf</code>, and <code>busytime.nsf</code>. The Notes client computers include databases with names such as <code>bookmark.nsf</code>, <code>cache.dsk</code>, and <code>homepage.nsf</code>. These databases are critical and cannot be taken offline when the Domino server and the Notes client are running. In addition, you should only restore these databases in disaster recovery situation.</p> <p>If the database is unlogged or local, the database is brought back online. If the database is logged and multiple databases are being restored, the database name is added to a list for recovery. During the restore process, Backup Exec assigns a unique name to databases and then before databases are brought online, reassigns the original name. Changing the name during the restore process has no effect on restored databases.</p>
Step 2	Restore DAOS-related NLO files that are missing.

Table F-3 Restoring a Lotus Domino database (*continued*)

Step	Description
Step 3	<p>Run transaction logs to bring the database up-to-date.</p> <p>The internal Domino recovery process automatically begins after the DAOS NLO files are restored to the server. The database is restored to a point in time using transactions from the required transaction logs. Required transaction logs that were backed up and recycled are also included in the recovery process. After the recovery process completes, the Lotus Domino database is brought online.</p> <p>If you back up your Lotus Domino databases regularly, then restoring the most recent backup set containing the Lotus Domino data is all that is required to restore the most recent backups of your Lotus Domino databases.</p> <p>Note: If circular logging is enabled and both the databases and the Domino transaction logs are lost, the database can only be recovered to the point of the last full backup.</p>

Use the same procedures to restore a server in a Microsoft Cluster Server cluster that you use to restore a server in a non-clustered environment.

When restoring a Lotus Domino database to a MCSC cluster and a failover occurs during the restore operation, active restore jobs are paused for 15 minutes as they wait for existing connections to resolve themselves. If the restore job does not restart before the failover time-out period expires, the job fails. If this occurs, the restore job must be resubmitted.

See [“About selecting Lotus Domino databases for restore”](#) on page 1079.


See [“Redirecting restore jobs for Lotus Domino databases”](#) on page 1083.

About selecting Lotus Domino databases for restore

When you view Lotus Domino databases in the restore selections pane, two backup sets appear for each Lotus Domino backup job. The first backup set contains the Lotus Domino databases and the second backup set contains the transaction logs.

To restore data, selections should be made from the backup set that contains the Lotus Domino databases; the required transaction logs are automatically restored with the selected database.

Figure F-2 Backup sets from a Lotus Domino backup job



Lotus Domino data is usually contained in the most recent backup set. However, some subsequent differential or incremental backup jobs run after a full backup job may not contain data in the backup set because only the transaction log was backed up. If the data you want to restore is not located in the most recent backup set, check the previous backup sets until you find the data.

Note: If a new DBIID has been assigned to databases and you run a differential or incremental backup, the data will be contained in the most recent backup set since transactions with the new DBIID will not match the old DBIID.

For example, the Domino server has a Full backup and a Differential backup. If you decide you want to restore data from Differential backup, you may select the Data1 directory and find that it is empty.

The following graphic shows an empty data directory.

Note: When restoring Lotus Domino databases to Microsoft Cluster Server cluster, the virtual computer name or the virtual IP address of the Domino server should be used when browsing or making Domino database selections in the **View by Resource** tab of the **Restore Job Properties** dialog box.

See [“Redirecting restore jobs for Lotus Domino databases”](#) on page 1083.

See [“Recovering a Lotus Domino server from a disaster”](#) on page 1087.

Selecting restore options for Lotus Domino databases

This procedure details how to select restore job properties for Lotus Domino databases, and provides definitions for Domino-specific restore options.

When you select a Lotus Domino backup set to restore, all database files and necessary transaction logs are automatically restored. You can also choose to restore specific database files.

See [“About Lotus Domino transaction logs”](#) on page 1074.

To select restore job properties for Lotus Domino databases

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Settings**, click **Lotus Domino**.
- 4 Select the appropriate options.

See [“Lotus Domino restore options”](#) on page 1082.

If your Lotus Domino database is replicated, the databases on each computer must have identical database and replica IDs. If you want to ensure that the databases continue to be replicated after the restore, select the Retain original IDs option.

- 5 Start the restore job or select other options from the **Properties** pane.

See [“Restoring data by setting job properties”](#) on page 605.

Lotus Domino restore options

You can set specific Lotus Domino-related restore options when you create a restore job.

See [“Selecting restore options for Lotus Domino databases”](#) on page 1082.

The Lotus Domino restore options are described in the following table:

Table F-4 Lotus Domino restore options

Item	Description
Seconds to wait for the database to go offline	Specifies the number of seconds for the restore process to wait for a database that is in use. When a Lotus database is restored it must first be taken offline. This will ensure that the database is not being accessed, closed, or deleted while the restore operation is being processed. If the database is still in use and cannot be taken offline after the specified wait time, the restore will fail.
Retain original IDs	Restores the original database IDs.
Assign new database ID	Assigns new IDs to the database.
Assign new database ID and replica ID	Assigns new IDs to the database. A replica ID is used to synchronize two or more databases that are being replicated in the Lotus Domino environment. You can assign a new replica ID during a restore to prevent other databases under replication from overwriting the restored database files.
Point in time restore	<p>Specifies the date and time to restore the database. The option is only available for logged databases when the archive logging style is set. Backup Exec will restore the Lotus Domino database you selected in the Restore selections dialog box and then automatically restore the necessary transaction logs required to bring the databases up to the date and time specified.</p> <p>If a point in time is not specified, the databases will be restored up to the last committed transactions in the log file.</p> <p>This option may require additional time since the archived transaction logs are also restored.</p>

Redirecting restore jobs for Lotus Domino databases

The Backup Exec logon account must have administrative credentials on the server to which you want to redirect the backup of the Lotus Domino server. Lotus Domino databases can only be redirected to a different directory on the local server from which the database was backed up. If you are restoring a database to a different location, it must reside in or under the Lotus Domino data directory. Point in time restores cannot be redirected.

Note: Redirecting the restore of a DAOS-enabled Domino database does not restore the nlo files.

See [“Restoring data by setting job properties”](#) on page 605.

See [“Creating a new Backup Exec System Logon Account”](#) on page 193.

To redirect the restore of a Lotus Domino database

- 1 Select the media that contains the data you want to restore.
- 2 On the navigation bar, click the arrow next to **Restore**.
- 3 Click **New Restore Job**.
- 4 Select the Lotus Domino databases.
- 5 After selecting options on the **Restore Job Properties** dialog box, on the **Properties** pane, under **Destination**, click **File Redirection**.
- 6 Select **Redirect file sets**.
- 7 Select the drive to which you are restoring in **Restore to drive**. You cannot enter the name of the drive, you must make a selection.
- 8 Enter the logon account for the server in **Server logon account**.
- 9 Enter the path to which you are restoring in **Restore to path**.
- 10 Enter the logon account for the path in **Path logon account**.
- 11 Start the redirection job or select other restore options from the **Properties** pane.

Redirecting the restore of DAOS NLO files

You can restore DAOS NLO files without restoring the entire DAOS-enabled Domino database. When you restore the DAOS NLO files you must specify a redirection destination path. In most cases, the path points to the DAOS internal repository that you set when you configured Lotus Domino.

See [“About the Lotus Domino Agent and the Domino Attachment and Object Service \(DAOS\)”](#) on page 1066.

To redirect the restore of DAOS NLO files

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the restore selections view, navigate to **Lotus Domino Databases**.
- 4 Expand **Lotus Domino Databases**.

5 Expand Domino Attachment and Object Service.

6 Select a backup set that contains the nlo files that you want to restore.

7 Do one of the following:

To restore all of the .nlo files in a folder Check the folder that contains all of the .nlo files that you want to restore.

To restore only the missing .nlo files Do the following in the order listed:

- Check the folder that contains all of the .nlo files that you want to restore.
- In the task pane, under **Settings**, click **General**.
- Check **Skip if file exists**.
When the restore job runs, Backup Exec restores only the missing .nlo files.

8 In the task pane, under Destination, click File Redirection.

9 Check Redirect file sets.

10 Click the ellipsis button next to the Restore to drive field.

11 Navigate to the DAOS folder using the path that you specified when you configured Domino.

If you changed the path after you backed up the Domino server, use the new path instead.

12 Click OK.

The correct path location entries should appear in the **Restore to drive** and the **Restore to path** fields.

13 Select other options, if appropriate.

See "[File Redirection restore options](#)" on page 634.

14 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the task pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
See “[Scheduling jobs](#)” on page 356.
- Click **Submit**.

- 15 When the redirection restore job completes, run a Domino DAOS catalog resynchronization operation at the Domino server.

For Domino DAOS catalog resynchronization information, see your Lotus Domino documentation.

How to prepare for disaster recovery on a Lotus Domino server

A Disaster Preparation Plan is necessary for restoring Lotus Domino databases efficiently and effectively in the event of a catastrophic failure. The goal is to minimize the time to recover. Developing a backup strategy for your Windows computers and Lotus Domino databases is the critical part of this plan.

When developing a strategy for backing up your Lotus Domino databases, consider the following recommendations:

- Keep linked databases on one volume. This allows Backup Exec to synchronize the databases before they are backed up.
- Back up active databases often. This reduces the amount of effort required to update the databases to the point following the most recent backup.
- Ensure that the notes.ini, cert.id, and <server>.id files are protected and available if a disaster occurs.
- Configure the DAOS prune period as recommended in the Lotus Domino documentation. However, Symantec recommends that you do not set the DAOS prune period for a period less than the time between two Domino backups.

See “[Recovering a Lotus Domino server from a disaster](#)” on page 1087.

See “[About disaster recovery of a Lotus Domino server using archive logging](#)” on page 1090.

See “[Recovering a Lotus Domino server that uses circular logging](#)” on page 1090.

Recovering a Lotus Domino server from a disaster

Lotus Domino system recovery can be performed in the following ways:

- **Manually**
See [“About manual disaster recovery of Windows computers”](#) on page 780.
- **By using Backup Exec’s Intelligent Disaster Recovery Options**
See [“About the the Intelligent Disaster Recovery Configuration Wizard”](#) on page 1810.

When you recover a DAOS-enabled Domino server from a disaster, all of the .nlo files that each recovered Domino database references are automatically restored.

Note: Disaster recovery of a Lotus Domino server in a Microsoft Cluster Server cluster uses the same steps as recovering a Domino server in a non-clustered environment.

Use the following steps as a guide when you want to do a disaster recovery operation on a Lotus Domino server.

Table F-5 Steps to take to recover a Lotus Domino server form a disaster

Step	Description
Step 1	Recover the Windows computer.
Step 2	Disable the monitor change journal. See “Disabling the monitor change journal” on page 1088.

Table F-5 Steps to take to recover a Lotus Domino server form a disaster
(continued)

Step	Description
Step 3	<p>Recover or re-install Lotus Domino to the same location as before the disaster occurred.</p> <p>All Lotus Domino system data must be recovered. System data includes log.nsf, names.nsf, template files, notes.ini, mail.box, and ID files.</p> <p>See “Recovering a Lotus Domino server and its databases” on page 1089.</p> <p>Note: If transaction logging is enabled, you must run a disaster recovery operation that is based on style of logging selected on the Lotus Domino server.</p> <p>See “Recovering a Lotus Domino server and its databases” on page 1089.</p> <p>After rebuilding the server, you can restore the databases from your most recent backup.</p>
Step 4	<p>Re-enable the monitor change journal.</p> <p>See “Re-enabling the monitor change journal” on page 1089.</p>

Disabling the monitor change journal

Use the following steps to disable and re-enable the monitor change journal in the registry. Then you can recover the Lotus Domino server, databases, and transaction logs.

See [“About disaster recovery of a Lotus Domino server using archive logging”](#) on page 1090.

To disable the monitor change journal

- 1 Open the registry and browse to the following key:
HKEY_LOCAL_MACHINE\SOFTWARE\SYMANTEC\BACKUP EXEC FOR WINDOWS\BACKUP EXEC\ENGINE\DOMINO
- 2 Disable monitor change journal by setting the value of the key Enable Change Journal to 0.
- 3 Restart the Remote Agent for Lotus Domino.

Recovering a Lotus Domino server and its databases

Use the following steps to recover a Lotus Domino server and its databases.

See [“Recovering a Lotus Domino server that uses circular logging”](#) on page 1090.

See [“Returning to the last known good configuration”](#) on page 777.

See [“Restoring data by setting job properties”](#) on page 605.

See [“Selecting restore options for Lotus Domino databases”](#) on page 1082.

To recover a Lotus Domino server and databases

- 1 Restore or re-install the Lotus Domino server program directory to the same location as before the disaster occurred.
- 2 Restore the notes.ini, cert.id, and <server>.id files from the last full backup of the Lotus Domino server program directory.
- 3 Use Backup Exec to restore the databases to the Domino data directory.
Backup Exec automatically restores all DAOS NLO files along with the DAOS-enabled databases. In addition, Domino automatically recreates both the daos.cfg file and the daoscat.nsf when you restart the Domino server.
- 4 Start the Lotus Domino Server.

Re-enabling the monitor change journal

Use the following steps to re-enable the monitor change journal.

See [“About disaster recovery of a Lotus Domino server using archive logging”](#) on page 1090.

See [“Selecting restore options for Lotus Domino databases”](#) on page 1082.

See [“Restoring data by setting job properties”](#) on page 605.

To re-enable monitor change journal

- 1 Open the registry and browse to the following key:
HKEY_LOCAL_MACHINE\SOFTWARE\SYMANTEC\BACKUP EXEC FOR WINDOWS\BACKUP EXEC\ENGINE\DOMINO
- 2 Enable monitor change journal by setting the value of the key Enable Change Journal to 1.
- 3 Restart the Remote Agent for Lotus Domino.

About disaster recovery of a Lotus Domino server using archive logging

If the active transaction log is lost, you can recover the database only up to the transactions contained in the last transaction log.

However, if all of the transaction logs are lost, you must have the following to recover the database:

- An up-to-date Notes.ini file from the Lotus Domino server.
- The backups of the database.
- All archived log extents.

In addition, if the monitor change journal is enabled, you must disable it in the registry before beginning the Lotus Domino server recovery.

See [“Disabling the monitor change journal”](#) on page 1088.

Recovering a Lotus Domino server that uses circular logging

If circular logging is enabled and the transaction log is lost, the Domino database can only be recovered to the point of the last backup.

Table F-6 The process to recover a Lotus Domino server that uses circular logging

Step	Description
Step 1	Restore or re-install the Lotus Domino server program directory (excluding the notes.ini, cert.id, and <server>.id files) to the same location as before the disaster occurred.
Step 2	Check that the log directory (logdir) is created and does not contain old files. If the log directory was not created, recreate the directory to the same location as before the disaster occurred. Do not start the Lotus Domino server after performing the previous steps.
Step 3	Restore the notes.ini, cert.id, and <server>.id files from the last full backup of the Lotus Domino server program directory to the same location as before the disaster occurred.

Table F-6 The process to recover a Lotus Domino server that uses circular logging (*continued*)

Step	Description
Step 4	To have Lotus Domino create the circular log file in the log directory when the server starts, set the following parameter in the notes.ini file: translog_path=logdir
Step 5	Use the Lotus Domino Agent to restore the databases to the Domino data directory. See “Restoring data by setting job properties” on page 605. See “Selecting restore options for Lotus Domino databases” on page 1082.

Recovering the Lotus Domino server, databases, and transaction logs when archive logging is enabled

Use the following steps to recover the Lotus Domino server, databases, and transaction logs.

See [“About disaster recovery of a Lotus Domino server using archive logging”](#) on page 1090.

To recover the Lotus Domino server, databases, and transaction logs when archive logging is enabled

- 1 Restore the non-database Domino server files (*.id and notes.ini).

If necessary, reinstall but do not configure the Domino server and then restore the non-database Domino files, which include the notes.ini and *.id files. Use the same directory structure, directory location, and logdir path as was created in the original installation. Do not launch the server after reinstalling it.
- 2 In the **Restore Job Properties** dialog box, under **Settings**, click **General**.
- 3 Check **Restore over existing files**.
- 4 Using a text editor, change the TRANSLOG_Status setting in the notes.ini file on the Domino server to 0.

For example, TRANSLOG_Status=0

- 5 Using the Backup Exec Agent for Lotus Domino, restore the last transaction log backed up prior to the loss of the active transaction log.
- 6 Verify the transaction log restore was successful.
- 7 Shutdown and then restart the Backup Exec Agent for Lotus Domino.
- 8 Delete all transaction logs except the transaction log restored in step 5, from the Domino transaction log directory.
- 9 Using a text editor, change the notes.ini file for the Domino server to match the following:

```
TRANSLOG_Recreate_Logctrl=1
```

```
TRANSLOG_Status=1
```

- 10 Run a full restore of the Domino databases or a point-in-time state within the archived log extents.

Backup Exec automatically restores all DAOS NLO files along with the DAOS-enabled databases. In addition, Domino automatically recreates both the `daos.cfg` file and the `daoscat.nsf` when you restart the Domino server.

After the full restore finishes, the `TRANSLOG_Logctrl` parameter in the `notes.ini` file is reset to 0.

- 11 Start the Domino server. Disaster recovery is complete.
- 12 If monitor change journal was disabled prior to beginning the disaster recovery process, you must re-enable it.

See [“Re-enabling the monitor change journal”](#) on page 1089.

Symantec Backup Exec Agent for Microsoft Exchange Server

This appendix includes the following topics:

- [About the Backup Exec Exchange Agent](#)
- [Requirements for using the Exchange Agent](#)
- [About installing the Exchange Agent](#)
- [Recommended configurations for Exchange](#)
- [Requirements for accessing Exchange mailboxes](#)
- [Backup strategies for Exchange](#)
- [How Granular Recovery Technology works with the Exchange Information Store](#)
- [Snapshot and offhost backups with the Exchange Agent](#)
- [About continuous protection for Exchange data](#)
- [Best practices for continuous protection of Exchange](#)
- [Setting default backup and restore options for Exchange data](#)
- [About backing up Exchange 2003/2007](#)
- [About backing up Exchange Server 2010 Databases](#)
- [Backing up Exchange](#)

- [About restoring Exchange data](#)
- [About redirecting Exchange restore data](#)
- [How to prepare for disaster recovery of Exchange Server](#)
- [Recovering from a disaster for Exchange 2000 or later](#)

About the Backup Exec Exchange Agent

The Exchange Agent lets you integrate backups of Microsoft Exchange Server databases with network backups without separate administration or dedicated hardware.

The Exchange Agent provides the following features:

- The ability to restore individual items from backups for which you enable Granular Recovery Technology.
- Continuous backup of the Exchange Server when Backup Exec Continuous Protection Server (CPS) is installed. The CPS Exchange backup job provides complete recovery to any point in time of the Information Store, including the latest complete transaction log. When you enable recovery points to run at intervals between the full backups, you can restore individual items at a point in time when the recovery point was created. Even without recovery points, you can restore individual items from a full backup.
- The ability to select storage groups for backup and restore jobs, or to select one or more databases within the storage group for backup and restore jobs.
- The ability to restore individual databases or storage groups from non-snapshot backups by using the Recovery Storage Group feature in Exchange Server 2003 and Recovery Database feature in Exchange Server 2010. For Exchange Server 2007/2010, you can restore snapshot backups to a recovery storage group or database.
- Seeding of an Exchange Server 2010 database copy. Seeding adds a database copy to a location on another mailbox server in a Database Availability Group.
- Snapshot backup and offhost backup on Exchange Server 2003 or Exchange Server 2007 instances that run on Windows Server 2003.
- Offhost backup with Granular Recovery Technology (GRT) for Exchange Server 2003/2007/2010

See [“About installing the Exchange Agent”](#) on page 1099.

See [“Backup strategies for Exchange”](#) on page 1102.

See [“Recommended configurations for Exchange ”](#) on page 1100.

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“About offhost backup”](#) on page 919.

See [“About the Advanced Open File Option”](#) on page 937.

Requirements for using the Exchange Agent

The media server must meet the following requirements:

Table G-1 Media server requirements for the Backup Exec Exchange Agent

Media server requirements	Description
To support the Exchange Agent	<ul style="list-style-type: none">■ A license key for Symantec Backup Exec Microsoft Exchange Server Agent (Exchange Agent) must be entered on the media server.■ The Backup Exec Remote Agent for Windows Systems (Remote Agent) must be installed on any remote Exchange Server that you want to back up.■ The media server must have access to the Exchange Server. Note: To back up Exchange Server 2010, you must install Backup Exec on a Microsoft Windows 2008 SP2 64-bit server or a Microsoft Windows 2008 R2 64-bit server.■ Symantec recommends that you use a Backup Exec services account that has domain and local administrator rights on the Exchange Server.

Table G-1 Media server requirements for the Backup Exec Exchange Agent
(continued)

Media server requirements	Description
To back up Exchange Server 2007/2010	<p>To back up Microsoft Exchange Server 2007/2010, you must install the Exchange Management Tools for Microsoft Exchange Server 2007/2010 on the media server. The management tools on the media server must be the same version or later as the management tools that are on the Exchange Server 2007/2010.</p> <p>You can install the management tools when you do a custom install of Microsoft Exchange Server 2007/2010. If you install the management tools and Backup Exec together on a media server, install the tools first. If you install Backup Exec before you install the management tools, you must restart the media server when the tools installation is complete.</p>
To support the Backup Exec Resource Discovery feature, which allows detection of new backup resources within a Windows domain	<p>For Exchange 2003, Microsoft Exchange System Manager utility must be installed.</p> <p>For Exchange 2007/2010, Exchange Management Tools must be installed. You can install both versions of the Exchange Management Tools on the media server.</p>
To back up Exchange data from any node of a Veritas Cluster Server	Microsoft Exchange System Manager utility must be installed on all nodes.

Table G-1 Media server requirements for the Backup Exec Exchange Agent
(continued)

Media server requirements	Description
To support Granular Recovery Technology (GRT) for the restore of individual items from Information Store backups	<p>One of the following versions of a Microsoft Windows operating system that supports minifilter drivers must be installed for Microsoft Exchange:</p> <ul style="list-style-type: none"> ■ Microsoft Windows 2000 Server (with Service Pack 4 and Update Rollup 1 for Service Pack 4) ■ Microsoft Windows Server 2003 (with at least Service Pack 1) ■ Microsoft Windows Server 2003 R2 Editions ■ Microsoft Windows Server 2008 SP2 ■ Microsoft Windows Server 2008 R2 Editions <p>Note: For Exchange Server 2010, you must use either Microsoft Windows 2008 SP2 or a Microsoft Windows Server 2008 R2.</p> <p>Devices that you use for GRT-enabled backups may have additional requirements.</p> <p>See “Recommended devices for backups that use Granular Recovery Technology” on page 324.</p> <p>See “About requirements for jobs that use Granular Recovery Technology” on page 325.</p>
To support Backup Exec Continuous Protection Server	<p>CPS components must be installed. For information on how to install the CPS components, see the <i>Symantec Backup Exec Continuous Protection Server Administrator's Guide</i>.</p> <p>See “Requirements for installing components for CPS Exchange backup jobs” on page 1113.</p>

The following are requirements for the Exchange Server with the Backup Exec Exchange Agent:

Table G-2 Exchange Server requirements

Exchange Server requirements	Description
To support Exchange Server 2007	<p>Download the Microsoft Exchange Server MAPI Client and Collaboration Data Objects package and install it on Exchange Server 2007.</p> <p>This package provides support for the following:</p> <ul style="list-style-type: none"> ■ The restore of individual mailboxes, mail messages, and public folders from an Information Store backup. ■ The collection of catalog information for a backup for which the Granular Recovery Technology option is enabled and the destination device is tape. <p>You can find this package on the Microsoft Web site.</p>
For operations on all Exchange resources	<p>The user account must be a member of the following groups:</p> <ul style="list-style-type: none"> ■ The Administrators group ■ The Domain Admins <p>You must also use the appropriate Exchange server management utility to assign the user account the Exchange Organization Administrators role (2007) or the Exchange Organization Management role (2010).</p>
To support the Granular Recovery Technology option for Exchange Server 2007	<p>You must use the appropriate Exchange server management utility to assign the user account the Exchange Organization Administrators role (2007) or the Exchange Organization Management role (2010).</p>
To support snapshot backups	<p>Use Microsoft Exchange Server that runs on Windows Server 2003 or later.</p> <p>Note: To select incremental or differential backup methods, Exchange Server 2003 Service Pack 1 or later must be installed.</p>

Table G-2 Exchange Server requirements (*continued*)

Exchange Server requirements	Description
To support Backup Exec Continuous Protection Server	<p>CPS components must be installed. For information on how to install the CPS components, see the <i>Symantec Backup Exec Continuous Protection Server Administrator's Guide</i>.</p> <p>See "Requirements for installing components for CPS Exchange backup jobs" on page 1113.</p>
To back up and restore Exchange Server 2010	<p>To back up the databases on a Database Availability Group (DAG) you must install the Remote Agent for Windows Systems on all the servers in the DAG.</p> <p>To support the Granular Recovery Technology option, you must install Remote Agent for Windows Systems on all Client Access Servers in the site.</p> <p>See "About the Remote Agent for Windows Systems" on page 1945.</p>

Backup Exec does not support the Granular Recovery Technology option when Outlook is installed on the same computer with any of the following:

- Exchange Server 2003
- Exchange Server 2007 on a version of Windows earlier than 2003

See the Microsoft Knowledge Base for information about installing Outlook and Exchange Server on the same computer.

See ["Using resource discovery to search for new resources"](#) on page 316.

See ["About selecting individual Exchange mailboxes for backup"](#) on page 1140.

See ["About continuous protection for Exchange data"](#) on page 1112.

About installing the Exchange Agent

The Exchange Agent is installed locally as a separate, add-on component of Backup Exec to protect local or remote Exchange Server databases.

To protect Exchange Server 2010, you must install Backup Exec on a Microsoft Windows 2008 SP2 64-bit media server or a Microsoft Windows 2008 R2 64-bit media server.

Note: When you install Microsoft Exchange Tools 2007/2010 and Backup Exec together on a media server, Exchange Tools 2007/2010 must be installed first. If you install Backup Exec before Exchange Tools, you must restart the media server after you finish the Exchange Tools installation.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

Recommended configurations for Exchange

Before starting backups for Exchange, read the following recommendations for configuring Exchange to make it easier to restore from backups:

Table G-3 Recommended configurations for Exchange

Recommendation	Description
Put transaction log files on a separate physical disk from the database.	This is the single most important configuration affecting the performance of Exchange. This configuration also has recovery implications, since transaction logs provide an additional recovery resource.
Make Write Cache unavailable on the SCSI controller.	The Windows operating system does not use buffers, so when Exchange receives a write complete notice from Windows, the write-to-disk has been completed. If Write Cache is enabled, Windows responds as though a write-to-disk has been completed, and will provide this information to Exchange (or other applications) incorrectly. The result could be data corruption if there is a system crash before the operation is actually written to disk.
Make circular logging unavailable if possible.	Circular logging minimizes the risk that the hard disk will be filled with transaction log files. But, if a solid backup strategy is in place, transaction log files are purged during the backup, thus freeing disk space. If circular logging is enabled, transaction log histories are overwritten, incremental and differential backups of storage groups and databases are disabled, and recovery is only possible up to the point of the last full or copy backup. Note: Continuous backup of Information Store transaction logs with the Backup Exec Continuous Protection Server is not supported if circular logging is enabled.

Table G-3 Recommended configurations for Exchange (*continued*)

Recommendation	Description
Avoid making the Exchange Server a domain controller.	For disaster recovery purposes, it is much easier to restore Exchange if you don't have to restore the Active Directory first.
Install Exchange into a domain that has at least two domain controllers.	Active Directory replication is not possible with only one domain controller in a domain. If the domain controller fails and corrupts the Active Directory, some transactions may not be recoverable if they were not included with the last backup. With at least two domain controllers in a domain, databases on the failed domain controller can be updated using replication to fill in missing transactions after the database backups have been restored.

See [“About the circular logging setting for Exchange”](#) on page 1106.

See [“Requirements for accessing Exchange mailboxes ”](#) on page 1101.

Requirements for accessing Exchange mailboxes

Backup Exec must have access to a uniquely named mailbox within the Exchange organization for Information Store operations, depending on how the backup and restore jobs are configured.

Access to a uniquely named mailbox is required when you do the following:

- Back up individual mailboxes separately from the Information Store (also called the legacy mailbox backup method).
- Configure a backup job that has all of the following settings:
 - A device other than a backup-to-disk folder is the destination device.
 - The Granular Recovery Technology option is enabled.
 - A backup method other than a snapshot method.
- You restore mailboxes and public folders.

You must use a Backup Exec logon account to connect to the Exchange server when you select mailboxes or public folders for backup. Backup Exec attempts to find a mailbox with the same name as the user name that is stored in the Backup Exec logon account.

If you use a Backup Exec logon account that stores a unique user name and has a corresponding mailbox with the same name, then you are not prompted for an additional logon account. Otherwise, you must choose or create a Backup Exec

logon account that stores the name of a unique mailbox within the Exchange organization.

A unique name does not share the first five characters in another mailbox name. For example, if EXCH1 is entered as the mailbox name, and there is another mailbox name such as EXCH1BACKUP, then Backup Exec cannot accept the name. You are prompted to choose another mailbox name.

You can choose or create a logon account that meets any of the following requirements:

- A logon account for which the user name matches a unique mailbox name.
- A logon account that uses a unique alias to a mailbox. The user account that connects to the Exchange Server must also have access to this mailbox.
- A logon account that uses the full computer name for a mailbox. The user account that connects to the Exchange server must also have access to this mailbox.

An example of a full computer name is:

`/O=Exchange_Organization/OU=Administrative_Group/CN=Recipients/CN=mailbox_name`

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“Creating a Backup Exec logon account”](#) on page 187.

Backup strategies for Exchange

Backup Exec incorporates online, nondisruptive Exchange database protection as part of everyday backup routines, which increases the chance of data recovery and minimizes data loss without inhibiting daily activity. Backup Exec protects Exchange data, including the individual storage group, database, mailbox, and public folder with full, copy, incremental, and differential backups.

To decide which backup methods to use, consider the following:

- In small office environments with relatively small numbers of messages passing through the system, a daily full backup will provide good data protection and the quickest recovery. If log file growth becomes an issue, consider using incremental online backups at midday to provide an added recovery point and manage the log file growth for you automatically.
- In large environments, incremental backups should be used to provide more frequent recovery point options throughout the day and to manage log file growth. Many shops run full backups on a weekly basis, preferring to run incremental backups throughout the week to keep backup run time to a

minimum. The trade-off with this technique occurs at recovery time when you must recover from the full backup and from each incremental backup as well.

What works best for you is based on the size of your environment, the number of transactions processed each day, and the expectations of your users when a recovery is required.

Consider the following backup strategies:

- Run Backup Exec Continuous Protection Server (CPS) jobs weekly or daily. The full backups and the replicated transaction logs provide complete recovery to any point in time of the Information Store, including the latest complete transaction log. You can also restore individual messages or folders from the CPS backup.

When you enable recovery points to run at intervals between the full backups, you can restore individual messages or folders at a point in time when the recovery point was created. Another advantage of recovery points is that log growth is controlled because the transaction logs are truncated after each recovery point runs.

Note: You cannot use CPS on an Exchange server that is in an Exchange Server 2010 Database Availability Group (DAG).

- Run full backups with the option to enable the restore of individual items selected so that you can restore individual mail messages and folders without restoring the entire database.

Depending on your environment, run full backups as follows:

- As frequently as possible, no less than once a day.
 - Daily with differential backups used at regular periods throughout the day.
 - Every few days (no less than weekly) with frequent incremental backups in between each full backup.
- Run Exchange backup jobs separately from other backup jobs.

In addition to backing up Exchange storage groups or databases, you should also back up the following on a regular basis:

Table G-4 Backup selections for Exchange configuration data

Recommended backup selections for configuration data	Description
File system	<p>Back up folders and drives containing files for Windows and Exchange. Usually, this is the root drive C:\ but may be different in each environment.</p> <p>Note: Back up the C:\ drive, but do not back up the virtual drive that is created by Exchange, if this virtual drive exists in your environment. It is intended only to provide Explorer access to the Exchange data, but all file system functions may not be replicated. Backup and restore operations are not recommended or supported.</p>
Windows registry	<p>Back up the registry by running a full backup.</p>
System State and/or Shadow Copy Components	<p>Select System State and run a full backup to back up the following:</p> <ul style="list-style-type: none"> ■ The Internet Information Service (IIS) metabase ■ The Windows registry <p>See “About selecting data to back up” on page 280.</p> <p>If the entire server must be restored, you must restore System State before restoring Exchange 2000. You must also restore both System State and Shadow Copy Components before restoring Exchange Server 2003/2007/2010.</p>

Table G-4 Backup selections for Exchange configuration data (*continued*)

Recommended backup selections for configuration data	Description
Active Directory	<p>To back up Active Directory, select System State on the domain controllers and run a full backup.</p> <p>When there are configuration changes on the Exchange server database, such as when objects are added, modified, or deleted, back up the Active Directory on the domain controllers.</p> <p>Note: Spread multiple domain controllers throughout each domain for efficient Active Directory replication, and so that if one domain controller fails, redundancy is still provided.</p>

Note: Configure an Information Store backup for which the Granular Recovery Technology (GRT) option is enabled to restore individual mailboxes, mail messages, and public folders. Backing up individual Exchange mailboxes separately from the Information Store uses legacy backup methods, and is no longer required for individual mailbox recovery.

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“About backing up Exchange 2003/2007”](#) on page 1129.

See [“About backing up Exchange Server 2010 Databases”](#) on page 1130.

See [“How to prepare for disaster recovery of Exchange Server”](#) on page 1165.

Automatic exclusion of Exchange data during volume-level backups

If you select a volume that contains Exchange data for backup, the Exchange Agent uses Active File Exclusion to automatically exclude Exchange data that should not be included in a volume-level backup. For example, .EDB and .STM files, as well as transaction log files, should not be part of a volume-level backup because they are opened for exclusive use by Exchange.

Without this exclusion, during a non-snapshot backup, these files appear as in use - skipped. During a snapshot backup, these files may be backed up in an inconsistent state, which could create restore issues.

While it is not recommended, if you want to include Exchange data in a volume-level backup, you must first dismount the storage groups or databases that you want backed up, and then run the backup job.

About the circular logging setting for Exchange

When circular logging is enabled, you cannot run incremental and differential backups of Exchange databases and storage groups, or run backup jobs for which continuous protection is enabled. These types of backups rely on a complete history of logs.

When circular logging is enabled, transaction log files that have already been committed to the database are overwritten, preventing the accumulation of logs. The log files are overwritten whether or not a full or incremental backup has been run, and a history of previous logs since the last full or incremental backup is not maintained.

When circular logging is disabled, transaction log files accumulate on the disk until the following occurs:

- A full or incremental backup is performed.
- A recovery point is run as part of a continuous backup of Exchange.

After these operations, the log files that have all transactions committed to the database are deleted.

See [“Backup strategies for Exchange”](#) on page 1102.

How Granular Recovery Technology works with the Exchange Information Store

Backup Exec Granular Recovery Technology (GRT) lets you restore individual items from an Information Store backup without having to restore the whole backup. You should review the requirements for a GRT-enabled backup before you configure it.

You can also enable GRT when you create an offhost backup for the Information Store. Offhost backup lets Backup Exec move the backup process from the host computer to the Backup Exec media server. The host computer is the remote computer that contains the volumes that you selected for backup. To run a GRT-enabled offhost backup, you must install the Backup Exec Advanced Disk-based Option on the media server.

Note: In previous versions of Backup Exec, individual Exchange mailboxes were backed up separately from the Information Store so that individual mailboxes could be restored. These legacy backup options are enabled by default only if you upgrade from a previous Backup Exec version and jobs for mailbox backups already exist.

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“Recommended devices for backups that use Granular Recovery Technology”](#) on page 324.

See [“About requirements for jobs that use Granular Recovery Technology”](#) on page 325.

See [“Configuring a GRT-enabled offhost backup for Exchange resources”](#) on page 928.

See [“About installing the Advanced Disk-based Backup Option”](#) on page 898.

About Backup Exec and Microsoft Exchange Web Services

Backup Exec uses Microsoft Exchange Web Services (EWS) to support the Granular Recovery Technology option. EWS provides support for the restore of individual mailboxes, mail messages, and public folders from an Exchange Server 2010 database backup.

Note: You do not need to install the MAPI Client and Collaboration Data Objects package if you use EWS.

To use EWS to restore individual items, Backup Exec disables the client throttling policy for the resource credentials you specify for the restore job. The client throttling policy is located on the Client Access Server and enforces connection bandwidth limits on the Exchange server.

Backup Exec also creates an impersonation role and a role assignment for Exchange Impersonation. Exchange Impersonation role assignment associates the impersonation role with the Backup Exec resource credentials you specify for the restore job.

Backup Exec creates and assigns the following roles:

- SymantecEWSImpersonationRole
- SymantecEWSImpersonationRoleAssignment

Snapshot and offhost backups with the Exchange Agent

The Exchange Agent supports the Microsoft Volume Shadow Copy Service (VSS), a snapshot provider service only available on Windows Server 2003 or later. Using VSS, a point in time view of the Exchange database is snapped and then backed up, leaving the actual Exchange database open and available for users.

Offhost backup enables the backup operation to be processed on a Backup Exec media server instead of on the Exchange server. Moving the backup from the Exchange server to a media server enables better backup performance and frees the remote computer as well.

If the Advanced Disk-based Backup Option (ADBO) is installed on the media server, you can use the Backup Exec Granular Recovery Technology (GRT) option when you create an offhost backup for the Information Store.

See [“Configuring a GRT-enabled offhost backup for Exchange resources”](#) on page 928.

The Exchange Agent snapshot does not support the following:

- NAS configurations
- The Exchange 2003 Recovery Storage Group feature
- Mixing snapshot backups and non-snapshot backups
Due to a Microsoft Exchange limitation, if non-snapshot backups are run as part of a data protection scheme, then snapshot backups should not be run. If snapshot backups are run, non-snapshot backups should not be done.

The type of backup method that is available when using VSS with the Exchange Agent depends on the version of Exchange Server, and are listed in the following table:

Table G-5 Available backup methods for Exchange snapshot versions

Exchange version	Available backup methods
Exchange Server 2003	The following backup methods are available: <ul style="list-style-type: none">■ Full■ Copy

Table G-5 Available backup methods for Exchange snapshot versions
(continued)

Exchange version	Available backup methods
Exchange Server 2003 with Service Pack 1 or later Exchange Server 2007/2010	The following backup methods are available: <ul style="list-style-type: none"> ■ Full ■ Copy ■ Differential ■ Incremental storage group level snapshot backup ■ Individual database restore
Exchange Server 2007	LCR/CCR - Back up from the passive copy or the active copy. Note: You cannot back up the passive copy of the Standby Continuous Replication (SCR) database with Exchange Server 2007. The SCR is not available for backup selection.

Troubleshooting Exchange Agent snapshot and offhost jobs

An Exchange Agent snapshot job fails on the following conditions:

- The Exchange Agent snapshot fails.
- You run a migrated or new snapshot backup for data on Exchange Server 2003 on Windows 2000. Snapshot backups of Exchange are only supported for Exchange Server 2003 or later on Windows Server 2003 or later. The job does not fall back to a non-snapshot backup because Exchange snapshot and non-snapshot backups are not interoperable.
The snapshot backup continues for resources that are supported, and the job can successfully complete with exceptions.
To allow the snapshot to continue for the supported resources, do one of the following:
 - Check Process logical volumes for backup one at a time on the Backup Job Properties for the Advanced Open File Option.
 - Check Process logical volumes for offhost backup one at a time on the Backup Job Properties for the Advanced Disk-based Backup Option.
- If incremental or differential backup methods are selected, and Exchange Server 2003 Service Pack 1 or later is not installed.
- If circular logging is enabled, and incremental or differential backup methods are selected.

- You run a snapshot job on Windows Small Business Server 2003. The Microsoft Exchange Server 2003 VSS Writer is disabled on Windows Small Business Server 2003, which causes snapshot backups for Exchange 2003 to fail. To successfully perform an Exchange 2003 snapshot backup, review the following Microsoft Knowledge Base article:
<http://support.microsoft.com/default.aspx?scid=kb;EN-US;Q838183>
 You must resolve this issue to perform a successful restore job by using the Intelligent Disaster Recovery option.

Configuring a snapshot backup for Exchange resources

Symantec recommends that you perform consistency checks before running a snapshot backup.

See [“About backing up Exchange 2003/2007”](#) on page 1129.

Table G-6 Configuring a snapshot backup for Exchange resources

Step	Action
Step 1	Create an Exchange backup job. See “About backing up Exchange 2003/2007” on page 1129.
Step 2	Set default options for the Advanced Open File Option. For Exchange Server 2007/2010 resources, Backup Exec automatically performs snapshot backups. You do not need to select options for the Advanced Open File Option. See “Setting default options for the Advanced Open File Option” on page 943. Ensure that either the Microsoft Volume Shadow Copy Service option or the Automatically select open file technology option is selected. If resources that are not supported for snapshot backup are included in the backup selection list, check Process logical volumes for backup one at a time to allow the job to complete with errors.

Table G-6 Configuring a snapshot backup for Exchange resources (*continued*)

Step	Action
Step 3	<p>Schedule or start the backup job.</p> <p>See “Creating a backup job by setting job properties” on page 332.</p> <p>See “Snapshot and offhost backups with the Exchange Agent” on page 1108.</p> <p>See “Troubleshooting Exchange Agent snapshot and offhost jobs” on page 1109.</p> <p>See “About restoring Exchange data from snapshot backups” on page 1149.</p>

Configuring an offhost backup with the Exchange Agent

Symantec recommends that you perform consistency checks before running an offhost backup.

See [“About backing up Exchange 2003/2007”](#) on page 1129.

If the Advanced Disk-based Backup Option (ADBO) is installed on the media server, you can use the Backup Exec Granular Recovery Technology (GRT) option when you create an offhost backup for the Information Store. You can restore individual mailboxes, mail messages, and public folders from Information Store backups for which GRT is enabled.

See [“Configuring a GRT-enabled offhost backup for Exchange resources”](#) on page 928.

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 1108.

See [“Troubleshooting Exchange Agent snapshot and offhost jobs”](#) on page 1109.

To configure an offhost backup with the Exchange Agent

- 1 Create an Exchange backup job.
See [“About backing up Exchange 2003/2007”](#) on page 1129.
- 2 In the **Properties** pane, under **Settings**, click **Advanced Disk-based Option**.
See [“Setting offhost backup options for a backup job”](#) on page 925.
- 3 Select the appropriate options, and then click **OK**.

- 4 Ensure that the snapshot provider that you select is **Automatic - Use hardware if available; otherwise use software**.
- 5 The options for job disposition are not available.
- 6 If resources that are not supported for offhost backup are included in the backup selection list, check **Process logical volumes for offhost backup one at a time** to allow the job to complete with errors.
- 7 Schedule or start the backup job.

About continuous protection for Exchange data

The Symantec Backup Exec Continuous Protection Server (CPS) combines data protection with replication technology and disk-based data protection. When CPS components are installed on the media server and on the Exchange server, you can continuously protect Exchange data.

When you enable continuous protection, Backup Exec protects the Information Store with a recurring full backup, which is called a CPS Exchange job. The CPS Exchange job is sent to a backup-to-disk folder on a local NTFS volume. You can also enable recovery points that run in between the full backups. Transaction logs that are created after the full backups are continuously replicated to a backup-to-disk folder.

The full backups and the replicated transaction logs provide recovery to any point in time of the Information Store, including the latest complete transaction log. The recovery points let you restore individual messages or folders at a point in time when the recovery point was created. Even without recovery points, you can restore individual messages or folders from a full backup.

The continuous protection feature does not support the following:

- The Microsoft Volume Shadow Copy Service (VSS) snapshot provider. Snapshot options that you select on the Advanced Open File Option backup job properties are ignored for CPS Exchange backup jobs. CPS Exchange backup jobs for Exchange Server 2007/2010 are always run as snapshot backups.

Note: CPS Exchange backup jobs for Exchange 2003 are always run as traditional streaming backups.

- For Exchange Server 2003 resources, the delegation of continuous protection-related jobs to managed media servers in a Central Admin Server Option (CASO) environment. Job delegation is supported for Exchange Server 2007 resources.

- Circular logging, if recovery points are used. Recovery points fail if circular logging is enabled.
- Clusters.
- Database Availability Groups.

See [“Requirements for using the Exchange Agent ”](#) on page 1095.

Requirements for installing components for CPS Exchange backup jobs

To use Backup Exec Continuous Protection Server (CPS) for continuous protection of Exchange data, you must install CPS components.

For information on how to install the Continuous Protection Server components, see the *Symantec Backup Exec Continuous Protection Server Administrator's Guide*.

Note: You cannot use CPS on an Exchange server that is in an Exchange Server 2010 Database Availability Group (DAG).

The following table lists requirements for installing the necessary components to create CPS Exchange backup jobs:

Table G-7 Requirements for installing components for CPS Exchange backup jobs

Component	Requirements
<p>If you install the Continuous Management Service (CMS) on the same server that hosts the Backup Exec media server</p>	<p>Do the following when you install CPS:</p> <ul style="list-style-type: none"> ■ Ensure that the path for the journal files is on a different drive than the CPS Exchange backup-to-disk folder. When you select the path for the journal files, review the current disk space availability on the server. Select a drive with enough space for the journal files. For information about required space for journal files, see the <i>Symantec Backup Exec Continuous Protection Server Administrator's Guide</i>. <p>Note: When you install CPS, you must restart the server.</p> <ul style="list-style-type: none"> ■ On the Push Install CPS Components Wizard panel, ensure that the Exchange Protection Agent is selected. If the Exchange Protection Agent is not available on the Push Install CPS Components Wizard panel, then check for DNS errors. <p>Do the following after you install CPS:</p> <ul style="list-style-type: none"> ■ Use Symantec LiveUpdate to update the server. For information on how to use LiveUpdate with CPS, see the <i>Symantec Backup Exec Continuous Protection Server Administrator's Guide</i>. ■ View the services to ensure that the CPS Exchange Agent is installed. If the Backup Exec Continuous Protection Broker Service is present, the CPS Exchange Agent is installed. ■ Push-install the CPS Continuous Protection Agent to the Exchange Server.
<p>Exchange Server</p>	<p>The following are required for the Exchange Server:</p> <ul style="list-style-type: none"> ■ It must be on a different server than the media server and the server on which CMS is installed. ■ It must be in the same domain as the media server unless the domains are trusted. ■ It must have the Backup Exec Remote Agent for Windows Systems installed. You can install the Remote Agent from either the media server or from the server on which CMS is installed.

Table G-7 Requirements for installing components for CPS Exchange backup jobs (continued)

Component	Requirements
If you install the Continuous Management Service (CMS) on a server that does not host the Backup Exec media server	Do the following: <ul style="list-style-type: none"> ■ You must push-install the CPS Continuous Protection Agent to the Backup Exec media server and to the Exchange Server. ■ Optionally, to view the CPS console from the media server, you can push-install the CPS Administration Console to the Backup Exec media server.
Backup Exec media server	The following Windows operating systems are supported: <ul style="list-style-type: none"> ■ Microsoft Windows 2000 Server (with at least Service Pack 4 and Update Rollup 1 for Service Pack 4) ■ Microsoft Windows Server 2003 (with at least Service Pack 1) ■ Microsoft Windows Server 2003 R2 Editions ■ Microsoft Windows Server 2008 The media server must have the following: <ul style="list-style-type: none"> ■ Minimum of 1 GB RAM ■ A processor with a minimum of 2 GHz
Backup Exec service account	The Continuous Protection Agent and the Exchange Protection Agent must use the Backup Exec service account. Ensure that the Backup Exec service account has the following: <ul style="list-style-type: none"> ■ Domain and local administrator rights. ■ Ability to query the location of the transaction logs on the local active Exchange server.

See [“Best practices for continuous protection of Exchange”](#) on page 1117.

See [“About managing the CPS Exchange backup job for Exchange data”](#) on page 1118.

Requirements for configuring continuous protection for Exchange data

To configure continuous protection jobs for the Exchange Information Store, the following are required:

Table G-8 Requirements for configuring continuous protection for Exchange

Requirement	Description
Install the Backup Exec Continuous Protection Server (CPS) components.	<p>For information on how to install the Continuous Protection Server components, see the <i>Symantec Backup Exec Continuous Protection Server Administrator's Guide</i>.</p> <p>See “Requirements for installing components for CPS Exchange backup jobs” on page 1113.</p>
Specify a recurring schedule for the full backup of the Information Store as part of the CPS job.	<p>Recovery points are dependent on the last full backup. If the full backup is not run often enough, or if it is not available, then the subsequent recovery points are not usable.</p> <p>If there is not enough disk space to maintain the full backup and subsequent recovery points, then consider one or both of the following actions:</p> <ul style="list-style-type: none"> ■ Schedule the full backup more often. ■ Schedule the recovery points less often. <p>You must balance the frequency of the full backup and the frequency of the recovery points to achieve efficient use of the available disk space.</p> <p>See “Scheduling jobs” on page 356.</p> <p>See “About using recovery points to restore individual Exchange items to a point in time” on page 1122.</p>
Ensure that the full backup job for the Information Store is not in a policy.	<p>You can create a backup job by setting the properties that you want. If you are new to Backup Exec or uncertain about how to set up a backup job, use the Backup Wizard.</p> <p>See “Creating a backup job by setting job properties” on page 332.</p>
Ensure that the Exchange mailbox stores are not included in any other backup job.	<p>If the Exchange mailbox store is backed up by a CPS Exchange backup job, do not include it in other backup selection lists.</p>

Table G-8 Requirements for configuring continuous protection for Exchange (continued)

Requirement	Description
Select a backup-to-disk folder as the destination device for the CPS Exchange backup job.	<p>Configure the backup-to-disk folder as follows:</p> <ul style="list-style-type: none"> ■ It must reside on an NTFS volume on the local media server. ■ It cannot be a removable backup-to-disk folder. ■ The option Allocate the maximum size for backup-to-disk files must not be selected. <p>See “Application settings for new backup-to-disk folders” on page 499.</p> <p>Note: You must select a specific backup-to-disk folder. If you select a device pool, even if it contains a backup-to-disk folder, the job fails.</p> <p>See “About backup-to-disk folders ” on page 496.</p> <p>Use the backup-to-disk folder exclusively for the CPS Exchange backup job. Do not back up other resources to the backup-to-disk folder that is the destination device for the CPS Exchange backup job.</p> <p>See “About reviewing disk space availability for CPS Exchange backup jobs” on page 1119.</p>
Make circular logging unavailable.	<p>If circular logging is enabled, the recurring full backup of the Information Store completes without errors, but recovery points fail.</p> <p>See “About the circular logging setting for Exchange” on page 1106.</p>
Select the Exchange Server name in the backup selections list.	<p>You cannot select an IP address for an Exchange server in the backup selections list on the media server.</p>

See [“About reviewing disk space availability for CPS Exchange backup jobs”](#) on page 1119.

See [“Best practices for continuous protection of Exchange”](#) on page 1117.

Best practices for continuous protection of Exchange

When you use continuous protection as part of your backup strategy, note the following best practices:

- Symantec recommends that you back up only one Exchange server for each continuous backup job. Create a separate selection list for each Exchange server resource.
- If you must copy backup sets to tape for off-site storage, create a job to duplicate backup sets. You can configure the job to copy the backup sets to tape after each occurrence of the full backup job.
If necessary, you can create a copy job to run before the full backup. This copies all of the transaction logs, as well as the full backup sets, to tape.
- If you duplicate Information Store backup sets to tape, and then back to disk, specify the same volume for the full and the incremental backups. The backup sets must be on the same volume to restore individual items from the incremental backup.
- You can create a custom filter to limit the display of recovery points in the Job History view.
- After you create and run a CPS Exchange backup job, do not change the backup-to-disk folder that it was run to. If you must change backup-to-disk folders, then create a new CPS Exchange backup job with a new backup-to-disk folder as the destination device. Delete the previous job.

See [“Creating selection lists”](#) on page 296.

See [“Adding a duplicate backup template to a policy”](#) on page 550.

See [“About managing custom filters”](#) on page 582.

See [“About backing up Exchange 2003/2007”](#) on page 1129.

See [“About reviewing disk space availability for CPS Exchange backup jobs”](#) on page 1119.

See [“Troubleshooting CPS Exchange backup jobs”](#) on page 1123.

About managing the CPS Exchange backup job for Exchange data

All backup operations that are related to the continuous protection of the Exchange Server are handled as a single job. This job is displayed in the Current Jobs view on the Job Monitor. The status of this job changes according to the operation that is running.

To view the continuous protection job for transaction log replication, or to view related errors, you must go to the CPS Administration Console. If the Continuous Protection Server Administration Console component is installed on the media server, you can view the CPS console.

The statuses for the different operations are listed in the following table:

Table G-9 Continuous protection job statuses

Continuous protection operation	Status on the Job Monitor in the Current Jobs view
When the recurring full backup for the Information Store is running	Active; CPS backup job running
When transaction logs are being replicated	Scheduled; CPS backup job running Note: You cannot right-click this job to edit the properties. You must click Job Setup, and then right-click the job to edit properties.
When a recovery point is running Note: The job name is displayed with Exchange Recovery Point appended to it.	Running

When the recovery point is complete, the recovery point appears in the Job History view. The recovery point appears as the name of the full job, with the description Exchange Recovery Point appended to it. If you enable error-handling rules, they apply to recovery points that fail.

If you place a CPS Exchange backup job on hold, the transaction log replication is stopped in CPS until you take the job off hold.

See [“About reviewing disk space availability for CPS Exchange backup jobs”](#) on page 1119.

See [“Stopping CPS Exchange backup jobs temporarily”](#) on page 1120.

See [“Viewing the CPS console from Backup Exec”](#) on page 1121.

See [“Troubleshooting CPS Exchange backup jobs”](#) on page 1123.

About reviewing disk space availability for CPS Exchange backup jobs

Hard links are a feature of the Microsoft operating systems, and are used when CPS Exchange backup jobs are processed. When you review disk space availability for CPS Exchange backup jobs, hard links affect the amount of reported available disk space.

The scheduled CPS Exchange full backup creates a media with a name such as IMG000060. The recovery points create VDB subfolders under the IMG media with names such as vdb_2007_03_08_1735_08. The Exchange transaction log files are put into the IMG media during backup. The logs in the VDB subfolders are hard links to the Exchange transaction log files. As VDB subfolders are created,

they contain all of the hard links from the previous recovery points until the next full backup runs.

For example:

VDB1 contains hard links to log files 1-5.

VDB2 contains hard links to log files 1-10.

VDB3 contains hard links to log files 1-15.

In this example, log files 1-5 are reported as having three times more space than they actually use.

A result of using hard links is that disk space usage appears to be greater than it is. For example, if 300 MB of disk space on a drive is used, it may appear that 500 MB of disk space is used. Be aware of this limitation when you review disk space availability for CPS Exchange backup jobs.

See [“Recommendations for using backup-to-disk folders with backup jobs that use Granular Recovery Technology”](#) on page 511.

Stopping CPS Exchange backup jobs temporarily

Stop the CPS Exchange backup jobs temporarily to perform any maintenance tasks that can affect the media server or the Exchange Server.

See the *Symantec Backup Exec Continuous Protection Server Administrator’s Guide* for more information on CPS procedures.

See [“Viewing the CPS console from Backup Exec”](#) on page 1121.

See [“About using recovery points to restore individual Exchange items to a point in time”](#) on page 1122.

See [“Best practices for continuous protection of Exchange”](#) on page 1117.

See [“About error-handling rules”](#) on page 590.

Table G-10 Stopping CPS Exchange backup jobs temporarily

Step	Action
Step 1	On the media server, put all of the scheduled occurrences of the active CPS Exchange backup job on hold. See “Placing all scheduled occurrences of an active job on hold” on page 563.

Table G-10 Stopping CPS Exchange backup jobs temporarily (*continued*)

Step	Action
Step 2	Stop the CPS service on the media server and on the Exchange Server, and then change the startup type to Manual.
Step 3	Stop the Backup Exec service on the media server, and then change the startup type to Manual.
Step 4	Perform the necessary maintenance on the media server or on the Exchange Server.
Step 5	After maintenance is complete, start the Backup Exec service on the media server, and then change the startup type to Automatic.
Step 6	Start the CPS service on the media server, and then change the startup type to Automatic.
Step 7	Start the CPS service on the Exchange Server, and then change the startup type to Automatic.
Step 8	Remove the scheduled occurrences of the active CPS Exchange backup job from hold. See “Placing all scheduled occurrences of an active job on hold” on page 563.

Viewing the CPS console from Backup Exec

If the Continuous Protection Server Administration Console component is installed on the media server, you can view the CPS console. The continuous backup job is displayed on the CPS console with ‘Backup Exec’ appended to the name. The job is listed as an Exchange Log Backup job type, with a status type of Running.

To view the CPS console from Backup Exec

- ◆ On the **Tools** menu, click **Continuous Protection Server**.

About using recovery points to restore individual Exchange items to a point in time

As part of the continuous protection of Exchange, you can enable Backup Exec to make recovery points at intervals that you specify. Recovery points create backup sets that you can browse from the Restore view. You can recover individual messages or folders from a point in time when either a full backup or recovery point was run. Each time a recovery point is made, it also truncates the transaction logs so that log growth is controlled.

Recovery points start to run at the specified intervals after the recurring full backup has started. However, recovery points do not run if the full backup is active. The recovery points start running again at the specified interval when the full backup has completed. Replication of the transaction logs is continuous, even when the full backup is active.

The recovery point only affects Exchange resources in the backup selection list. Resources that are not related to Exchange, but are in the same backup selection list, are not affected by recovery points.

Performance impacts of setting recovery point frequency for the Information Store transaction logs include the following:

Table G-11 Results of changing the default interval for recovery points

Making recovery points	Results
If recovery points are set to occur more frequently than the default rate of every 8 hours	The following may occur: <ul style="list-style-type: none">■ The Job Monitor view and the restore selections list may become crowded and difficult to read.■ The performance of the Exchange server may be slower than when the recovery points are set to occur at the default rate.
If recovery points are set to occur less frequently than the default rate of every 8 hours	Transaction logs are also deleted less frequently and therefore use more disk space.

Recovery points can only be created as part of the continuous protection strategy. If you choose not to use recovery points, individual mail messages and folders can only be recovered from the last full backup.

Caution: Transaction logs are deleted when recovery points occur. This may affect other Backup Exec jobs for the Exchange Server, or other jobs that are created by third-party applications.

See [“Best practices for continuous protection of Exchange”](#) on page 1117.

See [“About backing up Exchange 2003/2007”](#) on page 1129.

See [“About continuous protection for Exchange data”](#) on page 1112.

Troubleshooting CPS Exchange backup jobs

Use a log file named Trace_RBS_#####.txt on the Exchange Server to find out if transaction logs are replicated to the media server. The Backup Exec Continuous Protection Broker Service generates this log.

Entries that show that the transaction log was successfully replicated from the Exchange Server to the media server appear similar to the following example:

```
CRepServiceBroker::CheckLogReplica::CheckLocalLogReplica(target:MEDIA  
SERVER, EXCHANGE SERVER, First Storage Group, E00000F4.log, replicated:true)  
... hr(0x0)
```

MEDIA SERVER is the media server name, and EXCHANGE SERVER is the Exchange Server name. The text "E00000F4.log, replicated:true" means that the Exchange transaction log E00000F4.log was replicated.

Setting default backup and restore options for Exchange data

You can use the defaults set by Backup Exec during installation for all Exchange backup and restore jobs, or you can choose your own defaults.

See [“Backup strategies for Exchange”](#) on page 1102.

See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 1108.

See [“About restoring Exchange data”](#) on page 1144.

To set default backup and restore options for Exchange

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Microsoft Exchange**.
- 3 Select the appropriate options.

See [“Default backup and restore options for Exchange”](#) on page 1123.

Default backup and restore options for Exchange

You can set the following default options for all backup and restore jobs for Exchange.

See [“Setting default backup and restore options for Exchange data”](#) on page 1123.

Table G-12 Default backup and restore options for Exchange

Item	Description
<p>Information Store backup method</p>	<p>Specifies one of the following backup methods:</p> <ul style="list-style-type: none"> ■ Full - Database & Logs (flush committed logs). This method backs up the databases as well as their associated transaction log files. After the databases and transaction logs are backed up, the transaction log files for which all transactions committed to the database are then deleted. ■ Copy - Databases & Logs. This method backs up the databases as well as their associated transaction log files. However, the transaction logs are not deleted after being backed up. You can use the copy method to make a full backup of a database without disturbing the state of ongoing incremental or differential backups. ■ Differential - Logs. This method backs up all of the transaction logs that have been created or modified since the last full backup. However, the transaction logs are not deleted after being backed up. To restore from differential backups, the last differential backup and the last full backup are required. ■ Incremental - Logs (flush committed logs). This method backs up all of the transaction logs that have been created or modified since the last full or incremental backup, and then delete the transaction logs that have been committed to the database. To restore from incremental backups, the last full backup and all incremental backups done since the last full backup are required. See “Snapshot and offhost backups with the Exchange Agent” on page 1108. If circular logging is enabled, incremental, differential, and continuous protection backups cannot be performed.

Table G-12 Default backup and restore options for Exchange (*continued*)

Item	Description
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups (Incremental and Differential backups supported with policy-based jobs only)	<p>Restores individual items from Information Store backups. Ensure that you meet the requirements for Granular Recovery Technology.</p> <p>See “Recommended devices for backups that use Granular Recovery Technology” on page 324.</p>
Enable legacy mailbox support (Exchange 2003). This option is not recommended; use GRT instead.	<p>Lets you select individual mailboxes for backup with the Information Store.</p> <p>See “Backing up individual Exchange mailboxes” on page 1143.</p> <p>Note that a separate job to back up mailboxes is not necessary.</p> <p>See “How to restore individual items by using Granular Recovery Technology” on page 321.</p> <p>You must check Enable legacy mailbox support to enable the mailbox backup methods to appear on the Exchange Backup Properties page.</p> <p>If you are upgrading from a previous version of Backup Exec, Enable legacy mailbox support is checked by default. Jobs for mailbox backups can continue as scheduled.</p> <p>Uncheck Enable legacy mailbox support to make the mailbox resources unavailable on the backup selections tree.</p>

Table G-12 Default backup and restore options for Exchange (*continued*)

Item	Description
Mailbox backup method	<p>The following backup methods are available:</p> <ul style="list-style-type: none"> <p>■ Full - Back up messages - Reset archive bit. This method backs up all messages in the selected mailboxes. This option is set by default. A time/date stamp that is placed in each folder indicates that the messages have been backed up.</p> <p>■ Copy - Back up messages. This method backs up all messages in the selected mailboxes. A time/date stamp is not used, so incremental and differential backups are not affected. Use the copy method to make a full backup of the mailboxes without disturbing the state of ongoing incremental or differential backups.</p> <p>■ Differential - Back up changed messages. Differential - Back up changed messages. The time/date stamp that is placed on the folders during the last full backup is used to determine which messages have been modified since the last full backup. The time/date stamp is not updated during the differential backup.</p> <p>■ Incremental - Back up changed messages - Reset archive bit. This method backs up only the messages that have been modified in the selected mailboxes since the last full or incremental backup. The time/date stamp that is placed on the folders during the last full or incremental backup is used to determine which messages have been modified since the last full or incremental backup. The time/date stamp is updated during the incremental backup.</p>

Table G-12 Default backup and restore options for Exchange (*continued*)

Item	Description
Enable single instance backup for message attachments	<p>Backs up only a single copy of all identical message attachments. When an identical attachment is found, a reference to the attachment is retained. The actual attachment is backed up at the end of the backup set.</p> <p>Enabling single instance backup for message attachment increases backup performance since duplicate attachments are backed up only once.</p> <p>Uncheck Enable single instance backup for message attachments if you want each identical copy of a message attachment to be backed up and kept in order on the backup set.</p> <p>If the backup job does not complete, the message attachments may not be included in the backup set. Run the backup until it completes successfully .</p> <p>If the incremental backup method was used, running the job again will not back up the same messages and attachments. You must run a full or copy backup to ensure that all messages and attachments are backed up completely.</p>
Back up the information used to automatically recreate user accounts and mailboxes	Lets you automatically recreate user accounts and mailboxes during a restore.

Table G-12 Default backup and restore options for Exchange (*continued*)

Item	Description
<p>Temporary location for log and patch files</p>	<p>Specifies a location where the associated log and patch files are to be kept until the database is restored. The default location is \temp. If storage groups are being restored, a subdirectory in \temp is created for each storage group. The log and patch files for each storage group are kept in the corresponding subdirectory.</p> <p>If Commit after restore completes is selected for the restore job, the log and patch files in the temporary location are applied to the database, and then the current log files are applied. After the restore is complete, the log and patch files are automatically deleted from the temporary location (including any subdirectories).</p> <p>See “About restoring Exchange data” on page 1144.</p> <p>Make sure the temporary location for log and patch files is empty before you start a restore job. If a restore job fails, check the temporary location (including subdirectories) to make sure that any log and patch files from a previous restore job were deleted.</p>
<p>Automatically recreate user accounts and mailboxes</p>	<p>Recreates the user accounts and their mailboxes if they do not already exist on the destination server. The restore job fails if a mailbox that is being restored does not exist on the destination server.</p> <p>To restore any mailboxes that were backed up with the legacy backup method, the option Back up the information used to automatically recreate user accounts and mailboxes must have been selected for the backup job.</p> <p>See “Backing up individual Exchange mailboxes” on page 1143.</p> <p>When you check Automatically recreate user accounts and mailboxes, you must enter a password for accounts that are recreated.</p> <p>Automatically recreate user accounts and mailboxes applies only if mailboxes are being restored to their original location. If the mailbox restore is being redirected, the user account and mailbox must already exist on the target server.</p>
<p>Change password...</p>	<p>Specifies a password to use when user accounts and mailboxes are automatically recreated on the destination server.</p>

Table G-12 Default backup and restore options for Exchange (*continued*)

Item	Description
When restoring individual mail messages and folders, restore over existing messages and folders	<p>Replaces an existing item with the message or folder. Check this when you are restoring from a continuous protection recovery point backup or from a backup that uses Granular Recovery Technology (GRT). A new object ID is not created for the message or folder; only the contents and properties are replaced.</p> <p>If this check box is not checked, or if the original message or folder does not exist, then the message or folder is recreated as a new message or folder.</p> <p>If this check box is not checked and if the original message or folder does exist, then the message or folder is skipped.</p> <p>See “How to restore individual items by using Granular Recovery Technology” on page 321.</p> <p>See “About using recovery points to restore individual Exchange items to a point in time” on page 1122.</p>

About backing up Exchange 2003/2007

To back up Exchange data, you can select the following:

- Multiple storage groups
- Individual storage groups
- Individual databases (not supported if using snapshot technology)

Symantec recommends that you select individual storage groups for backup rather than selecting individual databases in storage groups. Although you can select individual databases in a storage group for backup, the transaction logs for the entire storage group are backed up for each database selected.

For example, if you select four databases in a storage group for backup, the entire collection of transaction logs for the storage group is also backed up four times. The transaction logs are not deleted until a full backup is run on every database in the storage group. You can still restore an individual database from a storage group backup.

Note: To perform incremental and differential backups of storage groups, or to perform a backup job for which continuous protection is enabled, make sure that circular logging is not enabled on the storage group.

See [“About backing up Exchange Server 2010 Databases”](#) on page 1130.

See [“Backing up Exchange”](#) on page 1132.

About backing up Exchange Server 2010 Databases

You can view a forest and the Database Availability Groups (DAG) that a forest contains in the backup selections pane. Backup Exec automatically adds the local forest that contains the Backup Exec media server to the **Microsoft Exchange Database Availability Groups** resource. All the DAG resources that the forest contains are also added to the list.

Note: You must enter an Exchange Agent license key on the media server for each Exchange server in the DAG that you want to back up. You must then install the Remote Agent for Windows Systems on all the servers in the DAG.

To back up the individual servers in a DAG, you can make backup selections from **Favorite Resources**, **Domains**, or **User-defined Selections**. If the DAG that contains the server is not in the list, Backup Exec automatically adds the DAG. You can also manually refresh the **Microsoft Exchange Database Availability Groups** resource to discover a DAG.

If the forest you want to use for backup selections is not in the list, you can manually add a forest to the **Microsoft Exchange Database Availability Groups** resource.

See [“About the Remote Agent for Windows Systems”](#) on page 1945.

See [“Adding licenses”](#) on page 175.

See [“Adding an Exchange Server 2010 forest to backup selections”](#) on page 1130.

See [“Backing up Exchange”](#) on page 1132.

Adding an Exchange Server 2010 forest to backup selections

You can add a forest to the **Microsoft Exchange Database Availability Groups** resource to make backup selections.

See [“About backing up Exchange Server 2010 Databases”](#) on page 1130.

See [“Backing up Exchange”](#) on page 1132.

To add an Exchange Server 2010 forest to backup selections

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.

- 3 In the **Properties** pane, under **Source**, click **Selections**.
- 4 On the backup selections list, right-click **Microsoft Exchange Database Availability Groups**.
- 5 Click **Add Forest**.
- 6 Select the appropriate options.
See [“Add Exchange Server 2010 forest options”](#) on page 1131.
- 7 Click **OK**.

Add Exchange Server 2010 forest options

You can set the following options for an Exchange Server 2010 forest.

See [“Adding an Exchange Server 2010 forest to backup selections”](#) on page 1130.

Table G-13 Exchange Server 2010 forest options

Item	Description
Domain controller or Database Availability Group server name	Lets you enter the name of the domain controller or the DAG node that contains the forest.
Logon Account	Select the name of the logon account that has rights to the forest.
New	Lets you create and add a logon account to the list. See “Creating a Backup Exec logon account” on page 187.

Managing an Exchange Server 2010 forest

You can add or remove a forest from the **Microsoft Exchange Database Availability Groups** resource.

See [“About backing up Exchange Server 2010 Databases”](#) on page 1130.

To manage an Exchange Server 2010 forest

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 On the backup selections tree, right-click **Microsoft Exchange Database Availability Groups**.

- 5 Click **Manage Forests**.
- 6 Select the appropriate options.
See [“Manage Exchange Server 2010 forest options”](#) on page 1132.
- 7 Click **Close**.

Manage Exchange Server 2010 forest options

You can use the following options to manage an Exchange Server 2010 forest.
See [“Managing an Exchange Server 2010 forest”](#) on page 1131.

Table G-14 Exchange Server 2010 manage forest options

Item	Description
Forests	Lists the forests that are available for backup selections.
Add	Lets you add a forest to the list of backup selections.
Delete	Lets you remove a forest from the list of backup selections.

Backing up Exchange

This procedure provides details on how to back up Exchange.

See [“About the circular logging setting for Exchange”](#) on page 1106.

See [“About backup-to-disk folders ”](#) on page 496.

See [“About backing up Exchange Server 2010 Databases”](#) on page 1130.

To back up Exchange

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 View the Exchange data that you want to back up by doing one of the following:

To display Exchange data on local or remote computers: Do the following in the order listed:

- Click the domain name icon or icons that contain the Exchange installations
- Expand the actual Windows computer icon that contains the Exchange installation.

To display Exchange data from a server cluster: Do the following in the order listed:

- On the virtual server, click the domain name icon or icons that contain the Exchange installations
- Expand the actual Windows computer icon that contains the Exchange installation.

To display an Exchange Server 2010 Database Availability Group (DAG): Do the following in the order listed:

- Expand the **Microsoft Exchange Database Availability Groups** resource.
- Expand the forest that contains the DAG.
- Expand the DAG that contains the Exchange installation.

5 Do one of the following:

To select all Exchange databases for backup Check **Microsoft Information Store**.

To select specific Exchange 2003/2007 storage groups or Exchange Server 2010 databases Expand the Microsoft Information Store icon, and then select individual storage groups or databases.

6 If prompted, select a logon account that you can use to connect to the Exchange Server.

7 On the **Properties** pane, under **Settings**, click **Microsoft Exchange**.

8 Select options for the backup job.

See "[Microsoft Exchange backup options](#)" on page 1134.

9 Start the backup job or select other backup options from the **Properties** pane, and then start the backup job.

See "[Creating a backup job by setting job properties](#)" on page 332.

Microsoft Exchange backup options

You can set the following options when you create a backup job for Exchange.

See [“Backing up Exchange”](#) on page 1132.

Table G-15 Exchange backup options

Item	Description
Information Store Backup method	<p>Specifies one of the following backup methods:</p> <ul style="list-style-type: none"> <p>■ Full - Database & Logs (flush committed logs). This method backs up the databases as well as their associated transaction log files. After the databases and transaction logs are backed up, the transaction log files for which all transactions are committed to the database are then deleted.</p> <p>■ Copy - Databases & Logs. This method backs up the databases as well as their associated transaction log files. However, the transaction logs are not deleted after being backed up. You can use the copy method to make a full backup of a database without disturbing the state of ongoing incremental or differential backups.</p> <p>■ Differential - Logs. This method backs up all of the transaction logs that have been created or modified since the last full backup. However, the transaction logs are not deleted after being backed up. To restore from differential backups, the last differential backup and the last full backup are required.</p> <p>■ Incremental - Logs (flush committed logs). This method backs up all of the transaction logs that have been created or modified since the last full or incremental backup, and then delete the transaction logs that have been committed to the database. To restore from incremental backups, the last full backup and all incremental backups done since the last full backup are required. See “Snapshot and offhost backups with the Exchange Agent” on page 1108.</p> <p>If circular logging is enabled, incremental, differential, continuous protection backups cannot be performed.</p>

Table G-15 Exchange backup options (*continued*)

Item	Description
Continuously back up transaction logs with Backup Exec Continuous Protection Server	<p>Enables a full restore of the Exchange database that includes transaction logs that are continuously protected between full backups.</p> <p>Before you check this check box, ensure that the Exchange server meets the requirements for configuring continuous protection.</p> <p>See “Requirements for configuring continuous protection for Exchange data” on page 1115.</p> <p>See “Best practices for continuous protection of Exchange” on page 1117.</p>
Make a recovery point that creates browsable backup sets and truncates logs every	<p>Creates backup sets that you can browse from the Restore view. You can recover individual messages or folders from a point in time when the last recovery point was run. Each time a recovery point is made, it also truncates the transaction logs so that log growth is controlled.</p> <p>If circular logging is enabled, the recurring full backup of the Information Store completes without errors, but recovery points fail.</p> <p>If you do not check Make a recovery point that creates browsable backup sets and truncates logs every, then individual mail messages and folders can be restored only from the CPS backup.</p> <p>The default interval is 8 hours. The minimum interval that you can set is every 15 minutes. The maximum interval is 1 year. If you change the specified interval for the recovery points, the new interval applies after the next full backup or recovery point is run.</p> <p>Before you change the default interval, review performance factors for setting recovery point intervals.</p> <p>See “About using recovery points to restore individual Exchange items to a point in time” on page 1122.</p>

Table G-15 Exchange backup options (*continued*)

Item	Description
<p>Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups (Incremental and Differential backups supported with policy-based jobs only)</p>	<p>Enables the restore of individual items from Information Store backups. Ensure that you meet the requirements for Granular Recovery Technology.</p> <p>See “Recommended devices for backups that use Granular Recovery Technology” on page 324.</p> <p>See “About requirements for jobs that use Granular Recovery Technology” on page 325.</p> <p>The GRT option is automatically enabled when you check the option Continuously back up transaction logs with Backup Exec Continuous Protection Server. The combined options let you restore individual items from a CPS Exchange backup.</p>
<p>Guide Me</p>	<p>Starts a wizard that helps you choose backup job properties for backing up Exchange data.</p>

Table G-15 Exchange backup options (*continued*)

Item	Description
<p>Perform consistency check before backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider</p>	<p>Performs a consistency check when the Microsoft Volume Shadow Copy Service option is selected. The option Microsoft Volume Shadow Copy Service is automatically used whenever a software backup is selected on the Advanced Disk-based Backup Option backup properties. You can also select the Microsoft Volume Shadow Copy Service on the Advanced Open File Option backup property page.</p> <p>The consistency check, which is run on the snapshot, determines if possible data corruption exists.</p> <p>If this option is selected, and the dependent option Continue with backup if consistency check fails is not selected, then data for specific Exchange objects that are determined to be corrupt are not backed up. All other non-corrupt Exchange objects are backed up.</p> <p>For example, if any transaction log file for a Storage Group is corrupt, then none of the transaction log files are backed up for that Storage Group when the Continue with backup if consistency check fails option is not selected. However, the Exchange database files are backed up if Backup Exec determines they are not corrupt. Similarly, if a specific Exchange database file is corrupt, then backup is skipped only for that corrupt database file. All other non-corrupt database files and transaction log files are backed up.</p> <p>When the option Continue with backup if consistency check fails is enabled, then all Exchange data is backed up regardless if corrupt files exist.</p> <p>See “Snapshot and offhost backups with the Exchange Agent” on page 1108.</p>
<p>Continue with backup if consistency check fails</p>	<p>Continues the backup job even if the consistency check fails. You may want the job to continue if you think a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database that may have only a small problem.</p>

Table G-15 Exchange backup options (*continued*)

Item	Description
<p>High Availability Server (Exchange 2007 or later)</p>	<p>Specifies one of the following backup sources for Exchange 2007/2010:</p> <ul style="list-style-type: none"> <p>■ Back up from the active copy only (job fails if not available)</p> <p>Lets you back up the active copy of the database. If Backup Exec cannot access the active copy, the job fails. Therefore, neither the active copy nor the passive copy is backed up.</p> <p>The active copy contains newer information than the passive copy. Therefore, when you back up the active copy, you have a backup of the most recent database data.</p> <p>Note: You cannot back up the passive copy of the Standby Continuous Replication (SCR) database with Exchange Server 2007. The SCR is not available for backup selection.</p> <p>■ Let Backup Exec automatically choose the best copy to back up (recommended)</p> <p>Lets you back up a passive copy of the database by default. Backup Exec selects the passive copy based on your selections in the Preferred Server settings. However, if the passive copy is not available, Backup Exec backs up the active copy of the database. During the backup, database performance degradation can occur if you have to back up the database over a WAN.</p> <p>■ Back up from the passive copy only, using Preferred Server settings if possible (job fails if not available)</p> <p>Lets you back up a passive copy of the database. If Backup Exec cannot access the passive copy, the job fails. In this case, neither the active nor the passive database is backed up. Select this option when you do not want to affect the performance of the active copy of the database. For Exchange Server 2010, Backup Exec selects the passive copy based on your selections in the Preferred Server settings.</p> <p>Note: You must have the preferred server settings configured to use this option.</p> <p>See “About preferred server configurations” on page 435.</p>

Table G-15 Exchange backup options (*continued*)

Item	Description
<p>Mailbox backup method</p> <p>Note: This option is available when you select the Enable legacy mailbox support (Exchange 2003). This option is not recommended; use GRT instead option.</p>	<p>The following backup methods are available:</p> <ul style="list-style-type: none"> ■ Full - Back up messages - Reset archive bit. This method backs up all messages in the selected mailboxes. This option is set by default. A time/date stamp that is placed in each folder indicates that the messages have been backed up. ■ Copy - Back up messages. This method backs up all messages in the selected mailboxes. A time/date stamp is not used, so incremental and differential backups are not affected. Use the copy method to make a full backup of the mailboxes without disturbing the state of ongoing incremental or differential backups. ■ Differential - Back up changed messages. This method backs up all of the messages that have been created or modified in the selected mailboxes since the last full backup. The time/date stamp that is placed on the folders during the last full backup is used to determine which messages have been modified since the last full backup. The time/date stamp is not updated during the differential backup. ■ Incremental - Back up changed messages - Reset archive bit. This method backs up only the messages that have been modified in the selected mailboxes since the last full or incremental backup. The time/date stamp that is placed on the folders during the last full or incremental backup is used to determine which messages have been modified since the last full or incremental backup. The time/date stamp is updated during the incremental backup.

Table G-15 Exchange backup options (*continued*)

Item	Description
<p>Enable single instance backup for message attachments (for mailbox and public folder backups only)</p> <p>Note: This option is available when you select the Enable legacy mailbox support (Exchange 2003). This option is not recommended; use GRT instead option.</p>	<p>Backs up only a single copy of all identical message attachments. When an identical attachment is found, a reference to the attachment is retained. The actual attachment is backed up at the end of the backup set.</p> <p>Enabling single instance backup for message attachment increases backup performance since duplicate attachments are backed up only once.</p> <p>Uncheck Enable single instance backup for message attachments if you want each identical copy of a message attachment to be backed up and kept in order on the backup set.</p> <p>If the backup job does not complete, the message attachments may not be included in the backup set. Run the backup until it completes successfully completed.</p> <p>If the incremental backup method was used, running the job again will not back up the same messages and attachments. You must run a full or copy backup to ensure that all messages and attachments are backed up completely.</p>
<p>Guide Me</p>	<p>Starts a wizard that helps you choose backup job properties for backing up Exchange mailboxes.</p>

About selecting individual Exchange mailboxes for backup

In previous versions of Backup Exec, you backed up individual Exchange mailboxes separately from the Information Store so that you could restore individual mailboxes. In this version of Backup Exec, you can enable the option to restore individual mail messages and folders from Information Store backups.

Restoring individual mailboxes from an Information Store backup that is on tape may be slower than restoring mailboxes from a legacy mailbox backup. If you have only a few mailboxes, and if you can only back up to tape, you may prefer to use the legacy mailbox backup methods.

Note: Exchange 2007/2010 does not support individual mailbox backup. You can use the Backup Exec Granular Recovery Technology option to restore individual items from Exchange 2007/2010.

If you must back up individual Exchange mailboxes separately from the Information Store, then consider the following:

Table G-16 Recommendations for legacy backup methods for mailboxes

Recommendation	Description
Use full and incremental backups	Consider running full backups of mailboxes or public folders on a regular basis. Supplement the full backups with incremental or differential backups to keep backup run time to a minimum.
Continue to back up the Information Store	Do not substitute mailbox backups for backups of the entire Information Store. You cannot perform a complete restore of Exchange Server from a mailbox backup. You can only perform a complete restore of Exchange Server from backups of the Information Store.
Recover deleted items rather than restoring them	Consider using the Exchange System Manager utility to adjust the deletion settings in each Store's properties. Deleted items can be retained for a period of time, allowing them to be recovered rather than restored. See your Microsoft Exchange Server documentation for details.
Exclude unwanted or unnecessary folders from the backup	<p>When you select a mailbox or public folder for backup, by default all folders and subfolders are included. You can exclude specific folders and subfolders.</p> <p>See “How to include or exclude files for backup” on page 355.</p> <p>For example, to exclude all mail in the Deleted Items folder, type:</p> <pre>**\Deleted Items*</pre> <p>To exclude all mail in the Sent Items folder, type:</p> <pre>**\Sent Items*</pre>
Enable single-instance backup for message attachments	<p>When backing up mailboxes and public folders, you can choose to back up only a single copy of all identical message attachments. When an identical attachment is found, a reference to that attachment is retained. The actual attachment is backed up only once at the end of the backup set.</p> <p>Enabling single instance backup for message attachment increases backup performance since duplicate attachments are backed up only once.</p> <p>Note: If the backup job does not complete, the message attachments may not be included in the backup set. Rerun the backup until it successfully completes. If you used the incremental backup method, running the job again does not back up the same messages and attachments. You must run a full or copy backup to ensure that all messages and attachments are backed up completely.</p>

Table G-16 Recommendations for legacy backup methods for mailboxes
(continued)

Recommendation	Description
<p>Do not back up special system mailboxes created by Exchange</p>	<p>Although these special system mailboxes can be backed up, it is not necessary or useful.</p> <p>The following are common examples of special system mailboxes:</p> <ul style="list-style-type: none"> ■ System Attendant ■ Any mailbox name starting with SMTP or System Mailbox <p>There may be others depending on the Exchange server configuration and environment.</p> <p>Also, when selecting objects from the mailbox tree, all objects are displayed as messages. Some non-message objects can be identified by the subject line. For example, if you create a Calendar event named Appointment1, that name is displayed in the subject line for that object. However, some objects such as Forms and Views do not have a subject line (even though they can be named) and may not be easily identifiable.</p>
<p>Select public folders from only one Exchange server</p>	<p>The same public folders may appear for multiple Exchange servers since public folders can be replicated. Selecting public folders on multiple Exchange servers only increases the time and media required for the backup and does not provide any additional protection.</p>

Note: Antivirus software may impact mailbox backup performance and could result in incorrect job log errors. Mail messages and the attachments will be fully backed up and fully restorable despite job log errors that may be generated when verifying the attached files.

Mailboxes are displayed in a server-centric view; that is, only mailboxes on the selected Exchange server are displayed in the backup selections list. Public folders may be displayed on more than one server since public folders can be replicated to many servers.

In versions of Backup Exec prior to 8.6, mailboxes could be selected for backup from a site-centric view, which listed all mailboxes in the Exchange Organization, not just mailboxes on the selected server. Even though only the server-centric view is now available for Exchange mailboxes, you can still restore mailbox backup sets that were created using a site-centric view.

If you select mailboxes from the same server, they are placed together in one backup set on the storage media. If you select mailboxes from more than one server, then the mailboxes are placed in separate backup sets according to the server.

Backing up individual Exchange mailboxes

If you have only a few mailboxes, and if you can only back up to tape, you may prefer to use this legacy mailbox backup method.

Note: Exchange 2007/2010 does not support individual mailbox backup.

See [“About selecting individual Exchange mailboxes for backup”](#) on page 1140.

See [“Requirements for accessing Exchange mailboxes”](#) on page 1101.

See [“About selecting individual Exchange mailboxes for backup”](#) on page 1140.

To back up individual Exchange mailboxes

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Microsoft Exchange**.
- 3 Check **Enable legacy mailbox support (Exchange 2003)**. **This option is not recommended; use GRT instead.**
- 4 If you want to recreate user accounts and mailboxes during a restore, check **Back up the information used to automatically recreate user accounts and mailboxes**.

This option is not checked by default.

- 5 Click **OK**.
- 6 On the navigation bar, click the arrow next to **Backup**.
- 7 Click **New Backup Job**.
- 8 On the **Properties** pane, under **Source**, click **Selections**.

- 9 View the Exchange mailboxes or public folders that you want to back up by doing one or all of the following:

To select Exchange mailboxes or public folders from local or remote computers:

Click the domain name icon or icons that contain the Exchange mailboxes. Expand the Microsoft Exchange Mailbox icon that contains the Exchange mailboxes or public folders.

To select Exchange mailboxes or public folders from a clustered Exchange server:

On the virtual server, click the domain name icon or icons that contain the Exchange mailboxes. Expand the Microsoft Exchange Mailbox icon that contains the Exchange mailboxes or public folders.

- 10 Select the mailboxes or individual messages, mailboxes, and folders that you want to back up.

When you select a mailbox or public folder, all folders and subfolders are included in the backup by default. For a faster backup, consider using Advanced File Selections to exclude some folders, such as Deleted Items or Sent Items, and subfolders from the backup.

Because public folders can be replicated on multiple Exchange servers, select public folders from only one Exchange server.

- 11 If prompted, select a logon account that allows you to connect to the Exchange mailboxes or public folders.

See [“Requirements for accessing Exchange mailboxes”](#) on page 1101.

- 12 On the **Properties** pane, under **Settings**, click **Microsoft Exchange**.

- 13 Select a backup method for the Information Store.

- 14 Select options for the backup job.

See [“Microsoft Exchange backup options”](#) on page 1134.

- 15 Start the backup job or select other backup options from the **Properties** pane, and then start the backup job.

See [“Creating a backup job by setting job properties”](#) on page 332.

About restoring Exchange data

You can use the defaults for all Exchange restore jobs, or you can choose your own defaults. You can also change the defaults for any specific restore job.

The requirements and procedures for restoring Exchange data vary depending on the backup strategy you used. Before you restore Exchange data, you should review the required configuration and tasks.

See [“Requirements for restoring Exchange 2000 or later”](#) on page 1145.

You can restore Exchange data in the following ways:

- Use the recovery storage group or recovery database to recover data from an older backup copy of the store without disturbing client access to current data. See [“About restoring data using the Exchange 2003/2007 recovery storage group or Exchange Server 2010 recovery database”](#) on page 1147.
- Restore Exchange data from snapshot backups. See [“About restoring Exchange data from snapshot backups”](#) on page 1149.
- Restore the Exchange database from continuous protection backup sets and restoring individual mail messages and folders. See [“About restoring Exchange data from continuous protection backups”](#) on page 1150.
- Restore individual Exchange items from a backup that uses Granular Recovery Technology (GRT). See [“Recommended devices for backups that use Granular Recovery Technology”](#) on page 324.
- Restore Exchange mailboxes and folders from a backup that is separate from the Information Store backup. See [“About restoring Exchange mailboxes and public folders from mailbox backups”](#) on page 1152.
- Restore individual Exchange public folder messages from tape. See [“Restoring individual Exchange public folder messages from tape by duplicating backup sets to disk”](#) on page 1153.
- Configure a restore job for Exchange data. See [“Restoring Exchange data”](#) on page 1154.
- Restore Exchange data to a server other than the one from which it was backed up. See [“About redirecting Exchange restore data”](#) on page 1159.

Requirements for restoring Exchange 2000 or later

Review the following before restoring Exchange 2000 or later:

- The storage groups and databases must already exist on the target server, and must have the same names as the original storage groups or databases.

- The target server must have the same Organization and Administrative Group name as the source server.

Before you start the restore, do the following:

- Configure the target databases so that they can be overwritten
- Dismount databases that are being restored

Configuring a database in Exchange

Before you restore Exchange, you should configure the target database.

To configure a database

- 1 Do one of the following:

For Exchange 2000/2003

Open the Exchange System Manager utility.

For Exchange 2007/2010

Open the Exchange Management Console utility.

- 2 Right-click the database that you want to overwrite.
- 3 Click **Properties**.
- 4 Do one of the following:

For Exchange 2000/2003/2007

On the **Database** tab, select **This database can be overwritten by a restore**.

For Exchange Server 2010

On the **Maintenance** tab, select **This database can be overwritten by a restore**.

Dismounting Exchange databases that are being restored

Before you restore Exchange, you should dismount the databases that are being restored.

To dismount databases that are being restored

- ◆ Do one of the following:
 - Use the Exchange System Manager utility or the Exchange Management Console utility.

- When creating a restore job, on the Backup Exec **Restore Job Properties** pane, under **Settings**, click **Microsoft Exchange**. Check **Dismount database before restore**.

About restoring data using the Exchange 2003/2007 recovery storage group or Exchange Server 2010 recovery database

The Recovery Storage Group (RSG) feature in Exchange 2003/2007 lets you to mount a second copy of an Exchange mailbox store on any Exchange server in the same Exchange Administrative Group as the original while the original store is still running and serving clients. This allows you to recover data from an older backup copy of the store without disturbing client access to current data.

Exchange Server 2010 uses recovery databases instead of RSGs. Each server has a recovery database and there cannot be more than one mounted recovery database.

See your Microsoft Exchange documentation for more information about RSGs and recovery databases.

After the RSG or recovery database is created, you can restore online backup sets to it. Then you can use the version of the EXMerge utility in Exchange 2003 or Exchange Management Shell in Exchange 2007/2010 to extract mailbox data from the stores into .PST files, and optionally merge the extracted data back into the online stores.

If the RSG or recovery database resides on a different Exchange server than the databases you are restoring, you should review the requirements for redirecting the restore of Exchange storage groups or recovery databases.

See [“About redirecting Exchange storage group and database restores”](#) on page 1160.

Following are requirements for restoring data using the Exchange 2003/2007 data Recovery Storage Group (RSG) or Exchange Server 2010 recovery database:

- For Exchange 2003, data cannot be restored from a snapshot backup.
- If multiple stores are selected for restore, mailbox stores in the RSG must come from the same storage group. You cannot add mailbox stores from different storage groups to the RSG at the same time.
- Public folder stores are not supported for restore using the RSG.
- Do not mount mailbox stores in the RSG before the restore. If you do mount the stores before the restore, then you must dismount them. Select the following option on the database property page in Exchange System Manager: This database can be overwritten by a restore

Then, delete any files created in the data path for the RSG and added stores prior to restoring them.

Any files created in the data path for the RSG and added store or stores, should be deleted as well, if stores were mounted prior to the restore.

- On the server that hosts the RSG, there must be a storage group with the same name as the original storage group for the data you are restoring. If no such storage group exists on the server, then you can use that name for the RSG when you create it.
- The Active Directory topology of the Exchange system must be intact and in the same state it was in when the backup was made. You cannot restore mailbox stores that were deleted and recreated. In addition, you cannot recover mailboxes from stores if the mailboxes were deleted and purged from the system or moved to other servers or mailbox stores.
- Only Exchange mailbox stores from Exchange 2000 Server with Service Pack 3 or later can be restored to the RSG. Restored mailbox stores are upgraded to the store version currently running on the RSG server.
- When the RSG exists on a server, the mailbox stores that it contains are the only stores that can be restored on that server by default. Symantec recommends that you create the RSG only when you intend to recover data using it, and remove the RSG from the server after the data recovery is complete.
- You can have more than one recovery database, however, you can only mount one recovery database to recover data.
- Do not mount the recovery database before the restore. If you do mount the recovery database before the restore, you must dismount it. Select the **This database can be overwritten by a restore** option on the database property page in Exchange Management Console utility.

Refer to your Microsoft Exchange Server documentation for more information on the requirements and restrictions of recovering Exchange data.

See [“Restoring Exchange data”](#) on page 1154.

See [“About redirecting Exchange storage group and database restores”](#) on page 1160.

Restoring a database to an Exchange 2007 recovery storage group

Use the following steps to restore a database to an Exchange 2007 recovery storage group on an Exchange 2007 server. After restoring the database, see your Microsoft Exchange documentation for further Exchange recovery information.

See [“Restoring Exchange data”](#) on page 1154.

See [“About redirecting Exchange storage group and database restores”](#) on page 1160.

To restore a database to an Exchange 2007 recovery storage group

- 1 Make sure that a recovery storage group exists on the destination Exchange 2007 server, and that a recovery database exists within the recovery storage group. If either of these do not exist, you must create them before continuing. To create a recovery storage group or a recovery database, see your Microsoft Exchange documentation.
- 2 On the navigation bar, click the arrow next to Restore.
- 3 Click **New Restore Job**.
- 4 On the **Restore Job Properties** pane, under **Source**, click **Selections**.
- 5 Select an Exchange database to restore.
- 6 On the **Restore Job Properties** pane, under **Destination**, click **Microsoft Exchange Redirection**.
- 7 Click **Redirect Exchange sets**.
- 8 In the **Restore to server or Database Availability Group** field, type the name of the destination Exchange server.
- 9 Click **Redirect using Volume Shadow Copy Service (VSS) snapshot provider**.
- 10 Click **Redirect to Recovery Storage Group (RSG) (Exchange 2007 only)**.
- 11 Click **Run Now**.

Backup Exec restores the Exchange 2007 database to the destination recovery storage group. After the restore job completes, see your Microsoft Exchange documentation for further Exchange 2007 recovery information.

About restoring Exchange data from snapshot backups

Note the following when restoring Exchange data from snapshot backups:

- If circular logging is enabled, only point-in-time, loss restores are possible. Roll-forward, no-loss restores cannot be performed.
- For Exchange 2003/2007, individual database restores cannot be performed. The job will fail.
- The following options are not applicable to restores of Exchange 2003 snapshot backups. Exchange will use the soft-recovery process when restored databases are mounted.
 - Restore all transaction logs; do not delete existing transaction logs (no loss restore)

- Restore all transaction logs until point-in-time; skip transaction logs after this time
- Purge existing data and restore only the databases and transaction logs from the backup sets
- Path on Exchange Server for temporary storage of log and patch files
- Commit after restore completes
- For Exchange 2003, data cannot be restored from a snapshot backup to a Recovery Storage Group (RSG).

See [“About restoring data using the Exchange 2003/2007 recovery storage group or Exchange Server 2010 recovery database”](#) on page 1147.

See [“About selecting individual Exchange mailboxes for backup”](#) on page 1140.

See [“Restoring Exchange data”](#) on page 1154.

About restoring Exchange data from continuous protection backups

The full backups and the replicated transaction logs provide a complete restore to any point in time of the Information Store, including the latest complete transaction log.

Note: After you restore a storage group or mailbox store from a CPS Exchange backup, you must restart the CPS Exchange backup job. Otherwise, the continuous protection job and any associated recovery points do not restart.

If you enabled recovery points to run at intervals between the full backups, you can restore individual messages or folders at a point in time when the recovery point was created. Even without recovery points, you can restore individual messages or folders from a full backup.

See [“Restoring the entire Information Store to the time of a full backup or recovery point from continuous protection backups”](#) on page 1151.

See [“Restoring up to the latest full transaction log from continuous protection backups”](#) on page 1151.

See [“Restoring up to a point in time between full backups or recovery points from continuous protection backups”](#) on page 1152.

See [“Restoring Exchange data”](#) on page 1154.

See [“About continuous protection for Exchange data”](#) on page 1112.

See [“About redirecting Exchange restore data”](#) on page 1159.

Restoring the entire Information Store to the time of a full backup or recovery point from continuous protection backups

Use the following steps to restore the entire Information Store to the time of a full backup or recovery point.

See [“About restoring Exchange data from continuous protection backups”](#) on page 1150.

To restore the entire Information Store to the time of a full backup or recovery point

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 Select backup sets from the full backup or the recovery point that contains the point in time that you want to restore to.
- 5 On the **Properties** pane, under **Settings**, click **Microsoft Exchange**.
- 6 Click **Purge existing data and restore only the databases and transaction logs from the backup sets**.
- 7 Click **Run now**.

Restoring up to the latest full transaction log from continuous protection backups

Use the following steps to restore up to the latest full transaction log.

See [“About restoring Exchange data from continuous protection backups”](#) on page 1150.

To restore up to the latest full transaction log

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 Select backup sets from the last full backup or recovery point.
- 5 On the **Properties** pane, under **Settings**, click **Microsoft Exchange**, and then click **Restore all transaction logs; do not delete existing transaction logs (no loss restore)**.
- 6 Click **Run now**.

Restoring up to a point in time between full backups or recovery points from continuous protection backups

Use the following steps to restore up to a point in time between full backups or recovery points.

See [“About restoring Exchange data from continuous protection backups”](#) on page 1150.

To restore to a point in time between full backups or recovery points

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 Select backup sets from any full backup or recovery point and specify the point in time.
- 5 On the **Properties** pane, under **Settings**, click **For continuous protection jobs only, restore all transaction logs until point in time; skip transaction logs after this time**.
- 6 Specify the point in time.
- 7 Click **Run now**.

About restoring Exchange mailboxes and public folders from mailbox backups

You can restore individual mailboxes, messages, and public folders from backups that were created by the following legacy methods:

- Individual mailboxes were backed up from the Microsoft Exchange Mailboxes selection.
- Individual public folders were backed up from the Microsoft Exchange Public Folders selection.

Following are notes about restoring mailboxes or public folders from legacy backup methods:

- Mailbox backup sets that were created using a site-centric view in versions of Backup Exec prior to 8.6 may require multiple jobs to be run in order to restore all the mailboxes:
Mailboxes that resided on the target server restore normally. Mailboxes that resided on other servers must be redirected to those servers.

- If you restore mailboxes from a backup created with a version of Backup Exec prior to 8.5, and mailboxes with duplicate Display Names reside on the server, select those mailboxes for restore separately and redirect the restore. See [“About redirecting restores of mailboxes or public folders that were backed up separately from the Exchange Information Store”](#) on page 1160.
 - The option Restore Over Existing Files in the Advanced Restore Job Properties does not apply to mailboxes or public folders. If other selections are restored in addition to mailboxes or public folders, and the option Restore Over Existing Files is selected, it applies only to the other selections; mailboxes and public folders are not restored over existing objects.
 - Do not restore special system mailboxes created by Exchange. The following are common examples of special system mailboxes, but there may be others depending on the Exchange server configuration and environment.
 - System Attendant
 - Any mailbox name starting with SMTP or System Mailbox (Exchange 2000 or later)
- See [“Requirements for restoring Exchange 2000 or later”](#) on page 1145.
- See [“Restoring Exchange data”](#) on page 1154.
- See [“About redirecting restores of mailboxes or public folders that were backed up separately from the Exchange Information Store”](#) on page 1160.
- See [“Creating selection lists”](#) on page 296.

Restoring individual Exchange public folder messages from tape by duplicating backup sets to disk

To restore individual public folder messages from tape, you must first duplicate the backup sets that contain the messages to a backup-to-disk folder. You can then restore the data from that folder.

The backup that you want to restore from must be a full backup or a copy backup. If there is an incremental or differential backup that is subsequent to the full backup, then you can restore individual items from the incremental or differential backup. The backup sets for the full backup and the incremental or differential backup must be on the same volume.

You cannot restore individual public folder messages from tape if the original backup is an incremental backup.

To restore individual Exchange public folder messages from tape by duplicating backup sets to disk

- 1 Insert the tape containing the required Exchange backup sets into a tape drive.
- 2 On the menu bar, click **File**, and then click **New > Duplicate Backup Sets Job**.
- 3 Click **Duplicate existing backup sets**.
- 4 Click **OK**.
- 5 Select the Exchange backup sets that you want to duplicate.
- 6 On the **Properties** pane, under **Destination**, click **Device and Media**.
- 7 In the **Device** list box, select a backup-to-disk folder.
- 8 Click **Run Now**.
- 9 After the job completes, run a restore job to restore the individual public folder messages from the Exchange backup sets that are duplicated in the backup-to-disk folder.

See [“Restoring Exchange data”](#) on page 1154.

Restoring Exchange data

This procedure details how to select restore job properties for Exchange, and provides definitions for restore options specific to Exchange.

Use the Exchange System Manager utility to manually dismount any databases that are being restored or check Dismount database before restore when creating the restore job.

See [“About restoring Exchange mailboxes and public folders from mailbox backups”](#) on page 1152.

See [“Requirements for restoring Exchange 2000 or later”](#) on page 1145.

Note: After you restore a storage group or mailbox store from a CPS Exchange backup, you must restart the CPS Exchange backup job. Otherwise, the continuous protection job and any associated recovery points do not restart.

To restore Exchange data

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.

- 4 In the restore selections list, select the backup sets that you want to restore, or expand the backup sets to select individual items for restore.
See [“About restoring Exchange data”](#) on page 1144.
- 5 On the **Properties** pane, under **Settings**, click **Microsoft Exchange**.
- 6 Select the appropriate options.
See [“Microsoft Exchange restore options”](#) on page 1155.
- 7 For backups that use Granular Recovery Technology, ensure that default staging locations are set correctly.
See [“Recommended devices for backups that use Granular Recovery Technology”](#) on page 324.
- 8 Select other restore options from the **Properties** pane as appropriate, and then start the restore job.
If you restore data from a CPS Exchange backup, you must restart the continuous protection job and any associated recovery points.
- 9 Do the following in the order listed after the job completes:
 - On the navigation bar, click **Job Monitor**.
 - Right-click the scheduled CPS Exchange backup job.
 - Click **Run now**.
- 10 Run a full backup of the restored databases.

Microsoft Exchange restore options

You can set the following options when you create a restore job for Exchange.

See [“Restoring Exchange data”](#) on page 1154.

Table G-17 Exchange restore options

Item	Description
<p>Automatically recreate user accounts and mailboxes</p>	<p>Recreates the user accounts and their mailboxes if they do not already exist on the target server. The restore job fails if a mailbox that is being restored does not exist on the destination server.</p> <p>Note: To restore any mailboxes that were backed up with the legacy backup method, the option Back up the information used to automatically recreate user accounts and mailboxes must have been selected for the backup job.</p> <p>See “Backing up individual Exchange mailboxes” on page 1143.</p> <p>When Automatically recreate user accounts and mailboxes is checked, the password that is entered on Tools > Options > Microsoft Exchange is used as the password for accounts that are recreated.</p> <p>See “Setting default backup and restore options for Exchange data” on page 1123.</p> <p>Automatically recreate user accounts and mailboxes applies only if mailboxes are being restored to their original location. If the mailbox restore is being redirected, the user account and mailbox must already exist on the target server.</p>
<p>When restoring individual mail messages and folders, restore over existing messages and folders</p>	<p>Replaces an existing item with the message or folder. A new object ID is not created for the message or folder; only the contents and properties are replaced.</p> <p>If this check box is not checked, or if the original message or folder does not exist, then it is recreated as a new message or folder; that is, a new object ID is created for it by Backup Exec.</p> <p>If this check box is not checked and if the original message or folder does exist, then the message or folder is skipped.</p>

Table G-17 Exchange restore options (*continued*)

Item	Description
<p>Restore all transaction logs; do not delete existing transaction logs (no loss restore)</p>	<p>Preserves the existing transaction logs on the Exchange server. Transaction logs from the storage media are then restored and added to the existing set of transaction logs on the Exchange server. When the restore operation finishes, Exchange automatically updates its databases with the uncommitted transactions found in the existing and newly restored transaction logs. This option is selected by default.</p> <p>Note: When you make restore selections, the backup sets include a transaction log number range. You can select the backup set that includes the transaction log you want to restore.</p> <p>If you are restoring individual databases into a storage group, you should select this option. If this option is not selected, uncommitted transactions for other databases in the storage group may be lost.</p> <p>If continuous protection is enabled, select backup sets from the last full backup or recovery point to restore up to the latest full transaction log.</p> <p>This option is not applicable to snapshot backups.</p>
<p>For continuous protection jobs only, restore all transaction logs until point-in-time; skip transaction logs after this time</p>	<p>Note: This option only supports restoring data from continuous protection backups. If you select this option to restore other types of backup data, the option is ignored, and a loss restore job runs.</p> <p>Restores transactions from a transaction log up to and including a point in time in the transaction log. After the point in time, recovery from the transaction log is stopped.</p> <p>Select the backup sets from a full backup or a recovery point, and then specify the point in time.</p> <p>In the date and time box, select the part of the date that you want to change. Enter a new date, or click the arrow to display a calendar from which you can select a date.</p> <p>Then, select the part of the time you want to change, and enter a new time.</p>

Table G-17 Exchange restore options (*continued*)

Item	Description
<p>Purge existing data and restore only the databases and transaction logs from the backup sets</p>	<p>Deletes the existing transaction logs. Only the databases and transaction logs from the backup sets are restored.</p> <p>If continuous protection is enabled, select the backup sets from a full backup or a recovery point that you want to restore the database to.</p> <p>When you restore an Exchange 2007 database and you enable this option, Backup Exec adds a .DELETE file name extension to all existing Exchange log files that it finds in the destination storage group.</p> <p>For example, Backup Exec renames e0001.log to e0001.log.delete.</p> <p>Backup Exec preserves the existing Exchange log files in the storage group until you manually delete them.</p>
<p>Path on Exchange Server for temporary storage of log and patch files</p>	<p>Specifies a location where the associated log and patch files are to be kept until the database is restored. The default location is \temp, and a subdirectory is created for each storage group. The log and patch files for each storage group are kept in the corresponding subdirectory.</p> <p>Make sure the temporary location for log and patch files is empty before you start a restore job. If a restore job fails, check the temporary location (including subdirectories) to make sure any previous log and patch files from a previous restore job were deleted.</p> <p>This option is not applicable to snapshot backups.</p>
<p>Dismount database before restore</p>	<p>Takes the Exchange database offline automatically before the restore job runs. If this option is not selected, you must manually take the database offline before the restore job can run.</p> <p>When restoring a snapshot backup, or when restoring an individual Exchange database from a snapshot, all databases in a storage group must be taken offline. Selecting this option automatically takes all databases in a storage group offline.</p>

Table G-17 Exchange restore options (*continued*)

Item	Description
Commit after restore completes	<p>Commits the last backup set if your selection contains the last backup set to be restored. This option directs the restore operation to replay the log files and roll back any uncompleted transactions. If this option is not selected, the database is left in an intermediate state and is not yet usable.</p> <p>If Commit after restore completes is checked when an intermediate backup is being applied, you cannot continue to restore backups. You must restart the restore operation from the beginning.</p> <p>After the database is restored, the log and patch files in the temporary location are applied to the database, and then the current log files are applied. After the restore is complete, the log and patch files are automatically deleted from the temporary location (including any subdirectories).</p> <p>This option is not applicable for snapshot backups.</p>
Mount database after restore	<p>Mounts the database so that it is available to users. This check box is only available if Commit after restore completes is selected.</p>
Guide Me	<p>Starts a wizard that helps you choose restore job properties for Exchange Server data.</p>

About redirecting Exchange restore data

With Backup Exec, you can restore Exchange data to the server from which it was backed up or redirect the Exchange data to another Exchange server. When redirecting Exchange data, the service pack on the Exchange server where data is being redirected should be the same as the service pack on the original Exchange server.

Review requirements for the following before redirecting restore data:

- Exchange 2000 or later storage group and databases, including public folders
See [“About redirecting Exchange storage group and database restores”](#) on page 1160.
- Exchange mailboxes or public folders backed up separately from the Information Store
See [“About redirecting restores of mailboxes or public folders that were backed up separately from the Exchange Information Store”](#) on page 1160.

Before starting the redirected restore job, review information on finding and viewing specific data to restore, as well as for details on restore options and submitting restore jobs.

See [“About restoring data”](#) on page 599.

After completing the restore, it is recommended that a full backup of the restored databases be performed.

See [“Redirecting Exchange restore data”](#) on page 1162.

About redirecting Exchange storage group and database restores

Following are requirements for redirecting Exchange 2000/2003/2007/2010 storage group and database restores:

- The storage groups and databases must already exist on the target server, and must have the same names as the original storage groups or databases.
- The destination server must have the same Organization and Administrative Group name as the source server.
- The destination databases must be configured so that they can be overwritten. See [“Configuring a database in Exchange”](#) on page 1146.

You cannot redirect the restore of the following:

- A version of Exchange server database to a different version of the database. Service packs for both Exchange servers should also be the same.
- Site Replication Service (SRS) and Key Management Service (KMS). These services are dependent on the computer they reside on; redirection to another computer is not supported and could result in the loss of functionality of these services.

Note: KMS is not available in Exchange 2003/2007/2010.

See [“About redirecting Exchange restore data”](#) on page 1159.

See [“Redirecting Exchange restore data”](#) on page 1162.

About redirecting restores of mailboxes or public folders that were backed up separately from the Exchange Information Store

You can redirect the restore of individual mailboxes or public folders from backup sets that were created when individual mailboxes or public folders were backed up from the Microsoft Exchange Mailboxes selection.

Note: Backup Exec does not support this feature for Exchange 2007/2010.

If you select a single mailbox or public folder, or one or more messages from a single mailbox, you can redirect that restore to another existing mailbox or public folder on the same server or to a different server. If any of the folders in the original mailbox do not exist in the destination mailbox, they are created during the restore.

If you select more than one mailbox or public folder, or folders and messages from more than one mailbox, you can only redirect the restore to another server. Mailboxes and public folders with the same names as those selected must already exist on the target server.

Following are requirements for redirecting the restore of mailboxes and public folders:

- If the mailboxes or public folders do not already exist on the target server, you must create them before redirecting the restore. Automatic recreation of mailboxes on the target server is not enabled for redirected restores.
- Ensure that Backup Exec can access mailboxes or public folders on the server that you are redirecting the restore to.
See [“About redirecting restores of mailboxes or public folders that were backed up separately from the Exchange Information Store”](#) on page 1160.
See [“Requirements for accessing Exchange mailboxes”](#) on page 1101.
- More than one mailbox can exist with the same Display Name. When the restore of a mailbox is redirected in Backup Exec, it is redirected to the target mailbox’s Display Name. If a duplicate Display Name exists, then the data may be restored to the wrong mailbox.

To prevent restoring the data to the wrong mailbox, type the name of the target mailbox exactly as it appears when browsing to the mailbox in the backup selections list, including the brackets surrounding the mailbox directory identifier (for example, "Mailbox Name [mailboxname]").

Following are requirements for redirecting the restore of individual messages to another mailbox.

- When redirecting the restore of mailbox data, all destination mailboxes must already exist before the restore begins. The contents of the restored mailboxes are placed in the destination mailboxes.
For example, Mailbox 1 consists of Top of Information Store, Inbox, and Folders 1 and 2, each containing some mail messages. If you back up Mailbox 1 and then you restore Mailbox 1 to the existing Mailbox 2, then all of Mailbox 1, including the Top of Information Store, Inbox, Folders 1 and 2, and messages,

are restored to Mailbox 2. Note that Mailbox 1 itself is not created under Mailbox 2.

If you redirect the restore of Mailbox 1\Top of Information Store\Folder 2 to Mailbox 2, the contents of Mailbox 1\Top of Information Store\Folder 2, Message 5 and Message 6, are placed in Mailbox 2 in the same folder as they were in Mailbox 1, as illustrated in the following graphic.

See [“About redirecting Exchange restore data”](#) on page 1159.

See [“About restoring Exchange mailboxes and public folders from mailbox backups”](#) on page 1152.

See [“Redirecting Exchange restore data”](#) on page 1162.

Enabling access to legacy mailboxes or public folders

Use the following steps to ensure that Backup Exec can access mailboxes or public folders.

See [“About redirecting restores of mailboxes or public folders that were backed up separately from the Exchange Information Store”](#) on page 1160.

To enable access to legacy mailboxes or public folders

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Microsoft Exchange**.
- 3 Check **Enable legacy mailbox support (Exchange 2003)**. **This option is not recommended; use GRT instead.**
- 4 Click **OK**.
- 5 Click the destination server in the backup selections tree.

The option Microsoft Exchange Mailboxes enables you to verify that the target mailbox id displayed in the list of mailboxes.
- 6 If prompted, enter a logon account that stores the credentials of a user account that is unique and has a corresponding mailbox of the same name.

Redirecting Exchange restore data

Use the following steps to redirect Exchange data.

See [“About redirecting Exchange restore data”](#) on page 1159.

See [“About redirecting Exchange storage group and database restores”](#) on page 1160.

See [“About redirecting restores of mailboxes or public folders that were backed up separately from the Exchange Information Store”](#) on page 1160.

To redirect the restore of Exchange data

- 1 Create a restore job.
See [“Requirements for restoring Exchange 2000 or later”](#) on page 1145.
- 2 After selecting options on the **Restore Job Properties** dialog box, on the **Properties** pane, under **Destination**, click **Microsoft Exchange Redirection**.
- 3 Select the appropriate options.
See [“Microsoft Exchange redirection options”](#) on page 1163.
- 4 Start the redirected restore job or select other restore options from the Properties pane.
- 5 After the restore is complete, Symantec recommends that you run a full backup of the restored databases.

Microsoft Exchange redirection options

You can set the following options when you redirect a restore job for Exchange. See [“Redirecting Exchange restore data”](#) on page 1162.

Table G-18 Exchange redirection options

Item	Description
Redirect Exchange sets	Enables the redirection of Exchange backup sets.
Restore to server or Database Availability Group	Specifies the name of the computer or Database Availability Group to which you are restoring, using the format <code>\\server name</code> .
Server logon account	Specifies a Backup Exec logon account that stores the credentials of a Windows user account. The default logon account is displayed by default. To use another logon account, click Change . See “Requirements for accessing Exchange mailboxes” on page 1101.
Redirect using Volume Shadow Copy Service (VSS) snapshot provider	Enables the redirection of backup sets using Volume Shadow Copy Service (VSS) snapshot provider.
Redirect to Storage Group and/or Database (Exchange 2007 or later only)	Enables the Restore to Database and Restore to Storage Group fields.

Table G-18 Exchange redirection options (*continued*)

Item	Description
Restore to Storage Group (Exchange 2007)	Specifies the name of an existing storage group. Use this option when you want to redirect the restore of one or more Exchange 2007 databases to a different storage group.
Restore to Database or Recovery Database	<p>Specifies the name of an Exchange 2007 database or an Exchange Server 2010 database or recovery database that you want to restore.</p> <p>For Exchange 2007, use this option when you want to redirect the restore of a single Exchange 2007 database.</p> <p>When you redirect a database restore, the storage group and database or the recovery database names you specify must already exist on the destination Exchange server.</p>
Redirect to Recovery Storage Group (RSG) (Exchange 2007 only)	<p>Redirects the restore of an Exchange mailbox database to a recovery storage group.</p> <p>You can use a recovery storage group and the Exchange 2003 or 2007 Mailbox Merge Wizard to help you restore individual user mailbox data.</p> <p>To use this option, you must create a recovery storage group. You must also create, within the recovery storage group, a database for each database you plan to restore. Each database you create in the recovery storage group must use the same name as the database that you want to restore. During a redirected restore operation, the Exchange Agent automatically detects and uses the recovery storage group.</p> <p>For more information on recovery storage groups, see your Microsoft Exchange Server 2007 documentation.</p>
Redirect to drive and path (Exchange 2003 and 2007)	Enables the Restore to drive and Restore to path fields.
Restore to drive	Specifies a destination drive on which to restore the Exchange database. Click the ellipsis button to view local and network drives.
Restore to path	Specifies a path to where you want to restore the Exchange database.
Redirect mailboxes or public folders	Enables the options to redirect mailbox backup sets and public folder backup sets.
Redirect mailbox sets	Enables the redirection of mailbox backup sets.

Table G-18 Exchange redirection options (*continued*)

Item	Description
Restore to mailbox	Specifies the name of the mailbox to which you are redirecting this restore. The mailbox must already exist on the target server.
Mailbox logon account	Specifies a logon account that has rights to the target mailbox. To select a logon account, click Change . To clear an existing logon account, click Clear .
Redirect public folder sets	Enables the redirection of public folder backup sets.
Restore to public folder	Specifies the name of the public folder to which you are redirecting this restore. The public folder must already exist on the target server.
Public folder logon account	Specifies a logon account that has rights to the public folder that you are redirecting to. To select a logon account, click Change . To clear an existing logon account, click Clear .

How to prepare for disaster recovery of Exchange Server

A disaster preparation plan is an absolute necessity for restoring Exchange efficiently and effectively in the event of a catastrophic failure. Because Exchange uses Windows security for authentication, disaster recovery of Exchange cannot be separated from the disaster recovery of Windows.

Planning ahead reduces the time needed to recover.

It is critical to build a kit that includes the following items:

- An operating system configuration sheet
- A hard drive partition configuration sheet
- Any RAID configuration
- A hardware configuration sheet
- EISA/MCA configuration disks
- An Exchange configuration sheet
- A Windows emergency repair diskette

To perform the actual recovery, you will need the following items:

- An installed copy of Backup Exec

- The latest full, incremental, and differential backups of the Exchange databases you want to recover. If CPS backups are enabled, you can use recovery points to recover the Exchange database.
- The Microsoft Exchange Server Installation CD
- Any service packs that were applied to the original installation

Recovering from a disaster for Exchange 2000 or later

This procedure guides you through a complete restoration of Exchange using Backup Exec. You should have already performed all the appropriate preparation.

See [“How to prepare for disaster recovery of Exchange Server”](#) on page 1165.

If the Exchange 2000 Server being recovered contains the Site Replication Service (SRS) and/or Key Management Service (KMS), then before you begin the disaster recovery, refer to your Exchange documentation for details on recovering those databases.

Always log on to Windows using the Administrator account (or an Administrator equivalent) during this procedure. Other requirements include:

- The storage groups and databases must already exist on the target server, and have the same names as the original storage groups or databases.
- The destination server must have the same Organization and Administrative Group name as the source server.
- The destination databases must be configured so that they can be overwritten. See [“Configuring a database in Exchange”](#) on page 1146.

You can use Intelligent Disaster Recovery to recover the Exchange server.

See [“Microsoft Exchange recovery notes”](#) on page 1844.

To perform disaster recovery for Exchange 2000 or later

- 1 Recover the Windows server first.

See [“Returning to the last known good configuration”](#) on page 777.

Make sure you restore the Exchange Server 2000 Server or later files that existed on all disk partitions.

When the Windows 2000 server disaster recovery procedure is complete (after the last reboot), you must recover the Exchange server.

- 2 From the Services applet, verify the Microsoft Exchange Information Store service is started.
- 3 Start Backup Exec.

- 4 Catalog the media that contains the latest backup of the Exchange 2000 Server or later storage groups you want to recover.
- 5 On the navigation bar, click the arrow next to Restore.
- 6 Click **New Restore Job**.
- 7 On the Properties pane, under Source, click **Selections**.
- 8 Select the latest full backups of each storage group or database for restore.
If continuous protection is enabled, you can select the backup sets from a full backup or a recovery point.
If the Exchange 2000 Server being recovered contains the Site Replication Service (SRS) and/or Key Management Service (KMS), then select those databases for restore as well.
- 9 If necessary, select all subsequent incremental storage group backups.
If differential backups are to be restored, only the most recent differential storage group backups need to be selected.
- 10 On the Properties pane, under Settings, click **Microsoft Exchange**.
- 11 Click the **Purge existing data and restore only the databases and transaction logs from the backup sets** option.
- 12 In the Path on Exchange Server for temporary storage of log and patch files field, type a location where the associated log and patch files are to be kept until the database is restored.
Make sure the temporary location for log and patch files is empty before you start a restore job. If a restore job fails, check the temporary location (including subdirectories) to make sure any previous log and patch files from a previous restore job were deleted.
- 13 If your selection contains the last backup set to be restored, check **Commit after restore completes**.
Do not check this if you still have backup sets to restore.
If Commit after restore completes is checked when an intermediate backup is being applied, you cannot continue to restore backups and you must restart the restore operation from the beginning.
After the database is restored, the log and patch files in the temporary location are applied to the database, and then the current log files are applied. After the restore is complete, the log and patch files are automatically deleted from the temporary location (including any subdirectories).
- 14 If you want the databases to be immediately available to users after the recovery, check **Mount database after restore**.

- 15** Start the restore job or select other restore options on the Properties pane.
- 16** After completing the restore, it is recommended that a full backup of the restored databases be performed.

Symantec Backup Exec Agent for Microsoft Hyper-V

This appendix includes the following topics:

- [About the Agent for Microsoft Hyper-V](#)
- [About installing the Agent for Microsoft Hyper-V](#)
- [Requirements for using the Agent for Microsoft Hyper-V](#)
- [About upgrading from the Agent for Microsoft Virtual Servers](#)
- [About backup selections for Microsoft Hyper-V](#)
- [Backing up data by using the Agent for Microsoft Hyper-V](#)
- [How Granular Recovery Technology works with the Agent for Microsoft Hyper-V](#)
- [About restore selections for Microsoft Hyper-V](#)
- [Restoring data to the Hyper-V host](#)
- [Restoring a virtual machine to a different host](#)
- [Setting default backup and restore options for the Agent for Microsoft Hyper-V](#)
- [About backing up and restoring highly available virtual machines](#)

About the Agent for Microsoft Hyper-V

The Symantec Backup Exec Agent for Microsoft Hyper-V (Agent for Microsoft Hyper-V) lets you back up and restore the following resources:

- Microsoft Windows Server 2008/2008 R2 Hyper-V hosts.

- All virtual machines that reside on the Hyper-V hosts.
- Clustered Hyper-V hosts, including virtual machines that reside on cluster shared volumes (CSV).

Backup Exec performs a single-pass backup to protect the host configuration data, all virtual machines, and VSS-aware applications that are installed on the virtual machines. Backup Exec's Granular Recovery Technology (GRT) is enabled by default for backup jobs. You can use a GRT-enabled backup to restore individual files and folders from a Windows virtual machine without restoring the entire virtual machine. In addition, you can restore individual items from Microsoft Exchange and Active Directory applications that reside on virtual machines. You can also restore individual databases from Microsoft SQL when it resides on virtual machines.

Note: You must have the appropriate Backup Exec agent for Microsoft Exchange, SQL, or Active Directory on the virtual machine to perform GRT.

Backup Exec can back up virtual machines that are online or that are in an offline state or a saved state. Virtual machines that use Microsoft Windows 2003 (with Hyper-V Integration Services) or later can be backed up while they are online. You can include both online and offline virtual machines in the same backup job. During the backup of an online virtual machine, Backup Exec takes a snapshot backup of the Hyper-V host. The host in turn takes a snapshot of the virtual machines on the host. This process enables Backup Exec to back up virtual servers without any downtime. If an online backup cannot be performed, then an offline backup is performed. With an offline backup, the virtual machine is placed briefly in a saved state. However, the virtual machine does not remain in the saved state for the entire backup job.

The amount of downtime for a saved state backup job depends on the following:

- The amount of memory that is allocated to the virtual machine.
- The current load on the host's operating system.

See [“Requirements for using the Agent for Microsoft Hyper-V”](#) on page 1171.

See [“Backing up data by using the Agent for Microsoft Hyper-V”](#) on page 1175.

See [“Restoring data to the Hyper-V host”](#) on page 1182.

About installing the Agent for Microsoft Hyper-V

The Symantec Backup Exec Agent for Microsoft Hyper-V (Agent for Microsoft Hyper-V) is installed as a separate, add-on component of Backup Exec. The Agent

for Microsoft Hyper-V is installed on the Microsoft Hyper-V host. If your Backup Exec media server is also your Microsoft Hyper-V host, you can install the Agent for Microsoft Hyper-V when you install Backup Exec. Or, you can install it after Backup Exec has been installed.

If Backup Exec is not installed on your Microsoft Hyper-V host, you must push-install the Remote Agent for Windows Systems to your Microsoft Hyper-V host. You do not need to install the Agent for Microsoft Hyper-V on virtual machines. However, a license key is required on the media server for the Agent for Microsoft Hyper-V. The Remote Agent for Windows Systems is included with the Agent for Microsoft Hyper-V.

See “[Installing a custom installation of Backup Exec](#)” on page 119.

See “[Installing additional Backup Exec options to the local media server](#)” on page 123.

See “[Push-installing Backup Exec to remote computers](#)” on page 126.

See “[Push-installing the Remote Agent and Advanced Open File Option to remote computers](#)” on page 134.

Requirements for using the Agent for Microsoft Hyper-V

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

The following items are required to use the Agent for Microsoft Hyper-V:

Table H-1 Requirements for the Agent for Microsoft Hyper-V

Software	Installed on
Microsoft Windows Server 2008 Hyper-V	Microsoft Hyper-V host
Backup Exec	Backup Exec media server
Agent for Microsoft Hyper-V	Microsoft Hyper-V host

Table H-1 Requirements for the Agent for Microsoft Hyper-V (*continued*)

Software	Installed on
VHDMount	Media server (if the media server is not the virtual server). Note: VHDMount is required only if the media server runs Microsoft Windows 2003 or Windows 2008 without the Hyper-V role installed. You can install the VHDMount component from Microsoft Virtual Server 2005 R2 SP1.

To run an online backup, the following requirements must be met:

- Microsoft Windows Server 2008/2003 SP2/Vista SP1/XP SP3 is installed on the virtual machine.
- Hyper-V Integration Services with Backup (Volume snapshot) is installed.
- The virtual machine is in a running state.

If these conditions are not met, the virtual machine is placed in a saved state if it is running. If the virtual machine is turned off, then that virtual machine is backed up only if you select the option **Back up virtual machines that are powered off**.

To enable Backup Exec to collect catalog data for Microsoft Exchange, Active Directory, and SQL on the virtual machine, the following items are required:

- A licensed version of the Backup Exec agent for the application (Agent for Microsoft Exchange, Agent for Microsoft SQL, or Agent for Microsoft Active Directory).
 A license is required for each application on each virtual machine. For example, you need five licenses for installation of Microsoft SQL on five virtual machines.
- The Remote Agent for Windows Systems must be installed on the virtual machine.
 The Agent for Microsoft Hyper-V includes a license for the Remote Agent for Windows Systems. The agents for Microsoft Exchange, Active Directory, and SQL also include a license for the Remote Agent for Windows Systems. No separate license is required for the Remote Agent for Windows Systems.
- The virtual machine must be capable of being backed up online.
- The credentials that you use to access the virtual machine must also have access to the application.

The Remote Agent for Windows Systems must be installed on the virtual machine to do the following:

- Enable individual files and folders to be restored back to the original virtual machine.
- Enable individual SQL databases to be restored back to the original virtual machine.
- Enable individual Exchange items to be restored back to the original virtual machine.
- Enable individual Active Directory objects to be restored back to the original virtual machine.

About upgrading from the Agent for Microsoft Virtual Servers

If you set up recurring jobs with Backup Exec 12, you must either recreate the job or change the selection list to use Microsoft virtual servers.

Backup Exec is not intended to be a tool to migrate from Microsoft Virtual Server to Microsoft Hyper-V. For information about how to migrate, see Microsoft's virtual machine migration guide.

[http://technet.microsoft.com/en-us/library/dd296684\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd296684(WS.10).aspx)

About backup selections for Microsoft Hyper-V

The following containers appear on the administration console's Selection pane for Hyper-V.

Table H-2 Microsoft Hyper-V backup selections

Container name	Items in the container	What is included in the backup job
Microsoft Hyper-V	This item includes Initial Store and Virtual Machines .	If you select the Microsoft Hyper-V container for backup, the backup job includes the application configuration settings and all virtual machines.
Initial Store	This item includes the virtual server application configuration settings.	If you select Initial Store for backup, the backup job includes a single XML file that contains the Hyper-V authorization configuration.

Table H-2 Microsoft Hyper-V backup selections (*continued*)

Container name	Items in the container	What is included in the backup job
Virtual Machines	<p>This item includes each virtual machine that resides on the virtual server.</p> <p>Note: When you select an individual virtual machine, the files that are on that virtual machine appear in the results pane. However, you cannot select individual files to include in or exclude from the backup.</p>	<p>If you select an individual virtual machine, the backup is a full image backup of the entire virtual machine, which includes the following items:</p> <ul style="list-style-type: none"> ■ .vhd files ■ .avhd files ■ Differencing disks ■ Hyper-V managed snapshots

How Backup Exec automatically protects new virtual machines during a backup job

Backup Exec's dynamic inclusion feature protects new virtual machines and folders that are found when a backup job runs. If new virtual machines are added between the time when the backup job is created and when the backup job runs, Backup Exec automatically backs up the new virtual machines. Because the backup job may include new virtual machines, the job may require more storage space and more time to run than you anticipated. The job history shows the number of virtual machines that were backed up.

In the backup selection list, dynamic inclusion is enabled for the following Hyper-V nodes:

- Microsoft Hyper-V
- Virtual Machines under Microsoft Hyper-V
- The Hyper-V host node

If you select the host node, then dynamic inclusion is enabled automatically for the Microsoft Hyper-V node.
- Microsoft Hyper-V HA Virtual Machines
- The cluster name node

If you select the cluster name node, then dynamic inclusion is enabled automatically for the Microsoft Hyper-V HA Virtual Machines node.

Backing up data by using the Agent for Microsoft Hyper-V

When you create a backup job for Microsoft Hyper-V, Full is the only available backup method. Even though a full image backup is created, Granular Recovery Technology (GRT) enables individual files and folders to be restored. GRT is enabled by default for individual files and folders on virtual machines and for individual items from VSS-aware applications that reside on virtual machines. VSS-aware applications include Microsoft Exchange, SQL, and Active Directory. By default, Backup Exec uses the resource credentials of the parent virtual machine.

Note: Only the files that reside on the virtual server are backed up. Virtual machines that have remote .vhd files are excluded from the backup job. You can use the Remote Agent for Windows Systems and the appropriate Backup Exec agent to protect virtual machines that have remote .vhd files.

To back up data by using the Agent for Microsoft Hyper-V

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Properties** pane, under **Source**, click **Selections**.
- 4 Select the resources you want to back up.
See [“About backup selections for Microsoft Hyper-V”](#) on page 1173.
- 5 In the **Properties** pane, under **Settings**, click **Microsoft Hyper-V**.
- 6 Select the appropriate options for this backup job.
See [“Microsoft Hyper-V backup options”](#) on page 1175.
- 7 To change the setting for granular recovery for VSS-aware applications that are installed on virtual machines, click **Edit**.
See [“Virtual Machine Application Granular Recovery Technology Settings”](#) on page 1176.
- 8 Start the backup job or select other backup options from the **Properties** pane.

Microsoft Hyper-V backup options

You can set the following options for each backup job that you create for Microsoft Hyper-V.

See [“Backing up data by using the Agent for Microsoft Hyper-V”](#) on page 1175.

Table H-3 Microsoft Hyper-V backup options

Item	Description
Backup method	Displays the Full backup method. Full is the only available backup method.
Exclude virtual machines that must be put in a saved state to back up	Excludes from the backup all offline virtual machines that do not support online backups and that are in a running state when the backup begins.
Back up virtual machines that are powered off	Enables Backup Exec to back up virtual machines that are turned off.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from Virtual Machines	Enables individual files and folders to be restored from the full backup. You must install the Remote Agent for Windows Systems on the virtual machine on which you want to restore the data. The Remote Agent for Windows Systems does not have to be installed on the virtual machine to back up the data.
Edit	Lets you change the GRT settings for Microsoft Active Directory, Exchange, and SQL.
Microsoft Active Directory	Indicates whether GRT is enabled or disabled for Microsoft Active Directory on the virtual machine. It is enabled by default.
Microsoft Exchange	Indicates whether GRT is enabled or disabled for Microsoft Exchange on the virtual machine. It is enabled by default.
Microsoft SQL	Indicates whether GRT is enabled or disabled for Microsoft SQL on the virtual machine. It is enabled by default.

Virtual Machine Application Granular Recovery Technology Settings

Use the following options to enable or disable granular recovery of individual items from Microsoft Active Directory, Exchange, and SQL.

Note: If you enable or disable granular recovery technology (GRT) for one of the following applications, the setting applies to both VMware and Hyper-V virtual machines. If you do not want to use the same settings, Symantec recommends that you set up separate backup jobs for each type of virtual machine.

See [“Backing up data by using the Agent for Microsoft Hyper-V”](#) on page 1175.

See [“How Backup Exec protects Microsoft Exchange, SQL, and Active Directory data on virtual machines”](#) on page 1178.

See [“Requirements for protecting Microsoft Exchange, SQL, and Active Directory data on virtual machines”](#) on page 1179.

Table H-4 Virtual Machine Application Granular Recovery Technology Settings

Item	Description
Enable GRT for Microsoft Active Directory objects on virtual machines	Enables Backup Exec to collect the information that is required to restore individual Active Directory objects on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Active Directory is installed.
Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines	Enables Backup Exec to collect the information that is required to restore individual Exchange databases and mailbox items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Exchange is installed.
Enable GRT for Microsoft SQL (database-level only) on virtual machines	Enables Backup Exec to collect the information that is required to restore individual SQL database items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which SQL is installed.

How Granular Recovery Technology works with the Agent for Microsoft Hyper-V

Backup Exec Granular Recovery Technology (GRT) lets you restore individual files and folders without having to restore the entire virtual machine. It also lets

you restore individual items from VSS-aware applications that are installed on virtual machines.

GRT is not intended to be used for system recovery. However, you can perform a complete system recovery by selecting the entire virtual machine as a restore selection in a restore job.

See [“How Backup Exec protects Microsoft Exchange, SQL, and Active Directory data on virtual machines”](#) on page 1178.

You should review the requirements for a GRT-enabled backup before you configure it.

See [“About requirements for jobs that use Granular Recovery Technology”](#) on page 325.

See [“Requirements for protecting Microsoft Exchange, SQL, and Active Directory data on virtual machines”](#) on page 1179.

To use GRT, you must select the individual files and folders that you want to restore from the list that appears when you expand the Netbios name or the computer name of the virtual machine. You cannot select individual folders and files from the virtual machines that appear when you expand the **Virtual Machines** node.

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“About restore selections for Microsoft Hyper-V”](#) on page 1180.

How Backup Exec protects Microsoft Exchange, SQL, and Active Directory data on virtual machines

Backup Exec can restore individual items from the following VSS-aware applications that reside on virtual machines:

Table H-5 Types of data that Backup Exec protects for VSS-aware applications on virtual machines

Application	Types of data that Backup Exec protects
Microsoft Exchange	Mailboxes, individual messages, calendar items, tasks, journal entries, and public folder data (disk-backups only)
Microsoft SQL	Databases
Microsoft Active Directory	Individual user accounts, printer objects, sites, and organizational units

When you create a backup job, Backup Exec automatically locates VSS-aware applications on virtual machines. During the backup job, Backup Exec backs up the data from the VSS-aware applications by using Granular Recovery Technology (GRT). By default, Backup Exec enables GRT using the same credentials that were used to connect to the virtual machine. You can disable GRT for any of the VSS-aware application types.

Note: If you enable or disable GRT for Microsoft Exchange, SQL, or Active Directory, the setting applies to both VMware virtual machines and Hyper-V virtual machines. If you do not want to use the same settings, Symantec recommends that you set up separate backup jobs for each type of virtual machine.

Note: Backup Exec supports the granular recovery of individual Exchange and SQL items only in non-clustered and non-distributed configurations.

During the backup job, Backup Exec collects metadata for the applications. If Backup Exec is unable to collect the metadata, then you cannot restore individual items for the applications. However, the backup job may otherwise complete successfully.

Backup Exec cannot collect metadata in the following situations:

- GRT is disabled for an application.
- Backup Exec cannot connect to the virtual machine.
- Incorrect credentials were entered for the virtual machine.

Note: Backup Exec uses the Microsoft Hyper-V writer during backups of VSS-aware applications on virtual machines. The Microsoft Hyper-V writer truncates application logs before data is moved to the storage device. Therefore, the application logs for the applications on the virtual machines are truncated if you use Microsoft Hyper-V.

See [“Requirements for protecting Microsoft Exchange, SQL, and Active Directory data on virtual machines”](#) on page 1179.

Requirements for protecting Microsoft Exchange, SQL, and Active Directory data on virtual machines

Backup Exec can back up and restore individual items from VSS-aware applications that are installed on virtual machines.

The following items are required to protect data for Microsoft Exchange, SQL, and Active Directory on virtual machines:

- The virtual machine must be turned on.
- You must enter the appropriate credentials for the virtual machine. Ensure that the credentials for the virtual machine allow access to the VSS-aware applications.
- The media server must be able to connect to the virtual machine using the network name or IP address.
- The Backup Exec Remote Agent for Windows Systems must be installed on the virtual machine.
- The correct number of licenses must be entered for the applications that you want to protect on the virtual machines.
- The operating system on the virtual machine must support VSS.

See [“How Backup Exec protects Microsoft Exchange, SQL, and Active Directory data on virtual machines”](#) on page 1178.

About restore selections for Microsoft Hyper-V

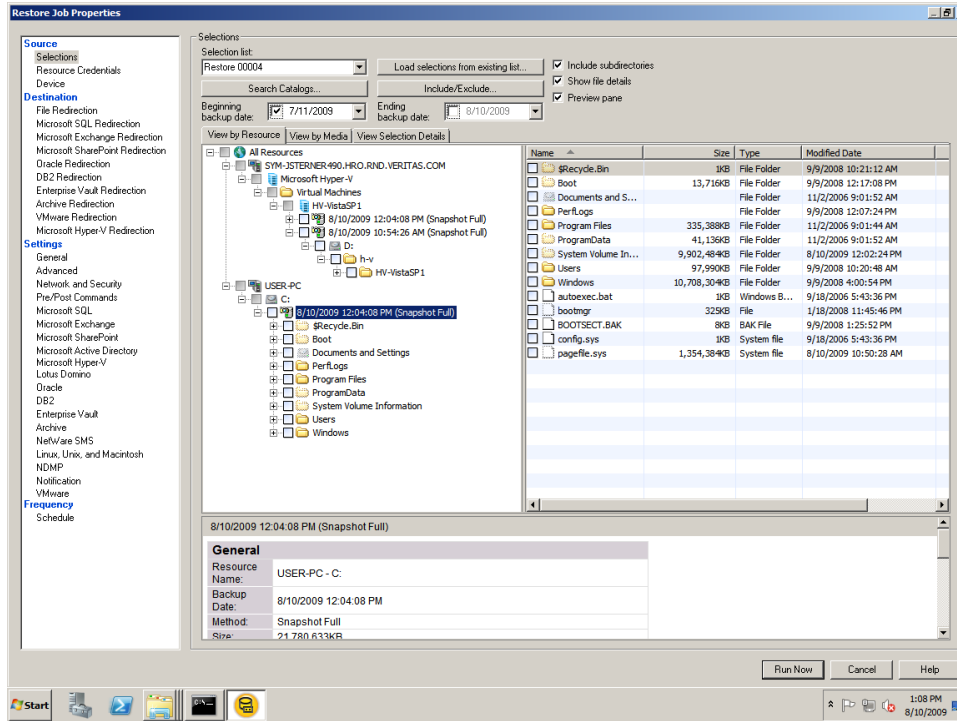
You can restore data from virtual machines in the following ways:

- Restore a complete virtual machine for disaster recovery purposes.
- Restore individual files or folders that were backed up from the virtual machine (if you selected the Granular Recovery Technology (GRT) option for the backup job).

Note: Linux virtual machines must be restored in their entirety at the .vhd level.


In the **Restore** view, a virtual machine that was backed up with GRT enabled appears under its NetBIOS name or computer name. If you expand the name, then individual files and folders appear.

Figure H-1 Restore with GRT enabled



Virtual machines also appear in the **Restore** view under **Virtual Machines**. In this view, either the display name appears or the name that you provided for the virtual machine during its creation appears. If you expand the display name for a virtual machine, its contents appear. If you select the virtual machine by its display name, you can recover the entire virtual machine.

Figure H-2 Restore without GRT enabled



Restoring data to the Hyper-V host

Follow these steps to restore configuration information to the Hyper-V host or to restore virtual machines to their original host.

If you want to restore a virtual machine to a different Hyper-V host, you must use the redirection feature.

See [“Restoring a virtual machine to a different host”](#) on page 1184.

To restore data to the Hyper-V host

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Select the items that you want to restore.

See [“About restore selections for Microsoft Hyper-V”](#) on page 1180.

- 4 In the **Properties** pane, under **Settings**, click **Microsoft Hyper-V**.

- 5 Select the appropriate options for this restore job.
See [“Microsoft Hyper-V restore options”](#) on page 1183.
- 6 Start the restore job or select other restore options from the **Properties** pane.

Microsoft Hyper-V restore options

You can set the following options for each restore job that you create for Microsoft Hyper-V.

See [“Restoring data to the Hyper-V host”](#) on page 1182.

Table H-6 Microsoft Hyper-V restore options

Item	Description
Overwrite powered on virtual machines	Restores virtual machines that are turned on. By default, the virtual machine is turned off before the restore job processes and the virtual machine is overwritten. To prevent virtual machines that are turned on from being overwritten, uncheck this option. If this option is unchecked and a virtual machine is running, the job fails. You must manually turn off the virtual machine before you attempt to run the restore job again.
Do not turn on the virtual machine	Leaves the virtual machine turned off after the restore job completes.
Turn on the virtual machine and resume from the available saved state	Turns on the virtual machine automatically after the restore job completes. The virtual machine resumes operations from the saved state from the time of the backup. Note: This option applies only to virtual machines that are backed up using a saved state. Virtual machines that are backed up online do not have a saved state.

Table H-6 Microsoft Hyper-V restore options (*continued*)

Item	Description
Turn on the virtual machine and discard available saved state	<p>Turns on the virtual machine automatically after the restore job completes. The virtual machine discards the available saved state.</p> <p>Note: This option applies only to virtual machines that are backed up using a saved state. Virtual machines that are backed up online do not have a saved state.</p>

Restoring a virtual machine to a different host

You can restore a virtual machine to a different Microsoft Hyper-V server. You also can redirect flat files from the virtual machine to any computer that has a Remote Agent for Windows Systems installed.

To restore a virtual machine to a different host

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Select the virtual machine that you want to restore.
 See [“About restore selections for Microsoft Hyper-V”](#) on page 1180.
- 4 Select the appropriate options.
 See [“Microsoft Hyper-V Redirection options”](#) on page 1184.
- 5 Start the restore job or select other restore options from the **Properties** pane.
- 6 In the **Properties** pane, under **Destination**, click **Microsoft Hyper-V Redirection**.

Microsoft Hyper-V Redirection options

You can set the following options when you restore a virtual machine to a different Microsoft Hyper-V host.

See [“Restoring a virtual machine to a different host”](#) on page 1184.

Table H-7 Microsoft Hyper-V Redirection options

Item	Description
Redirect Hyper-V sets	Restores a virtual machine to a different location.
Restore to server	Specifies the name of the virtual server where you want to restore the data.
Server logon account	Specifies the logon account for the virtual server where the data is being restored.
Restore to drive	Specifies the destination for the restored data. You can browse to local drives and network drives.
Restore to path	<p>Specifies the target path on the device that is listed in the Restore to Drive field. To retain the original directory structure, make sure that the Preserve tree option is selected in the Restore Job Properties - Settings - General dialog box.</p> <p>See “General options for restore jobs” on page 612.</p> <p>If the Preserve tree option is not selected, all of the data is restored to the path that is designated in this field.</p>
Redirect to a different Hyper-V host and register the virtual machine	Redirects the restored data from the virtual machine to another host. The entire virtual machine is restored.
Redirect to a folder	<p>Redirects file sets. The restore recreates the folder hierarchy that is associated with each file.</p> <p>You can restore these files to one of the following locations:</p> <ul style="list-style-type: none">■ The same virtual server■ Another virtual server■ An external hard drive

Setting default backup and restore options for the Agent for Microsoft Hyper-V

You can use the default options that Backup Exec sets during installation for all Microsoft Hyper-V backup and restore jobs. Or, you can choose your own default settings. You can change these options for individual jobs.

To set default backup and restore options for the Agent for Microsoft Hyper-V

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Settings**, click **Microsoft Hyper-V**.
- 3 Select the appropriate options.
See “[Microsoft Hyper-V default options](#)” on page 1186.
- 4 Click **OK**.

Microsoft Hyper-V default options

You can use the default options that Backup Exec sets during installation for all Microsoft Hyper-V backup and restore jobs. Or, you can choose your own default settings.

See “[Setting default backup and restore options for the Agent for Microsoft Hyper-V](#)” on page 1186.

Table H-8 Microsoft Hyper-V default options

Item	Description
Exclude virtual machines that are in a saved state at backup time	Excludes from the backup all offline virtual machines that do not support online backups and that are in a running state when the backup begins.
Back up virtual machines that are powered off	Enables Backup Exec to back up virtual machines that are turned off.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	Enables individual files and folders to be restored from the full backup. You must install the Remote Agent for Windows Systems on the virtual machine on which you want to restore the data. The Remote Agent for Windows Systems does not have to be installed on the virtual machine to back up the data.

Table H-8 Microsoft Hyper-V default options (continued)

Item	Description
Edit	Lets you change the GRT settings for Microsoft Active Directory, Exchange, and SQL.
Microsoft Active Directory	Indicates whether GRT is enabled or disabled for Microsoft Active Directory on the virtual machine. It is enabled by default.
Microsoft Exchange	Indicates whether GRT is enabled or disabled for Microsoft Exchange on the virtual machine. It is enabled by default.
Microsoft SQL	Indicates whether GRT is enabled or disabled for Microsoft SQL on the virtual machine. It is enabled by default.
Overwrite powered on virtual machines	<p>Restores virtual machines that are turned on.</p> <p>By default, the virtual machine is turned off before the restore job processes and the virtual machine is overwritten. To prevent virtual machines that are turned on from being overwritten, uncheck this option. If this option is unchecked and a virtual machine is running, the job fails. You must manually turn off the virtual machine before you attempt to run the restore job again.</p>
Do not turn on the virtual machine	Leaves the virtual machine turned off after the restore job completes.
Turn on the virtual machine and resume from the available saved state	<p>Turns on the virtual machine automatically after the restore job completes. The virtual machine resumes operations from the saved state from the time of the backup.</p> <p>Note: This option applies only to virtual machines that are backed up using a saved state. Virtual machines that are backed up online do not have a saved state.</p>

Table H-8 Microsoft Hyper-V default options (continued)

Item	Description
Turn on the virtual machine and discard available saved state	Turns on the virtual machine automatically after the restore job completes. The virtual machine discards the available saved state. Note: This option applies only to virtual machines that are backed up using a saved state. Virtual machines that are backed up online do not have a saved state.

About backing up and restoring highly available virtual machines

When virtual machines are configured for high availability, they are moved to a new node in the backup selection list. Clustered virtual machines appear under the name of the cluster, in the **Highly Available Hyper-V Machines** node. Non-clustered virtual machines remain in the **Microsoft Hyper-V** node. When you make a backup selection, Backup Exec checks for highly available virtual machines. If highly available virtual machines are discovered, Backup Exec reminds you to select those virtual machines for backup.

The restore selection list is similar to the backup selection list. Clustered virtual machines appear under the name of the cluster. Non-clustered virtual machines appear under the **Microsoft Hyper-V** node. You can restore a highly available virtual machine in the same way you restore any other virtual machine. The virtual machine maintains its high availability. However, if you redirect the restore to another Hyper-V host, then the virtual machine is no longer highly available when the restore job completes. You must reconfigure the virtual machine to be highly available.

Symantec Backup Exec Agent for Microsoft SharePoint

This appendix includes the following topics:

- [About the SharePoint Agent](#)
- [Requirements for the SharePoint Agent](#)
- [About installing the SharePoint Agent](#)
- [Adding a SharePoint server farm to the backup selections list](#)
- [How Backup Exec automatically protects new SharePoint resources during a backup job](#)
- [Changing the name of a SharePoint server farm](#)
- [Deleting a farm from the Microsoft SharePoint Server Farms node](#)
- [Disabling or enabling communication between a SharePoint Web server and Backup Exec](#)
- [Setting default options for SharePoint Server 2003 and later](#)
- [About using the SharePoint Agent with SharePoint Server 2010 and Windows SharePoint Foundation 2010](#)
- [About using the SharePoint Agent with SharePoint Server 2007 and Windows SharePoint Services 3.0](#)
- [About using the SharePoint Agent with SharePoint Portal Server 2003 and Windows SharePoint Services 2.0](#)

- [Microsoft SharePoint backup options](#)
- [Microsoft SharePoint restore options](#)
- [Microsoft SharePoint redirection options](#)

About the SharePoint Agent

The Symantec Backup Exec Agent for Microsoft SharePoint (SharePoint Agent) is an optional, add-on component to Backup Exec. The SharePoint Agent enables network administrators to perform backup and restore operations on Microsoft SharePoint installations that are connected to a network. SharePoint backups can be integrated with network backups without separate administration or dedicated hardware.

The SharePoint Agent supports installations of the following items:

- SharePoint Portal Server 2003
- SharePoint Server 2007
- SharePoint Server 2010
- Windows SharePoint Services 2.0 and 3.0
- SharePoint Foundation 2010

See [“About using the SharePoint Agent with SharePoint Portal Server 2003 and Windows SharePoint Services 2.0”](#) on page 1235.

See [“About using the SharePoint Agent with SharePoint Server 2007 and Windows SharePoint Services 3.0”](#) on page 1218.

See [“About using the SharePoint Agent with SharePoint Server 2010 and Windows SharePoint Foundation 2010”](#) on page 1200.

Requirements for the SharePoint Agent

The SharePoint Agent has the following requirements:

- The SharePoint Agent must be installed on the media server.
- The Backup Exec Remote Agent for Windows Systems (Remote Agent) must be installed on each remote SharePoint Server that will be protected. In addition, for SharePoint Server 2003/2007/2010, the Remote Agent must be installed on all servers in the server farm.

- The credentials specified by the logon account used for backing up and restoring SharePoint Server data must have local administrative rights on the servers where SharePoint components are installed.
For more information on granting permissions on folders in the workspace or backward-compatible document libraries, see your SharePoint Server documentation.
- The credentials specified by the logon account used for backing up and restoring the Single Sign-on database must be either the account name or a member of the group that is specified in the "Account name" field in the Single Sign-on Settings section of the Manage Settings for Single Sign-on administration page in SharePoint Server.
- The logon account that you use to restore content into an existing site collection must have appropriate rights to create objects in that site collection. If you restore into a site collection that does not exist, the logon account becomes the primary site collection owner.
- Internet Information Services (IIS) rights can affect database backups and restores. Ensure that the logon account used for backup and restore has rights to access the IIS sites. Integrated Windows Security should be enabled within the IIS rights.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

About installing the SharePoint Agent

The SharePoint Agent must be installed on the media server.

See “[Installing additional Backup Exec options to the local media server](#)” on page 123.

See “[Push-installing the Remote Agent and Advanced Open File Option to remote computers](#)” on page 134.

Adding a SharePoint server farm to the backup selections list

If you create a shortcut to a SharePoint server farm as user-defined selection, it appears in the backup selections list. Any SharePoint server farms that publish to Backup Exec also appear in the backup selections list. If a farm that you want to back up does not display under the **Microsoft SharePoint Server Farms** node

on the **Backup Job Properties** dialog box, you can manually add that farm to the list.

Backup Exec adds the new server farm to the **Microsoft SharePoint Server Farms** node under **Backup Exec Agents** and contacts the specified Web server to retrieve the remainder of the farm topology.

When you create jobs to protect the SharePoint resources for the server farm, make backup selections from this server farm node. In addition, back up the default Microsoft SQL databases (master, model, msdb, pubs) for each Microsoft SQL instance that hosts SharePoint databases.

Please note that after a server farm is added to Microsoft SharePoint Server Farms, the SharePoint databases hosted on Microsoft SQL instances can no longer be selected for backup directly from the Microsoft SQL Server resource nodes.

If you change the SharePoint server farm topology after it has been added to Microsoft SharePoint Server Farms, you must browse the server farm node so that Backup Exec can recognize and save the changes.

Backup Exec's dynamic inclusion feature automatically protects any new resources that were added after a backup job was created.

See [“How Backup Exec automatically protects new SharePoint resources during a backup job”](#) on page 1194.

Note: In a CASO environment, you must add the server farm to each managed media server individually. You should repeat this procedure for each managed media server on which you want the server farm to appear in the backup selections list.

To add a SharePoint server farm to the backup selections list

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the selection tree, right-click **Microsoft SharePoint Server Farms** and then click **Add Server Farm**.
- 4 In the **Web server name** field, type the name of a Web server that belongs to the farm you want to add.

- 5 In the **Server farm name** field, type a name for the farm or use the default name. The name you type here will display under the **Microsoft SharePoint Server Farms** node. Backup Exec lists the names of discovered Web servers in parentheses after the server farm name in the Selection tree.

The following characters cannot be used in farm names: ‘ ~ ^ * () { } \ ; : ’ ” , < > / ?
- 6 Click **OK**.

Add Server Farm options

You can add a Microsoft SharePoint server farm to the backup selections list manually.

See [“Adding a SharePoint server farm to the backup selections list”](#) on page 1191.

Table I-1 Add Server Farm options

Item	Description
Web server name	Specifies the Web server name that belongs to the server farm.
Server farm name	Specifies the name of the server farm. The name you enter here displays under Microsoft SharePoint Server Farms in the backup selections list.

Manage SharePoint Server Farms options

You can edit the properties of any server farm in the backup selections list. You can also delete a server farm from the backup selections list if you no longer use it.

See [“Disabling or enabling communication between a SharePoint Web server and Backup Exec”](#) on page 1196.

See [“Changing the name of a SharePoint server farm”](#) on page 1195.

See [“Deleting a farm from the Microsoft SharePoint Server Farms node”](#) on page 1195.

Table I-2 Manage SharePoint Server Farms options

Item	Description
Delete	Deletes a server farm from Microsoft SharePoint Server Farms in the backup selections list.
Properties	Lets you edit the properties of a server farm.

Server Farm Properties

You can prevent Backup Exec from communicating with one or more Web servers in a server farm when Backup Exec attempts to retrieve the server farm topology.

See [“Disabling or enabling communication between a SharePoint Web server and Backup Exec”](#) on page 1196.

Table I-3 Server Farm Properties options

Item	Description
Server farm name	Specifies the server farm name that contains a Web server for which you want to disable communications with Backup Exec.
Web server	Check the Web servers for which you want to disable communications with Backup Exec. To enable communications, clear the check box.

How Backup Exec automatically protects new SharePoint resources during a backup job

Backup Exec's dynamic inclusion feature automatically protects any new resources that were added after a backup job was created. If Backup Exec discovers that you added a new resource as a child to a protected resource, it automatically backs up the new resource. Because the backup job may include new resources, the job may require more storage space and more time to run than you anticipated.

See [“Adding a SharePoint server farm to the backup selections list”](#) on page 1191.

See [“Creating a backup job by setting job properties”](#) on page 332.

Changing the name of a SharePoint server farm

When Backup Exec adds a farm, it creates a default name for the farm. You can change the default farm name to a name that is meaningful to you.

Note: If you change the name of a farm that is already protected, the restore view contains entries for both the original name and the new name. Changing the name of a farm does not change it in existing restore selections.

To change the name of a farm

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the selection tree, right-click **Microsoft SharePoint Server Farms**, and then click **Manage Server Farms**.
- 4 Select the farm whose name you want to change, and then click **Properties**.
- 5 Type the new farm name in the **Server farm name** field. Then name you type here will display under the **Microsoft SharePoint Server Farms** node. Backup Exec lists the names of discovered Web servers in parentheses after the server farm name in the Selection tree.

The following characters cannot be used in farm names: ‘ ~ ^ * () { } \ ; : ' " , < > / ?

- 6 Click **OK**, and then click **Close**.

Deleting a farm from the Microsoft SharePoint Server Farms node

If a server farm is no longer in use or is no longer valid, you can remove it from the **Microsoft SharePoint Server Farms** node.

Note: If Backup Exec is installed on the same server that is used as a Web server in a farm, you cannot delete that farm.

To delete a farm from the Microsoft SharePoint Server Farms node

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.

- 3 In the selection tree, right-click **Microsoft SharePoint Server Farms**, and then click **Manage Server Farms**.
- 4 Select the farm you want to delete, and then click **Delete**.

Disabling or enabling communication between a SharePoint Web server and Backup Exec

Backup Exec communicates with the Web servers that participate in SharePoint server farms to discover the farm topology. This process may take some time if Backup Exec attempts to communicate with a Web server that is unavailable. If you know that a particular Web server in a farm will be unavailable for a period of time, you can disable the communication between the Web server and Backup Exec.

To disable or enable communication between a Web server and Backup Exec

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the selection tree, right-click **Microsoft SharePoint Server Farms**, and then click **Manage Server Farms**.
- 4 Select the farm that contains the Web server you want to disable or enable, and then click **Properties**.
- 5 To prevent Backup Exec from communicating with a Web server, clear the check box next to that Web server's name. If the Web server is now available to communicate with Backup Exec, select the check box next to the Web server's name.

When you disable communication with a Web server, Backup Exec removes the name of that Web server from the server farm name under the **Microsoft SharePoint Server Farms** node.

- 6 Click **OK**, and then click **Close**.

Setting default options for SharePoint Server 2003 and later

You can set default options to use for all backup and restore jobs for SharePoint Server 2003 and later.

To set default options for SharePoint Server 2003 and later

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Microsoft SharePoint**.
- 3 Select the appropriate options.

See “[Microsoft SharePoint default options](#)” on page 1197.

Microsoft SharePoint default options

You can set default options to use for all backup and restore jobs for SharePoint Server 2003 and later.

See “[Setting default options for SharePoint Server 2003 and later](#)” on page 1196.

Table I-4 Microsoft SharePoint default options

Item	Description
Backup method	<p>Allows you to select one of the following as the default backup method:</p> <ul style="list-style-type: none">■ Full - Back up entire database Backs up the entire database.■ Differential - Back up database changes only Backs up only the changes that were made to the database since the last full backup. The differential backup method cannot be used to back up Index databases or Document Libraries. The Full backup method must be used to back up these resources.■ Log - Back up and truncate transaction log Backs up the data that is contained in the transaction log. This method does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated). If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery model before a log backup is run. Alternatively, you can run full backups only, or run full and differential backups of the SQL databases that are associated with SharePoint.

Table I-4 Microsoft SharePoint default options (*continued*)

Item	Description
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual documents from the database backup (available for full backups only)	<p>Enables the restore of individual documents from database backups. This option is only available when performing full backups. The option is not available if Differential - Back up database changes only has been selected as the backup method. You must have a current version of the Remote Agent for Windows Systems on the SharePoint server when you run the GRT-enabled backup job.</p> <p>See “How to restore individual items by using Granular Recovery Technology” on page 321.</p>
Perform consistency check before backup of Microsoft SQL databases used by Microsoft SharePoint	<p>Runs a full consistency check (including indexes) of the Microsoft SQL databases used by Microsoft SharePoint before you back up the database.</p>
Continue with backup if consistency check fails	<p>Continues with the backup operation even if the consistency check fails.</p>
Bring restored databases online	<p>Brings the databases online after a restore job.</p>
Reconnect previous database links	<p>Re-establishes the link between the restored databases and their corresponding sites when you restore Web applications or Windows SharePoint Services sites.</p>
Preserve existing Internet Information Services (IIS) Web site and application pool (SharePoint 2007 and 2010 only)	<p>Preserves the web site and application pool for the SharePoint Web application that you restore if it already exists in IIS. If you do not check this option, the Web site and application pool are deleted from IIS during the restore. After they are deleted, they are recreated in the default location that SharePoint specifies. This option is for SharePoint 2007 and 2010 only.</p>
Replace existing database	<p>Specifies that you want to replace existing databases on the specified SQL server instance with the databases being restored. Otherwise, an error occurs if databases with the same names as those being restored exist on the specified SQL server instance.</p>

Table I-4 Microsoft SharePoint default options (*continued*)

Item	Description
If versioning is enabled on the restore destination	<p>Allows you to choose from the following options if versioning is enabled on the destination to which you want to restore an individual item:</p> <ul style="list-style-type: none"> ■ Add as a new version Backup Exec restores the existing item as a new version, which makes it the most recent version of the existing item. ■ Skip if the item exists Backup Exec does not restore the item if an identical item exists in the restore destination. Backup Exec notes that the file was skipped in the job log. ■ Restore over existing items Backup Exec replaces the existing item with the restored item.
If versioning is not enabled on the restore destination	<p>Allows you to choose from the following options if versioning is not enabled on the destination to which you want to restore an individual item:</p> <ul style="list-style-type: none"> ■ Skip if the item exists Backup Exec does not restore the item if an identical item exists in the restore destination. Backup Exec notes that the file was skipped in the job log. ■ Restore over existing items Backup Exec replaces the existing item with the restored item.
Restore only the most recent version of an item	Restores only the most recent version of an item.
Include security information	<p>Restores any applicable security information with the item. You can restore different levels of security based on the SharePoint item you restore:</p> <ul style="list-style-type: none"> ■ Sites - User and SharePoint Group information and security ACL are restored for top-level sites ■ Sub-sites - Security ACL is restored ■ Lists - Security ACL and other security-related information is restored ■ List items - Granular security information is restored to individual list items for Microsoft SharePoint Server 2007 and Windows SharePoint Services 3.0 only

About using the SharePoint Agent with SharePoint Server 2010 and Windows SharePoint Foundation 2010

The Symantec Backup Exec Agent for Microsoft SharePoint (SharePoint Agent) includes support for Microsoft Office SharePoint Server 2010 and Windows SharePoint Foundation 2010.

SharePoint 2010 offers new metadata features including tags, social bookmarks, and content ratings. These new types of metadata are stored in the service applications that reside outside of the content database. For example, enterprise managed tags reside in the Managed Metadata Service application. You should make sure to back up all of your service applications to ensure that all metadata is protected.

Any metadata that is stored outside of the content database cannot be restored using Granular Recovery Technology (GRT). You can, however, use GRT to restore SharePoint data with metadata attached to it. As long as the metadata resides in the same service application, SharePoint maintains the link between the data and the metadata.

Backup Exec provides a hierarchical tree view of SharePoint resources in the **Backup Selections** pane.

In the **Backup Selections** pane, **Microsoft SharePoint Server Farms** shows a logical view of the topology of each server farm on your network.

Figure I-1 Microsoft Office SharePoint 2010 example (Backup Selections pane - View by Resource tab)



See [“Backing up a farm for Microsoft Office SharePoint Server 2010 or Windows SharePoint Foundation 2010”](#) on page 1202.

See [“Backing up individual SharePoint 2010 Web applications in a Microsoft SharePoint server farm”](#) on page 1203.

See [“About restoring resources for SharePoint Server 2010 and SharePoint Foundation 2010”](#) on page 1203.

See [“Redirecting a restore job for SharePoint 2010”](#) on page 1208.

See [“Redirecting the restore of individual SharePoint 2010 objects to a file path”](#) on page 1210.

See [“Redirecting the restore of a Microsoft Office SharePoint Server 2010 Web application”](#) on page 1211.

See [“Redirecting the restore of a Microsoft Office SharePoint Server 2010 database to an alternate SQL instance”](#) on page 1213.

See [“Recovering SharePoint 2010 data after a disaster”](#) on page 1214.

Backing up a farm for Microsoft Office SharePoint Server 2010 or Windows SharePoint Foundation 2010

Before you can back up a Microsoft SharePoint Server 2010 server farm, you must add the farm to the backup selections list.

See [“Adding a SharePoint server farm to the backup selections list”](#) on page 1191.

Use the following steps to back up a Microsoft Office SharePoint Server 2010 farm or a Windows SharePoint Foundation 2010 farm.

To back up a farm for Microsoft Office SharePoint Server 2010 or Windows SharePoint Foundation 2010

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Backup Selections** pane, expand **Microsoft SharePoint Server Farms**.
- 4 Expand a server farm that contains the Microsoft Office SharePoint Server 2010 components that you want to back up.
- 5 Select the SharePoint resources that you want to back up.
See [“Selections options for backup jobs”](#) on page 336.
- 6 In the **Backup Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 7 Select the appropriate options.
See [“Microsoft SharePoint backup options ”](#) on page 1246.
- 8 In the **Properties** pane, select other backup options as needed.
- 9 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See [“Scheduling jobs”](#) on page 356.

Backing up individual SharePoint 2010 Web applications in a Microsoft SharePoint server farm

Use the following steps to back up individual Microsoft Office SharePoint Server 2010 Web applications.

To back up individual SharePoint 2010 Web applications in a Microsoft SharePoint server farm

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Backup Selections** pane, expand **Microsoft SharePoint Server Farms**.
- 4 Expand the server farm that contains the SharePoint Web application that you want to back up.
- 5 Expand **Microsoft SharePoint Foundation Web Application**.
- 6 Select the Web applications that you want to back up.
- 7 In the **Backup Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 8 Select the appropriate options.
See "[Microsoft SharePoint backup options](#)" on page 1246.
- 9 In the **Backup Job Properties** pane, select other backup options as needed.
- 10 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See "[Scheduling jobs](#)" on page 356.

About restoring resources for SharePoint Server 2010 and SharePoint Foundation 2010

You can restore the following SharePoint Server 2010 and SharePoint Foundation 2010 resources:

- Web applications and their associated databases.

- Individual documents and pictures that are contained in libraries.
- Sites and sub-sites.
Individual objects and their versions can be restored from full database backups.
- Lists and list items.
Individual objects and their versions can be restored from full database backups.
See the Microsoft SharePoint documentation for more information about lists and list items.
- Configuration database.
A configuration database contains all of the configuration information for the entire SharePoint Server farm. Use caution when restoring this database. Any changes that you make to the farm topology before you restore from the backup are lost. The configuration database can only be restored to its original location.
- Service applications.

See [“Restoring resources for SharePoint Server 2010 and SharePoint Foundation 2010”](#) on page 1204.

See [“Restoring individual SharePoint 2010 objects from full database backups to their original locations”](#) on page 1205.

See [“Restoring a Microsoft Office SharePoint Server 2010 Web application to its original location”](#) on page 1207.

Restoring resources for SharePoint Server 2010 and SharePoint Foundation 2010

You can restore SharePoint Server 2010 and SharePoint Foundation 2010 resources.

See [“About restoring resources for SharePoint Server 2010 and SharePoint Foundation 2010”](#) on page 1203.

To restore resources for SharePoint Server 2010 and SharePoint Foundation 2010

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Expand the server farm that contains the SharePoint components that you want to restore.
- 4 Expand the SharePoint resources that you want to restore.
- 5 Select the backup sets for the SharePoint resources that you want to restore.

- 6 In the **Restore Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 7 Select the appropriate options.
See "[Microsoft SharePoint restore options](#)" on page 1247.
- 8 In the **Restore Job Properties** pane, select other restore options as needed.
- 9 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.
See "[Scheduling jobs](#)" on page 356.

Restoring individual SharePoint 2010 objects from full database backups to their original locations

You can restore individual documents, images, sites, sub-sites, lists, and list items from full SharePoint database backup jobs if you selected the following option during the backup job:

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual documents from the database backup (available for full backups only)

See "[About restoring resources for SharePoint Server 2010 and SharePoint Foundation 2010](#)" on page 1203.

See "[How to restore individual items by using Granular Recovery Technology](#)" on page 321.

Note: If you restore into an existing site collection, the logon account that you use must have appropriate rights to create objects in that site collection. If you restore into a site collection that does not exist, the logon account becomes the primary site collection owner.

The option to enable the restore of individual objects is not available for differential backup jobs.

To restore individual objects from full database backups to their original locations

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 On the **View by Resource** tab, expand the server farm that contains the Web application where the individual objects you want to restore reside.
- 4 Expand **Microsoft SharePoint Foundation Web Application**.
- 5 Expand the Web application that contains the content database from which you want to restore objects.
- 6 Expand the content database that contains the objects you want to restore.
- 7 Expand the backup set that contains the objects you want to restore.
- 8 Expand the content database.
- 9 Expand the folder that contains the objects you want to restore.
- 10 In the **Results** pane, select the objects that you want to restore.
- 11 In the **Restore Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 12 Do one of the following:

If versioning is enabled on the restore destination

Select one of the following options:

- **Add as a new version**
Backup Exec restores the existing item as a new version, making it the most recent version of the existing item.
- **Skip if the item exists**
Backup Exec does not restore the item if an identical item exists in the restore destination. Backup Exec notes that the file was skipped in the job log.
- **Restore over existing items**
Backup Exec replaces the existing item with the restored item.

If versioning is not enabled on the restore destination Select one of the following options:

- **Skip if the item exists**
Backup Exec does not restore the item if an identical item exists in the restore destination. Backup Exec notes that the file was skipped in the job log.
- **Restore over existing items**
Backup Exec replaces the existing item with the restored item.

13 In the **Restore Job Properties** pane, select other restore options as needed.

14 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.
See "[Scheduling jobs](#)" on page 356.

Restoring a Microsoft Office SharePoint Server 2010 Web application to its original location

Use the following steps to restore a Microsoft Office SharePoint Server 2010 Web application to its original location.

See "[About restoring resources for SharePoint Server 2010 and SharePoint Foundation 2010](#)" on page 1203.

Note: When you restore a Microsoft Office SharePoint Server 2010 Web application, all objects that are contained in the Web application's content databases are overwritten.

To restore a Microsoft Office SharePoint Server 2010 Web application to its original location

- 1** On the navigation bar, click the arrow next to **Restore**.
- 2** Click **New Restore Job**.

- 3 Expand the server farm that contains the Web application that you want to restore.
- 4 Expand **Microsoft SharePoint Foundation Web Application**.
- 5 Expand the Web application that you want to restore.
- 6 Expand the content database, and then select the backup set that contains the content database that you want to restore.

Note: If the Web application contains multiple content databases, Symantec recommends that you expand the other content databases and select the corresponding backup sets for those databases as well.

- 7 In the **Restore Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 8 Check **Bring restored databases online**.
- 9 Check **Reconnect previous database links**.
- 10 In the **Restore Job Properties** pane, select other restore options as needed.
- 11 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
 - Set the scheduling options.
 - Click **Submit**.
- See "[Scheduling jobs](#)" on page 356.

Redirecting a restore job for SharePoint 2010

Follow these steps to redirect a restore job to an existing Web application on a Web server in a farm.

To bring the databases online after you complete the redirected restore job, verify that **Bring restored databases online** and **Reconnect previous database links** are selected in the Microsoft SharePoint settings. When you restore Web applications or Windows SharePoint Services sites, these options also re-establish the link between the restored databases and their corresponding sites.

Note: If you restore full or differential backup sets in separate restore jobs, clear these options for all jobs except the last job. You should select these options for the last restore job in the sequence. You may be prompted to insert any media that you have already used.

To redirect a restore job for SharePoint 2010

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Select the SharePoint resources you want to restore.
See “[Selections options for restore jobs](#)” on page 609.
You can restore the Configuration database and Single Sign-on database back to the original location only.
- 4 On the **Properties** pane, under **Destination**, select **Microsoft SharePoint Redirection**.
- 5 Check **Redirect Microsoft SharePoint sets**.
- 6 Click **SharePoint 2003 portal sites or SharePoint 2007/2010 web applications**.
- 7 In the **URL or web application name** field, type the URL of the site or the web application name to which you want to restore the data.
For example: http://portalsite1 or https://portalsite1
- 8 In the **Front-end web server name** field, type the name of the Web server on which the site resides.
You must create the target SharePoint Server 2010 site or Windows SharePoint Services site on the specified Web server with the same database structure as the source site before you run the restore job.
- 9 Do one of the following:
 - Use the default logon account as indicated.
 - Click **Change** to select a different logon account.
- 10 Set additional restore options on the **Properties** pane as needed.
- 11 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
 - Set the scheduling options.
 - Click **Submit**.
- See “[Scheduling jobs](#)” on page 356.

See “[About using the SharePoint Agent with SharePoint Server 2010 and Windows SharePoint Foundation 2010](#)” on page 1200.

Redirecting the restore of individual SharePoint 2010 objects to a file path

You can redirect the restore of file-based SharePoint data that has been uploaded to a document library or is attached to list items. Examples of file-based SharePoint data include documents and images. Individual objects and their versions can only be redirected to a file location, not another content database. Use the following steps to restore individual objects to a file path.

Note: You cannot restore SharePoint security information when you restore an item to a file path.

To redirect individual SharePoint 2010 objects to a file path

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 On the **View by Resource** tab, expand the server farm that contains the Web application where the individual objects you want to restore are located.
- 4 Expand **Microsoft SharePoint Foundation Web Application**.
- 5 Expand the Web application that contains the content database from which you want to restore objects.
- 6 Expand the content database that contains the objects you want to restore.
- 7 Expand the backup set that contains the objects you want to restore.
- 8 Expand the content database.
- 9 Expand the folder that contains the objects you want to restore.
- 10 In the **Results** pane, select the objects that you want to restore.

- 11 In the **Restore Job Properties** pane, under **Destination**, click **Microsoft SharePoint Redirection**.
- 12 Check **Redirect Microsoft SharePoint sets**.
- 13 Click **Individual SharePoint sites, documents, lists, or items**.
- 14 Click **Redirect to path**.
- 15 In the **Restore to drive or UNC path** and **Restore to path** fields, enter the drive letter and path to which you want to direct the restore.
Use the following format for a UNC path \\servername\share.
- 16 Use the default logon account as indicated, or click **Change** to select a different one.
- 17 In the **Restore Job Properties** pane, select other restore options as needed.
- 18 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See "[Scheduling jobs](#)" on page 356.

See "[About using the SharePoint Agent with SharePoint Server 2010 and Windows SharePoint Foundation 2010](#)" on page 1200.

Redirecting the restore of a Microsoft Office SharePoint Server 2010 Web application

Before you can redirect the restore of a Microsoft Office SharePoint Server 2010 Web application, the SharePoint software must be installed on the destination server.

To redirect the restore of a Microsoft Office SharePoint Server 2010 Web application

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Expand the server farm that contains the Web application that you want to restore.
- 4 Expand **Microsoft SharePoint Foundation Web Application**.

- 5 Expand the Web application that contains the content database that you want to restore.
- 6 Expand the content database, and then select the backup set that contains the content database that you want to restore.

If the Web application contains multiple content databases, expand the other content databases and select the corresponding backup sets for those databases as well. You must restore all content databases for the Web application together.

- 7 In the **Restore Job Properties** pane, under **Destination**, click **Microsoft SharePoint Redirection**.
- 8 Check **Redirect Microsoft SharePoint sets**.
- 9 Click **SharePoint 2003 portal sites or SharePoint 2007/2010 web applications**.
- 10 In the **URL or web application name** field, enter the name of the Web application that you want to restore. You also can enter the Web application's URL.

Use the following format: <Web application name> or http://production1.

The target Web application must exist and it must be configured with the same number of content databases as the original Web application.

- 11 In the **Front-end web server name** field, enter the name of the Microsoft IIS server that hosts the Web server. You also can enter the Web server's IP address.
- 12 Use the default logon account as indicated, or click **Change** to select a different one.
- 13 In the **Restore Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 14 Check **Bring restored databases online**.
- 15 Check **Reconnect previous database links**.
- 16 In the **Restore Job Properties** pane, select other restore options as needed.
- 17 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See “[Scheduling jobs](#)” on page 356.

See “[About using the SharePoint Agent with SharePoint Server 2010 and Windows SharePoint Foundation 2010](#)” on page 1200.

Redirecting the restore of a Microsoft Office SharePoint Server 2010 database to an alternate SQL instance

You can restore SharePoint databases to an alternate SQL instance. For example, you can use SharePoint administrative tools to harvest granular information or attach a database to a new Web application.

To redirect the restore of a Microsoft Office SharePoint Server 2010 database to an alternate SQL instance

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Navigate to and select the SharePoint database that you want to restore.
- 4 On the **Properties** pane, under **Destination**, click **Microsoft SharePoint Redirection**.
- 5 Check **Redirect Microsoft SharePoint sets**.
- 6 Click **Restore Sharepoint databases to an alternate SQL instance**.
- 7 In the **Restore to SQL server\instance** field, enter the SQL server and instance to which you want to redirect the SharePoint database.
- 8 In the **Restore to database** field, enter the name of the database to which you want to redirect the SharePoint database.
- 9 Select **Replace existing database** to overwrite an existing database with the same name.

If you do not select this field and a database with the same name exists on the SQL instance, you receive an error.
- 10 Use the default logon account as indicated, or click **Change** to select a different one.

11 On the **Properties** pane, select other job properties that might be appropriate for your environment.

12 Do one of the following:

To run the job now

Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
 - Set the scheduling options.
 - Click **Submit**.
- See “[Scheduling jobs](#)” on page 356.

See “[About using the SharePoint Agent with SharePoint Server 2010 and Windows SharePoint Foundation 2010](#)” on page 1200.

Recovering SharePoint 2010 data after a disaster

You can use the SharePoint Agent to recover a SharePoint server after a hard drive failure. Before you recover SharePoint data, you must recover the SharePoint server's operating system.

You can use Backup Exec's Intelligent Disaster Recovery Option or you can manually recover the server's operating system.

See “[About the Intelligent Disaster Recovery Option](#)” on page 1806.

See “[About manual disaster recovery of Windows computers](#)” on page 780.

After the Windows server is recovered, you can recover SharePoint 2010 data. Complete the actions in the table in successive order to recover the SharePoint data.

Table I-5 To recover SharePoint 2010 data after a disaster

Step	Action	Notes
Step 1	Recover the master database and the model database for the SQL instances that SharePoint uses. You must perform this step if you manually recovered the server's operating system.	Skip this step if you used Backup Exec's Intelligent Disaster Recovery Option to recover the server's operating system. See “ About recovering SQL manually ” on page 1312.

Table I-5 To recover SharePoint 2010 data after a disaster (*continued*)

Step	Action	Notes
Step 2	Inventory the media to be recovered.	See “Inventorying media in a device” on page 448.
Step 3	Catalog the media to be recovered.	See “Creating a new catalog” on page 248.
Step 4	Restore the msdb databases for any SQL instances that SharePoint uses.	<p>Select the backup sets that contain the msdb databases for the SQL instances that SharePoint uses.</p> <p>Configure the following Microsoft SQL restore options:</p> <ul style="list-style-type: none"> ■ Select Leave the database ready to use; additional transaction logs or differential backups cannot be restored. ■ Select Overwrite the existing database. <p>See “Restoring data by setting job properties” on page 605.</p>
Step 5	Restore all Web applications.	<p>Select the backup sets for all SharePoint Web applications.</p> <p>Select the following Microsoft SharePoint restore option:</p> <p>Replace existing databases</p> <p>See “Restoring a Microsoft Office SharePoint Server 2010 Web application to its original location” on page 1207.</p>

Table I-5 To recover SharePoint 2010 data after a disaster (*continued*)

Step	Action	Notes
Step 6	Restore Shared Services Applications databases.	<p>Restore the following Shared Services Applications databases:</p> <ul style="list-style-type: none"> ■ Business Data Connectivity Service ■ Managed Metadata Service ■ PerformancePoint Service Application ■ Search Service Application ■ Secure Store Service ■ User Profile Service Application ■ Web Analytics Service Application ■ Word Automation Services ■ Services\State Services\Service DB 1 <p>Select the following Microsoft SharePoint restore option:</p> <p>Replace existing databases</p> <p>See “Restoring data by setting job properties” on page 605.</p> <p>Note: Some of the remaining restore jobs may fail because communication with the SharePoint server has not been fully established yet. This behavior is expected. Proceed with the recovery process until all steps are complete.</p>

Table I-5 To recover SharePoint 2010 data after a disaster (*continued*)

Step	Action	Notes
Step 7	Restore search services.	<p>Restore the following services:</p> <ul style="list-style-type: none"> ■ SharePoint Foundation Help Search\Search Instance\Index Files 1 ■ Search-DB 1 <p>Select the following Microsoft SharePoint restore option:</p> <p>Replace existing databases</p> <p>See “Restoring data by setting job properties” on page 605.</p> <p>Note: You may get a message in the Job Log to restart your computer, you can ignore this.</p>
Step 8	Restore the SharePoint Configuration V4-DB resource.	<p>Select the backup sets for the SharePoint Configuration V4-DB resource.</p> <p>Configure the following Microsoft SharePoint restore options:</p> <ul style="list-style-type: none"> ■ Select Bring restored databases online. ■ Select Reconnect previous database links. ■ Select Replace existing databases. <p>See “Restoring data by setting job properties” on page 605.</p>
Step 9	Restart the SharePoint server.	<p>After the restore job is complete, restart the SharePoint server. Then proceed to the next step.</p>

Table I-5 To recover SharePoint 2010 data after a disaster (*continued*)

Step	Action	Notes
Step 10	Restore remaining SharePoint resources.	Select the backup sets for the SharePoint Global Settings resources, if necessary. See “Restoring data by setting job properties” on page 605.
Step 11	Back up the SharePoint server.	When the disaster recovery is complete, Symantec recommends that you perform a backup job as soon as possible. See “Creating a backup job by setting job properties” on page 332.

About using the SharePoint Agent with SharePoint Server 2007 and Windows SharePoint Services 3.0

The Symantec Backup Exec Agent for Microsoft SharePoint (SharePoint Agent) includes support for Microsoft Office SharePoint Server 2007 and Windows SharePoint Services 3.0.

Backup Exec provides a hierarchical tree view of SharePoint resources in the **Backup Selections** pane.

In the **Backup Selections** pane, Microsoft SharePoint Server Farms shows a logical view of the topology of each server farm on your network.

Figure I-2 Microsoft Office SharePoint Server 2007 example (Backup Selections pane - View by Resource tab)



Figure I-3 Mixed SharePoint versions example (Backup Selections pane - View by Resource tab)



About adding a SharePoint 2007 server farm to the backup selections list

Before you can back up a Microsoft Office SharePoint Server 2007 server farm, you must add the farm to the backup selections list.

See [“Adding a SharePoint server farm to the backup selections list”](#) on page 1191.

Backing up a farm for Microsoft Office SharePoint Server 2007 or a Windows SharePoint Services 3.0

Use the following steps to back up a Microsoft Office SharePoint Server 2007 or a Windows SharePoint Services 3.0 farm.

To back up a farm for Microsoft Office SharePoint Server 2007 or Windows SharePoint Services 3.0

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Backup Selections** pane, expand **Microsoft SharePoint Server Farms**.
- 4 Expand a server farm that contains the Microsoft Office SharePoint Server 2007 components that you want to back up.
- 5 Select the SharePoint resources you want to back up.
See [“Selections options for backup jobs”](#) on page 336.
- 6 In the **Backup Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.

- 7 Select the appropriate options.
See [“Microsoft SharePoint backup options”](#) on page 1246.
- 8 In the **Properties** pane, select other backup options as needed.
- 9 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See [“Scheduling jobs”](#) on page 356.

Backing up individual SharePoint 2007 Web applications in a Microsoft SharePoint server farm

Use the following steps to back up individual Microsoft Office SharePoint Server 2007 Web applications.

To back up individual SharePoint 2007 Web applications in a Microsoft SharePoint server farm

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Backup Selections** pane, expand **Microsoft SharePoint Server Farms**.
- 4 Expand the server farm that contains the SharePoint Web application that you want to back up.
- 5 Expand the Windows SharePoint Services Web Application that contains the Web applications that you want to back up.
- 6 Select the Web applications that you want to back up.
- 7 In the **Backup Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 8 Select the appropriate options.
See [“Microsoft SharePoint backup options”](#) on page 1246.
- 9 In the **Backup Job Properties** pane, select other backup options as needed.
- 10 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See [“Scheduling jobs”](#) on page 356.

About restoring resources for SharePoint Server 2007 and SharePoint Services 3.0

You can restore the following SharePoint Server 2007 and SharePoint Services 3.0 resources:

- Web applications and their associated databases. Symantec recommends that you restore all Web application databases together to preserve the topology.
- Individual documents that are contained in libraries.
- Sites and sub-sites.
Individual objects and their versions can be restored from full database backups.
- Lists and list items.
Individual items and their versions can be restored from full database backups. See the Microsoft SharePoint documentation for more information about lists and list items.
- Configuration database. A configuration database contains all of the configuration information for the entire SharePoint Server farm. Use caution when restoring this database. Any changes that you make to the farm topology before you restore from the backup are lost. The configuration database can only be restored to its original locations.
- Single Sign-on databases. Single Sign-on databases can only be restored to their original locations.
- Shared Service Providers. A Shared Service Provider is a grouping of shared services and related shared resources.

See [“Restoring resources for SharePoint Server 2007 and SharePoint Services 3.0”](#) on page 1222.

See [“Restoring individual SharePoint 2007 objects from full database backups to their original locations”](#) on page 1223.

See “[Restoring a Microsoft Office SharePoint Server 2007 Shared Services Provider](#)” on page 1224.

See “[Restoring a Microsoft Office SharePoint Server 2007 Web application to its original location](#)” on page 1225.

Restoring resources for SharePoint Server 2007 and SharePoint Services 3.0

You can restore SharePoint Server 2007 and SharePoint Services 3.0 resources.

See “[About restoring resources for SharePoint Server 2007 and SharePoint Services 3.0](#)” on page 1221.

To restore resources for SharePoint Server 2007 and SharePoint Services 3.0

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Expand the server farm that contains the SharePoint components that you want to restore.
- 4 Expand the SharePoint resources that you want to restore.
- 5 Select the backup sets for the SharePoint resources that you want to restore.
- 6 In the **Restore Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 7 Select the appropriate options.
See “[Microsoft SharePoint restore options](#)” on page 1247.
- 8 In the **Restore Job Properties** pane, select other restore options as needed.
- 9 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See “[Scheduling jobs](#)” on page 356.

Restoring individual SharePoint 2007 objects from full database backups to their original locations

You can restore individual documents, images, sites, sub-sites, lists, and list items from full SharePoint database backup jobs if you selected the following option during the backup job:

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from the database backup (available for full backups only)

See [“About restoring resources for SharePoint Server 2007 and SharePoint Services 3.0”](#) on page 1221.

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

The option to enable the restore of individual documents is not available for differential backup jobs.

To restore individual objects from full database backups to their original locations

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 On the **View by Resource** tab, expand the server farm that contains the Web application where the individual documents you want to restore reside.
- 4 Expand **Windows SharePoint Services Application**.
- 5 Expand the Web application that contains the content database from which you want to restore objects.
- 6 Expand the content database that contains the objects you want to restore.
- 7 Expand the backup set that contains the objects you want to restore.
- 8 Expand the content database.
- 9 Expand the folder that contains the objects you want to restore.
- 10 In the **Results** pane, select the objects that you want to restore.
- 11 In the **Restore Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 12 Do one of the following:

If versioning is enabled on the restore destination

Select one of the following options:

■ **Add as a new version**

Backup Exec restores the existing item as a new version, making it the most recent version of the existing item.

■ **Skip if the item exists**

Backup Exec does not restore the item if an identical item exists in the restore destination. Backup Exec notes that the file was skipped in the job log.

■ **Restore over existing items**

Backup Exec replaces the existing item with the restored item.

If versioning is not enabled on the restore destination

Select one of the following options:

■ **Skip if the item exists**

Backup Exec does not restore the item if an identical item exists in the restore destination. Backup Exec notes that the file was skipped in the job log.

■ **Restore over existing items**

Backup Exec replaces the existing item with the restored item.

13 In the **Restore Job Properties** pane, select other restore options as needed.

14 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See [“Scheduling jobs”](#) on page 356.

Restoring a Microsoft Office SharePoint Server 2007 Shared Services Provider

You can restore a Microsoft Office SharePoint Server 2007 Shared Services Provider.

See [“About restoring resources for SharePoint Server 2007 and SharePoint Services 3.0”](#) on page 1221.

Use the following steps to restore a Microsoft Office SharePoint Server 2007 Shared Services Provider.

To restore a Microsoft Office SharePoint Server 2007 Shared Services Provider

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Expand the server farm that contains the Shared Service Provider that you want to restore.
- 4 Expand the node for the Shared Services Provider that you want to restore.
- 5 Select the backup sets for all of the components of the Shared Service Provider that you want to restore.

Symantec recommends that you restore all Shared Service Provider components together.

- 6 In the **Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 7 Select the **Bring restored databases online** check box.
- 8 Select the **Reconnect previous database links** check box.
- 9 In the **Restore Job Properties** pane, select other restore options as needed.
- 10 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See [“Scheduling jobs”](#) on page 356.

Restoring a Microsoft Office SharePoint Server 2007 Web application to its original location

You can restore a Microsoft Office SharePoint Server 2007 Web application to its original location.

See [“About restoring resources for SharePoint Server 2007 and SharePoint Services 3.0”](#) on page 1221.

Use the following steps to restore a Microsoft Office SharePoint Server 2007 Web application to its original location.

Note: When you restore a Microsoft Office SharePoint Server 2007 Web application, all documents that are contained in the Web application's content databases are overwritten.

To restore a Microsoft Office SharePoint Server 2007 Web application to its original location

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Expand the server farm that contains the Web application that you want to restore.
- 4 Expand Windows SharePoint Services Application.
- 5 Expand the Web application that you want to restore.
- 6 Expand the content database, and then select the backup set that contains the content database that you want to restore.

If the Web application contains multiple content databases, expand the other content databases and select the corresponding backup sets for those databases as well.
- 7 In the **Restore Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 8 Check **Bring restored databases online**.
- 9 Check **Reconnect previous database links**.
- 10 In the **Restore Job Properties** pane, select other restore options as needed.
- 11 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See "[Scheduling jobs](#)" on page 356.

Redirecting a restore job for SharePoint 2007

Follow these steps to redirect a restore job to an existing site on a Web server in a farm.

To bring the databases online after you complete the redirected restore job, verify that the **Bring restored databases online** and **Reconnect previous database** links options are selected in the Microsoft SharePoint settings. When you restore Web applications or Windows SharePoint Services sites, these options also re-establish the link between the restored databases and their corresponding sites.

Note: If you restore full or differential backup sets in separate restore jobs, clear these options for all jobs except the last job. You should select these options for the last restore job in the sequence. You may be prompted to insert any media that you have already used.

To redirect a restore job for SharePoint 2007

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Select the SharePoint resources you want to restore.
See [“Selections options for restore jobs”](#) on page 609.
You can restore the Configuration database and Single Sign-on database back to the original location only.
- 4 On the **Properties** pane, under **Destination**, select **Microsoft SharePoint Redirection**.
- 5 Check **Redirect Microsoft SharePoint sets**.
- 6 Click **SharePoint 2003 portal sites or SharePoint 2007 web applications**.
- 7 In the **URL or web application name** field, type the URL of the site to which you want to restore the data.
For example: `http://portalsite1` or `https://portalsite1`
- 8 In the **Front-end web server name** field, type the name of the Web server on which the site resides.
You must create the target SharePoint Server 2007 site or Windows SharePoint Services site on the specified Web server with the same database structure as the source site before you run the restore job.
- 9 Do one of the following:
 - Use the default logon account as indicated.

- Click **Change** to select a different logon account.
- 10 Set additional restore options on the **Properties** pane, or start the restore job.

Redirecting the restore of individual SharePoint 2007 objects to a file path

You can redirect the restore of SharePoint file-based data such as documents and images that have been uploaded to a document library or are attached to list items. Individual items and their versions can only be redirected to a file location, not another content database. Use the following steps to restore individual objects to a file path.

Note: You cannot restore SharePoint security information when you restore an item to a file path.

To redirect individual SharePoint 2007 objects to a file path

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **View by Resource** tab, expand the server farm that contains the Web application where the individual objects you want to restore reside.
- 4 Expand Windows SharePoint Services Application.
- 5 Expand the Web application that contains the content database from which you want to restore objects.
- 6 Expand the content database that contains the objects you want to restore.
- 7 Expand the backup set that contains the objects you want to restore.
- 8 Expand the content database.
- 9 Expand the folder that contains the objects you want to restore.
- 10 In the **Results** pane, select the objects that you want to restore.
- 11 In the **Restore Job Properties** pane, under **Destination**, click **Microsoft SharePoint Redirection**.
- 12 Check **Redirect Microsoft SharePoint sets**.
- 13 Click **Individual SharePoint sites, documents, lists, or items**.
- 14 Click **Redirect to path**.

- 15 In the **Restore to drive or UNC path** and **Restore to path** fields, enter the drive letter and path to which you want to direct the restore.
Use the following format for a UNC path `\\servername\share`.
- 16 Use the default logon account as indicated, or click **Change** to select a different one.
- 17 In the **Restore Job Properties** pane, select other restore options as needed.
- 18 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
 - Set the scheduling options.
 - Click **Submit**.
- See "[Scheduling jobs](#)" on page 356.

Redirecting the restore of a Microsoft Office SharePoint Server 2007 Web application

Before you can redirect the restore of a Microsoft Office SharePoint Server 2007 Web application, the SharePoint software must be installed on the destination server.

To redirect the restore of a Microsoft Office SharePoint Server 2007 Web application

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Expand the server farm that contains the Web application that you want to restore.
- 4 Expand **Windows SharePoint Services Application**.
- 5 Expand the Web application that contains the content database that you want to restore.
- 6 Expand the content database, and then select the backup set that contains the content database that you want to restore.

If the Web application contains multiple content databases, expand the other content databases and select the corresponding backup sets for those databases as well. You must restore all content databases for the Web application together.

- 7 In the **Restore Job Properties** pane, under **Destination**, click **Microsoft SharePoint Redirection**.
- 8 Check **Redirect Microsoft SharePoint sets**.
- 9 Click **SharePoint 2003 portal sites or SharePoint 2007 web applications**.
- 10 In the **URL or web application name** field, enter the name of the Web application that you want to restore. You also can enter the Web application's URL.

Use the following format: <Web application name> or http://production1.

The target Web application must exist and it must be configured with the same number of content databases as the original Web application.
- 11 In the **Front-end web server name** field, enter the name of the Microsoft IIS server that hosts the Web server. You also can enter the Web server's IP address.
- 12 Use the default logon account as indicated, or click **Change** to select a different one.
- 13 In the **Restore Job Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 14 Check the **Bring restored databases online** check box.
- 15 Check the **Reconnect previous database links** check box.
- 16 In the **Restore Job Properties** pane, select other restore options as needed.
- 17 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See "[Scheduling jobs](#)" on page 356.

Redirecting the restore of a Microsoft Office SharePoint Server 2007 database to an alternate SQL instance

You can restore SharePoint databases to an alternate SQL instance. For example, you can attach a database to a new Web application.

To redirect the restore of a Microsoft Office SharePoint Server 2007 database to an alternate SQL instance

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Navigate to and select the SharePoint database that you want to restore.
- 4 On the **Properties** pane, under **Destination**, click **Microsoft SharePoint Redirection**.
- 5 Check **Redirect Microsoft SharePoint sets**.
- 6 Click **Restore Sharepoint databases to an alternate SQL instance**.
- 7 In the **Restore to SQL server\instance** field, enter the SQL server and instance to which you want to redirect the SharePoint database.
- 8 In the **Restore to database** field, enter the name of the database to which you want to redirect the SharePoint database.
- 9 Select **Replace existing database** to overwrite an existing database with the same name.

If you do not select this field and a database with the same name exists on the SQL instance, you receive an error.
- 10 Use the default logon account as indicated, or click **Change** to select a different one.
- 11 On the **Properties** pane, select other job properties that might be appropriate for your environment.
- 12 Start the restore job.

See [“About using the SharePoint Agent with SharePoint Server 2007 and Windows SharePoint Services 3.0”](#) on page 1218.

Recovering SharePoint 2007 data after a disaster

You can use the SharePoint Agent to recover a SharePoint server after a hard drive failure. Before you recover SharePoint data, you must recover the SharePoint server's operating system.

You can use Backup Exec's Intelligent Disaster Recovery Option or you can manually recover the server's operating system.

See [“About the Intelligent Disaster Recovery Option”](#) on page 1806.

See [“About manual disaster recovery of Windows computers”](#) on page 780.

After the Windows server is recovered, you can recover SharePoint 2007 data. Complete the actions in the table in successive order to recover the SharePoint data.

Table I-6 To recover SharePoint 2007 data after a disaster

Step	Action	Notes
Step 1	Recover the master database and the model database for the SQL instances that SharePoint uses. You must perform this step if you manually recovered the server's operating system.	Skip this step if you used Backup Exec's Intelligent Disaster Recovery Option to recover the server's operating system. See "About recovering SQL manually" on page 1312.
Step 2	Inventory the media to be recovered.	See "Inventorying media in a device" on page 448.
Step 3	Catalog the media to be recovered.	See "Creating a new catalog" on page 248.
Step 4	Restore the msdb databases for any SQL instances that SharePoint uses.	Select the backup sets that contain the msdb databases for the SQL instances that SharePoint uses. Configure the following Microsoft SQL restore options: <ul style="list-style-type: none"> ■ Select Leave the database ready to use; additional transaction logs or differential backups cannot be restored. ■ Select Overwrite the existing database. See "Restoring data by setting job properties" on page 605.

Table I-6 To recover SharePoint 2007 data after a disaster (*continued*)

Step	Action	Notes
Step 5	Restore SharePoint farm resources.	<p>Select the backup sets for all SharePoint resources in the farm except for the following:</p> <ul style="list-style-type: none"> ■ Configuration V3-DB database resource ■ Global Settings resource ■ Single Sign-on database resource (if applicable) <p>Configure the following Microsoft SharePoint restore options:</p> <ul style="list-style-type: none"> ■ Select Bring restored databases online. ■ Select Reconnect previous database links. ■ Select Replace existing database. <p>See “Restoring data by setting job properties” on page 605.</p> <p>Note: Some of the remaining restore jobs may fail because communication with the SharePoint server has not been fully established yet. This behavior is expected. Proceed with the recovery process until all steps are complete.</p>

Table I-6 To recover SharePoint 2007 data after a disaster (*continued*)

Step	Action	Notes
Step 6	Restore the SharePoint Configuration V3-DB resource.	<p>Select the backup sets for the SharePoint Configuration V3-DB resource.</p> <p>Configure the following Microsoft SharePoint restore options:</p> <ul style="list-style-type: none"> ■ Select Bring restored databases online. ■ Select Reconnect previous database links. ■ Select Replace existing database. <p>See “Restoring data by setting job properties” on page 605.</p>
Step 7	Restart the SharePoint server.	<p>After the restore job is complete, restart the SharePoint server. Then proceed to the next step.</p>
Step 8	Restore remaining SharePoint resources.	<p>Select the backup sets for the SharePoint Global Settings and Single Sign-on database resources, if necessary.</p> <p>See “Restoring data by setting job properties” on page 605.</p>
Step 9	Back up the SharePoint server.	<p>When the disaster recovery is complete, Symantec recommends that you perform a backup job as soon as possible.</p> <p>See “Creating a backup job by setting job properties” on page 332.</p>

About using the SharePoint Agent with SharePoint Portal Server 2003 and Windows SharePoint Services 2.0

You can use the SharePoint Agent to back up and restore SharePoint Portal Server 2003 farm components, which include the following:

- Configuration database
- Portal sites and their associated databases, which include the following:
 - Content database
 - User Profile database
 - Services database
 - Index databases
 - Team databases
- Windows SharePoint Services sites and their associated databases
- Single Sign-on database
- Document Library Store (Web Storage System-based)
- Document Libraries (Web Storage System-based)
Individual documents and their versions can be backed up from and restored to Web Storage System-based document libraries, or redirected to file paths.
- Document Libraries / Picture Libraries (Microsoft SQL Server based)
Individual documents and their versions can be restored from full database backups.
- Sites and sub-sites
Individual objects and their versions can be restored from full database backups.
- Lists and list items
Individual objects can be restored from full database backups.
See the Microsoft SharePoint documentation for more information about lists and list items.

In addition, you can back up and restore Windows SharePoint Services components, which include the following:

- Configuration database
- Team sites and their associated Content database

- Document Libraries / Picture Libraries (Microsoft SQL Server based)
Individual documents and their versions can be restored from full database backups.

About selecting SharePoint Server 2003 resources for backup

Backup Exec provides a hierarchical tree view of SharePoint resources in the following locations in the Selection tree:

- In the selection tree, a node titled **Microsoft SharePoint Server Farms** displays a logical view of the topology of each SharePoint server farm on your network. Backup Exec automatically discovers SharePoint farms when you browse to a SharePoint front-end Web server and adds the farms to this node. Additionally, you can add farms manually using the **Add Server Farm** menu option.
- A node titled **Microsoft SharePoint Resources** displays for any server that has locally installed SharePoint resources. For single-server SharePoint deployments, all of the SharePoint resources are listed and are selectable for backup. For SharePoint server farm deployments, this node lists only the SharePoint resources that reside locally on this server and that can be selected for backup from this node. On front-end Web servers, this node lists the entire farm topology, but only the resources that reside locally can be selected for backup.

Backing up resources from SharePoint 2003

Each portal site has a minimum of three databases: Content databases, Services databases, and User Profile databases. Symantec recommends that you back up these databases together to preserve the topology.

To back up SharePoint resources

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 Select the SharePoint resources you want to back up.
See [“About selecting SharePoint Server 2003 resources for backup”](#) on page 1236.
- 4 On the **Properties** pane, under **Settings**, select **Microsoft SharePoint**.
- 5 Select the appropriate options.
See [“Selections options for backup jobs”](#) on page 336.
- 6 Start the backup job or select other backup options from the **Properties** pane.

About selecting SharePoint 2003 resources for restore

Backup Exec provides the following hierarchical tree views of SharePoint resources in the Selections tree when using the **Resource View**:

Table I-7 Hierarchical tree view

Node name	Node description
Server Farm node	Represents a logical view of the topology for the SharePoint resources that have been backed up from the farm. The name that displays for this node matches the name you defined for the server farm in Backup Selections under the Microsoft SharePoint Server Farms node. If you expand the nodes for each SharePoint component that appears in this view, the backup sets for that component are displayed and can be selected for restore.
Individual Server nodes	Displays the SharePoint components that resided locally on the server when they were backed up. If you expand the nodes for each SharePoint component that appears in this view, the backup sets for that component are displayed and can be selected for restore. Each server from which SharePoint components were backed up contains a Microsoft SharePoint Resources node.

You can restore the following resources:

- Portal sites and their associated databases: each portal site has a minimum of three databases: Content databases, Services databases, and User Profile databases. Symantec recommends that you restore these databases together to preserve the topology.
- Windows SharePoint Services sites and their associated databases
- Document library stores (Web Storage System-based)
- Individual documents that are contained in Document or Picture libraries (Web Storage System-based or Microsoft SQL Server-based)
- Sites and sub-sites
Individual objects and their versions can be restored from full database backups.
- Lists and list items
Individual objects can be restored from full database backups.
See the Microsoft SharePoint documentation for more information about lists and list items.
- Configuration database: the Configuration database contains all of the configuration information for the entire SharePoint Server farm. Use caution

when restoring this database because any changes made to the farm topology after the backup from which you are restoring will be lost. For more information, refer to the Microsoft SharePoint Portal Server 2003 documentation. You can only restore the Configuration database back to the original location.

- Single Sign-on databases: you can only restore Single Sign-on databases back to the original location.

Restoring SharePoint 2003 resources

You can restore SharePoint 2003 resources.

See [“About selecting SharePoint 2003 resources for restore”](#) on page 1237.

To restore SharePoint resources

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 Select the full and differential backup sets that correspond to the SharePoint data you want to restore.

If you restore the SharePoint resources for a portal site in one job, the Index database is restored last. If you are restoring in separate jobs, you must restore the Index database last.

- 5 On the **Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 6 Select the appropriate options.
See [“Microsoft SharePoint restore options”](#) on page 1247.
- 7 Set additional restore options on the **Properties** pane or start the restore job.

Restoring individual SharePoint 2003 objects (Microsoft SQL Server-based) from full database backups

You can restore individual documents, images, sites, sub-sites, lists, and list items from full SharePoint database backup jobs if you selected the following option during the backup job:

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from the database backup (available for full backups only)

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

The option to enable the restore of individual objects is not available for differential backup jobs.

To restore individual objects from full database backups

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Navigate to the backup set for the appropriate content database, and then select the objects that you want to restore.

See “[Selections options for restore jobs](#)” on page 609.

- 4 On the **Properties** pane, under **Settings**, click **Microsoft SharePoint**.
- 5 Do one of the following:

If versioning is enabled on the device to which you are restoring individual documents

Select one of the following options:

- **Add as a new version**
Backup Exec restores the existing item as a new version, making it the most recent version of the existing item.
- **Skip if the item exists**
Backup Exec does not restore the item if an identical item exists in the restore destination. Backup Exec notes that the file was skipped in the job log.
- **Restore over existing items**
Backup Exec replaces the existing item with the restored item.

If versioning is not enabled on the device to which you are restoring individual documents

Select one of the following options:

- **Skip if the item exists**
Backup Exec does not restore the item if an identical item exists in the restore destination. Backup Exec notes that the file was skipped in the job log.
- **Restore over existing items**
Backup Exec replaces the existing item with the restored item.

- 6 Select the **Restore only the most recent version of an item** check box if you only want to restore the most recent version of each object you restore.
- 7 Check the **Include security information** option to restore any applicable security information with the item.

You can restore different levels of security based on the SharePoint item you restore:

- Sites - User and SharePoint Group information and security ACL are restored for top-level sites
 - Sub-sites - Security ACL is restored
 - Lists - Security ACL and other security-related information is restored
- 8 If you are restoring from tape, do the following steps in the order listed:
- On the **Properties** pane, under **Settings**, select **Advanced**.
 - Specify the path for a temporary staging location in the option titled **Path on an NTFS volume that is local to the media server for temporary storage of restore data**.
The path must reside on the Backup Exec media server. Symantec recommends that you avoid using system volumes for temporary staging locations.
- 9 Set additional restore options on the **Properties** pane, or start the restore job.

Restoring SharePoint 2003 document libraries (Web storage system-based)

Individual SharePoint documents are always restored to SharePoint document libraries as checked out to the credentials specified by the logon account used for the restore. The documents must be checked in or published by that user before others can use them.

If you try to restore over a document that is published or checked in, the restore will fail. If you try to restore over a document that is checked out, the restore will fail if the document is checked out to a user that differs from the logon account credentials used for the restore.

To restore SharePoint 2003 document libraries (Web storage system-based)

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Select the SharePoint Document Library data you want to restore.
See [“Selections options for restore jobs”](#) on page 609.
- 4 Set additional restore options on the **Properties** pane or start the restore job.

Restoring previous versions of SharePoint 2003 documents from document library (Web storage system-based) backups

The SHADOW folder, at the root of the Document Library, contains previous versions of the documents that existed in the Document Library at the time of backup. If you select the SHADOW folder to include in a Document Library backup, you can have access to the previous versions of the documents. However, you cannot restore the previous versions directly back into the Document Library. You must restore them to an alternate location and then manually copy them into the Document Library.

To restore previous versions of SharePoint 2003 documents from document library (Web storage system-based) backups

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Navigate to the SharePoint Document Library that contains the documents you want to restore.
- 4 Under the backup set, expand the SHADOW folder and then select the documents you want to restore.
- 5 Redirect the restore job for individual documents to a file path.

See [“Redirecting the restore of individual SharePoint 2003 objects to a file path”](#) on page 1243.

Redirecting a restore job for SharePoint 2003

You can redirect a restore job to an existing site on a Web server in a farm.

To bring the databases online after you complete the redirected restore job, verify that the **Bring restored databases online** and **Reconnect previous database links** options are selected in the Microsoft SharePoint settings. When you restore portal sites or Windows SharePoint Services sites, these options also re-establish the link between the restored databases and their corresponding sites.

Note: If you restore full or differential backup sets in separate restore jobs, clear these options for all jobs except the last job. You should select these options for the last restore job in the sequence. You may be prompted to insert any media that you have already used.

Follow these steps to redirect a restore job for SharePoint 2003 data.

To redirect a restore job for SharePoint 2003

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Select the SharePoint resources you want to restore.
See "[Selections options for restore jobs](#)" on page 609.
You can restore the Configuration database and Single Sign-on database back to the original location only.
- 4 On the **Properties** pane, under **Destination**, select **Microsoft SharePoint Redirection**.
- 5 Check **Redirect Microsoft SharePoint sets**.
- 6 Click **SharePoint 2003 portal sites or SharePoint 2007 web applications**.
- 7 In the **URL or web application name** field, type the URL of the site to which you want to restore the data.
For example, `http://portalsite1` or `https://portalsite1`. To restore to a SharePoint 2003 site, the site must already exist.
- 8 In the **Front-end web server name** field, type the name of the Web server on which the site resides.
You must create the destination SharePoint Portal Server 2003 portal site or Windows SharePoint Services site on the specified Web server with the same database structure as the source site before you run the restore job.
- 9 Do one of the following:
 - Use the default logon account as indicated.
 - Click **Change** to select a different logon account.
- 10 Set additional restore options on the **Properties** pane, or start the restore job.

Redirecting the restore of SharePoint 2003 document library (Web storage system-based) data to another document library

Before redirecting the restore of SharePoint 2003 document library data, the SharePoint Portal Server software must be installed on the target server. If any of the folders in the original document library do not exist in the destination document library, they will be created during the restore.

Caution: When you restore SharePoint Portal document library data, any documents that exist in the target location and that have the same name as the documents being restored may be overwritten, depending on the Backup Exec overwrite properties for the restore job.

To redirect the restore of SharePoint 2003 document library data to another document library

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Navigate to and select the SharePoint Portal document library data you want to restore.
See "[Selections options for restore jobs](#)" on page 609.
- 4 On the **Properties** pane, under **Destination**, click **Microsoft SharePoint Redirection**.
- 5 Check **Redirect Microsoft SharePoint sets**.
- 6 Click **Individual SharePoint sites, documents, lists, or items**.
- 7 Select **Restore to workspace or document library (Web Storage System-based only)**.
- 8 In the **Restore to server** field, type the name of the SharePoint server to which you want to restore.
Use the following format: \\servername.
- 9 In the **Restore to workspace or document library** field, type the name of the document library to which you want to restore.
If you have not yet created the document library, you must do so before starting the restore operation.
- 10 Use the default logon account as indicated, or click **Change** to select a different one.
- 11 On the **Properties** pane, select other job properties that might be appropriate for your environment.
- 12 Start the restore job.

Redirecting the restore of individual SharePoint 2003 objects to a file path

You can redirect the restore of SharePoint file-based data such as documents and images that have been uploaded to a document library or are attached to list items.

Individual items and their versions can only be redirected to a file location, not another content database. Use the following steps to restore individual objects to a file path.

Note: You cannot restore SharePoint security information when you restore an item to a file path.

To redirect individual SharePoint 2003 objects to a file path

- 1 Place the media containing the data you want to restore in the storage device.
- 2 On the navigation bar, click the arrow next to **Restore**.
- 3 Click **New Restore Job**.
- 4 Navigate to and select the SharePoint objects that you want to restore.
See "[Selections options for restore jobs](#)" on page 609.
- 5 On the **Properties** pane, under **Destination**, click **Microsoft SharePoint Redirection**.
- 6 Check **Redirect Microsoft SharePoint sets**.
- 7 Click **Individual SharePoint sites, documents, lists, or items**.
- 8 Select **Redirect to path**.
- 9 In the **Restore to drive or UNC path** and **Restore to path** fields, enter the drive letter and path to which you want to direct the restore, or click the ellipsis (...) button to browse to the location.
Use the following format for a UNC path: \\servername\share.
- 10 Use the default logon account as indicated, or click **Change** to select a different one.
- 11 On the **Properties** pane, select other job properties that might be appropriate for your environment.
- 12 Start the restore job.

Redirecting the restore of a Microsoft Office SharePoint Server 2003 database to an alternate SQL instance

You can restore SharePoint databases to an alternate SQL instance.

To redirect the restore of a Microsoft Office SharePoint Server 2003 database to an alternate SQL instance

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Navigate to and select the SharePoint database that you want to restore.
- 4 On the **Properties** pane, under **Destination**, click **Microsoft SharePoint Redirection**.
- 5 Check **Redirect Microsoft SharePoint sets**.
- 6 Click **Restore Sharepoint databases to an alternate SQL instance**.
- 7 In the **Restore to SQL server\instance** field, enter the SQL server and instance to which you want to redirect the SharePoint database.
- 8 In the **Restore to database** field, enter the name of the database to which you want to redirect the SharePoint database.
- 9 Select **Replace existing database** to overwrite an existing database with the same name.

If you do not select this field and a database with the same name exists on the SQL instance, you receive an error.
- 10 Use the default logon account as indicated, or click **Change** to select a different one.
- 11 On the **Properties** pane, select other job properties that might be appropriate for your environment.
- 12 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later

Do the following in the order listed:

- In the **Properties** pane, under **Frequency**, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See [“Scheduling jobs”](#) on page 356.

See [“About using the SharePoint Agent with SharePoint Portal Server 2003 and Windows SharePoint Services 2.0”](#) on page 1235.

Microsoft SharePoint backup options

You can set specific options for SharePoint when you run a backup job.

Table I-8 Microsoft SharePoint backup options

Item	Description
<p>Backup method</p>	<p>Allows you to select from the following backup methods:</p> <ul style="list-style-type: none"> ■ Full - Back up entire database Backs up the entire database. ■ Differential - Back up database changes only Backs up only the changes made to the database since the last full backup. The differential backup method cannot be used to back up Index databases or Document Libraries. The Full backup method must be used to back up these resources. ■ Log - Back up and truncate transaction log Backs up the data that is contained in the transaction log. This method does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated). If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery model before a log backup is run. Alternatively, you can run full backups only, or run full and differential backups of the SQL databases that are associated with SharePoint.
<p>Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual documents from the database backup (available for full backups only)</p>	<p>Enables the restore of individual documents, images, sites, sub-sites, lists, and list items from database backups. This option is only available when you perform full backups. The option is not available if Differential - Back up database changes only has been selected as the backup method. You must have a current version of the Remote Agent for Windows Systems on the SharePoint server when you run the GRT-enabled backup job.</p> <p>See “How to restore individual items by using Granular Recovery Technology” on page 321.</p>
<p>Release the lock on the SharePoint farm topology if it is set (SharePoint 2003 only)</p>	<p>Releases the lock on the SharePoint farm topology before you run the backup or restore operation. Because another application may have locked the topology, you should check with your SharePoint administrator before you select this option.</p>

Table I-8 Microsoft SharePoint backup options (*continued*)

Item	Description
Perform consistency check before backup of Microsoft SQL databases used by Microsoft SharePoint	Runs a full consistency check (including indexes) of the Microsoft SQL databases used by Microsoft SharePoint before you back up the databases.
Continue with backup if consistency check fails	Continues with the backup operation even if the consistency check fails.

See [“Backing up resources from SharePoint 2003”](#) on page 1236.

See [“Backing up a farm for Microsoft Office SharePoint Server 2007 or a Windows SharePoint Services 3.0”](#) on page 1219.

See [“Backing up individual SharePoint 2007 Web applications in a Microsoft SharePoint server farm”](#) on page 1220.

See [“Backing up a farm for Microsoft Office SharePoint Server 2010 or Windows SharePoint Foundation 2010”](#) on page 1202.

See [“Backing up individual SharePoint 2010 Web applications in a Microsoft SharePoint server farm”](#) on page 1203.

Microsoft SharePoint restore options

You can set specific options for SharePoint when you run a restore job.

Table I-9 Microsoft SharePoint restore options

Item	Description
Bring restored databases online	Brings the databases online after a restore job.
Reconnect previous database links	Re-establishes the link between the restored databases and their corresponding sites when you restore Web applications or Windows SharePoint Services sites.

Table I-9 Microsoft SharePoint restore options (*continued*)

Item	Description
Release the lock on the SharePoint farm topology if it is set (SharePoint 2003 only)	Releases the lock on the SharePoint farm topology before you run the backup or restore operation. Because another application may have locked the topology, you should check with your SharePoint administrator before you select this option. This option is for SharePoint 2003 only.
Preserve existing Internet Services (IIS) Web site and application pool (SharePoint 2007 and 2010 only)	Preserves the web site and application pool for the SharePoint Web application that you restore if it already exists in IIS. If you do not check this option, the Web site and application pool are deleted from IIS during the restore. After they are deleted, they are recreated in the default location that SharePoint specifies. This option is for SharePoint 2007 and 2010 only.
Replace existing database	Specifies that you want to replace existing databases on the specified SQL server instance with the databases being restored. Otherwise, an error occurs if databases with the same names as those being restored exist on the specified SQL server instance.
If versioning is enabled on the restore destination	<p>Allows you to choose from the following options if versioning is enabled on the destination to which you want to restore an individual item:</p> <ul style="list-style-type: none"> ■ Add as a new version Backup Exec restores the existing item as a new version, which makes it the most recent version of the existing item. ■ Skip if the item exists Backup Exec does not restore the item if an identical item exists in the restore destination. Backup Exec notes that the file was skipped in the job log. ■ Restore over existing items Backup Exec replaces the existing item with the restored item.
If versioning is not enabled on the restore destination	<p>Allows you to choose from the following options if versioning is not enabled on the destination to which you want to restore an individual item:</p> <ul style="list-style-type: none"> ■ Skip if the item exists Backup Exec does not restore the item if an identical item exists in the restore destination. Backup Exec notes that the file was skipped in the job log. ■ Restore over existing items Backup Exec replaces the existing item with the restored item.

Table I-9 Microsoft SharePoint restore options (*continued*)

Item	Description
Restore only the most recent version of an item	Restores only the most recent version of an item.
Include security information	Restores any applicable security information with the item. You can restore different levels of security based on the SharePoint item you restore: <ul style="list-style-type: none"> ■ Sites - User and SharePoint Group information and security ACL are restored for top-level sites ■ Sub-sites - Security ACL is restored ■ Lists - Security ACL and other security-related information is restored

The procedures for restoring SharePoint data vary depending on the type of data you want to restore.

- See [“Restoring SharePoint 2003 resources”](#) on page 1238.
- See [“Restoring SharePoint 2003 document libraries \(Web storage system-based\)”](#) on page 1240.
- See [“Restoring previous versions of SharePoint 2003 documents from document library \(Web storage system-based\) backups”](#) on page 1241.
- See [“Restoring individual SharePoint 2003 objects \(Microsoft SQL Server-based\) from full database backups”](#) on page 1238.
- See [“Restoring resources for SharePoint Server 2007 and SharePoint Services 3.0”](#) on page 1222.
- See [“Restoring a Microsoft Office SharePoint Server 2007 Shared Services Provider”](#) on page 1224.
- See [“Restoring a Microsoft Office SharePoint Server 2007 Web application to its original location”](#) on page 1225.
- See [“Restoring individual SharePoint 2007 objects from full database backups to their original locations”](#) on page 1223.
- See [“Restoring resources for SharePoint Server 2010 and SharePoint Foundation 2010”](#) on page 1204.
- See [“Restoring individual SharePoint 2010 objects from full database backups to their original locations”](#) on page 1205.

- See [“Restoring a Microsoft Office SharePoint Server 2010 Web application to its original location”](#) on page 1207.

Microsoft SharePoint redirection options

You can set specific options for SharePoint when you run a redirected restore job.

Table I-10 Microsoft SharePoint redirection options

Item	Description
Redirect Microsoft SharePoint sets	Redirects SharePoint restore jobs to a new location.
SharePoint 2003 portal sites or SharePoint 2007/2010 web applications	Enables redirection for SharePoint 2003 portal sites or SharePoint 2007/2010 web applications.
URL or web application name	Specifies the URL of the site or web application to which you want to restore data.
Front-end web server name	Specifies the name of the Web server on which the site you want to restore resides.
Individual SharePoint sites, documents, lists or items	Enables redirection for SharePoint sites, documents, lists, or items.
Redirect to path	Redirects individual SharePoint sites, documents, lists, or items to a file path.
Restore to drive or UNC path	Specifies the drive or UNC path to which you want to direct the restore. Use the following format for a UNC path: <code>\\servername\share</code>
Restore to path	Specifies the path to which you want to direct the restore job.
Redirect to workspace or document library (Web Storage System-based only)	Redirects individual SharePoint sites, documents, lists, or items to a workspace or a document library.
Restore to server	Specifies the SharePoint server to which you want to direct the restore job. Use the following format: <code>\\servername</code> .

Table I-10 Microsoft SharePoint redirection options (*continued*)

Item	Description
Restore to workspace or document library	Specifies the name of the document library to which you want to direct the restore job. If you have not yet created the document library, you must do so before starting the restore operation.
Restore SharePoint databases to an alternate SQL instance	Lets you restore SharePoint databases to a different SQL instance.
Restore to SQL server\instance	Specifies the SQL server and instance to which you want to restore the SharePoint database.
Restore to database	Specifies the name of the alternate database to which you want to restore the SharePoint database.
Replace existing database	Specifies that you want to replace existing databases on the specified SQL server instance with the databases being restored. Otherwise, an error occurs if databases with the same names as those being restored exist on the specified SQL server instance.
SharePoint logon account	Specifies the logon account you use to access SharePoint data. Click Change to select a different account.

The procedures for redirecting SharePoint data vary depending on the type of data you select and the location to which you want to redirect it.

- See [“Redirecting a restore job for SharePoint 2003”](#) on page 1241.
- See [“Redirecting the restore of a Microsoft Office SharePoint Server 2003 database to an alternate SQL instance”](#) on page 1244.
- See [“Redirecting a restore job for SharePoint 2007”](#) on page 1227.
- See [“Redirecting the restore of a Microsoft Office SharePoint Server 2007 Web application”](#) on page 1229.
- See [“Redirecting the restore of individual SharePoint 2007 objects to a file path”](#) on page 1228.
- See [“Redirecting the restore of a Microsoft Office SharePoint Server 2007 database to an alternate SQL instance”](#) on page 1230.

- See [“Redirecting a restore job for SharePoint 2010”](#) on page 1208.
- See [“Redirecting the restore of individual SharePoint 2010 objects to a file path”](#) on page 1210.
- See [“Redirecting the restore of a Microsoft Office SharePoint Server 2010 Web application”](#) on page 1211.
- See [“Redirecting the restore of a Microsoft Office SharePoint Server 2010 database to an alternate SQL instance”](#) on page 1213.

Symantec Backup Exec Agent for Microsoft SQL Server

This appendix includes the following topics:

- [About the Agent for Microsoft SQL Server](#)
- [Requirements for using the SQL Agent](#)
- [About installing the SQL Agent](#)
- [How to use Backup Exec logon accounts for SQL resources](#)
- [About backup strategies for SQL](#)
- [About consistency checks for SQL](#)
- [How to use snapshot technology with the SQL Agent](#)
- [Setting default backup and restore options for SQL](#)
- [Setting backup options for SQL](#)
- [Setting restore options for SQL](#)
- [About restoring SQL databases and file groups](#)
- [About disaster recovery of a SQL Server](#)

About the Agent for Microsoft SQL Server

The Symantec Backup Exec Agent for Microsoft SQL Server (SQL Agent) enables network administrators to perform backup and restore operations on installations of SQL that are connected to a network. SQL database backups can be integrated with network backups without separate administration or dedicated hardware.

The SQL Agent provides support for the following:

- Database, transaction log, differential, and filegroup backups, as well as database recovery and replacement.
- An automated restore of the master database.
- The Intelligent Disaster Recovery option, which automates the disaster recovery process of SQL Servers.
- Restores of SQL databases to alternate locations.
- Automated restore selections and options checking, which tests the validity of your current SQL Server restore selections and job options before the restore job runs.
- Hot backup copies of SQL databases during backup operations. This feature enables you to direct a copy of the actual data streams being sent to media by a SQL database to a local directory for later use.
- Backups of multiple instances.
- Integration with the Symantec Backup Exec Advanced Disk-based Backup Option (ADBO) and the Advanced Open File Option (AOFO). ADBO and AOFO are separate, add-on components of Backup Exec. The use of ADBO and AOFO can reduce both restore time and backup impact on the server.
- Standby database. If the primary SQL server fails, or is shut down for maintenance, another database called a standby database can be brought online. The standby database contains a copy of the primary server's databases so that users can continue to access the database even though the primary server is unavailable. When the primary server is available again, the changes on the standby database must be restored back to the primary server or the changes will be lost. The databases on the primary server should then be backed up and restored on the standby database again.

Backup Exec provides a backup option that enables you to put the database in standby mode when the log file backup completes, and a recovery completion state of Leave the database in read-only mode to create and maintain a standby database.

- Database Consistency Checks (DBCC) for each backup and restore job, including a fast database consistency check of only the physical consistency of the database.
- Full, bulk-logged, and simple recovery models. With the simple recovery model, copies of the transactions are not stored in the log file, which prevents transaction log backups from being run. Therefore, you can recover the database to the point of the last backup, but you cannot restore the database to the point of failure or to a specific point in time.
- Restores of transaction logs to a specific point in time or to a named transaction when log marks are used.

In SQL 2005 or later installations, the SQL Agent provides support for the following:

- Database snapshots.
- New copy backup jobs, which enable you to copy a SQL 2005 or later database without having to run a full SQL database backup job.
- Maintaining replication settings during redirected restores.
- Verify only restore jobs, which enable you to determine both the validity of the SQL data on the media and the ability of the destination SQL database to accept this data before the database is deleted or overwritten during a restore job.
- Back up with checksum generation. Used as a redundancy check, this option works with the Verify Only Restore Job option.
- Continuation of restore jobs when errors are detected. This feature enables you to restore as much data as possible from a corrupt database backup.

In SQL Server 2008 Enterprise Edition installations, the SQL Agent provides support for the following:

- In SQL Server 2008 Enterprise Edition installations, you can use SQL software compression for backup jobs.

See [“About installing the SQL Agent”](#) on page 1256.

Requirements for using the SQL Agent

The following are required for the SQL Agent:

- Backup Exec must have access rights to read both of the following SQL registry keys:
 - HKEY_LOCAL_MACHINE\Software\Microsoft\Microsoft SQL Server

- HKEY_LOCAL_MACHINE\Software\Microsoft\MSSQLServer

If Backup Exec does not have access to these registry keys, a restore to the default directory may not work, and the Automate master database restore option on the Restore Job Properties for SQL dialog box will not work.

To ensure that Backup Exec has access rights, verify that the logon account used has administrator rights to the Windows server that the SQL instance is installed on.

- The media server must have access to the SQL installation.
- The credentials stored in the Backup Exec logon account used for backing up and restoring SQL must have been granted the System Administrator role on the SQL instance.
- The Backup Exec Remote Agent for Windows Systems (Remote Agent) must be installed on any remote SQL Server that you want to back up.

About installing the SQL Agent


The SQL Agent is installed locally as a separate, add-on component of Backup Exec and can protect local or remote SQL Server databases.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

How to use Backup Exec logon accounts for SQL resources

To back up SQL, use a Backup Exec logon account that stores the credentials of a Windows user account. The Windows user account must have been granted the System Administrator role on the SQL instance.

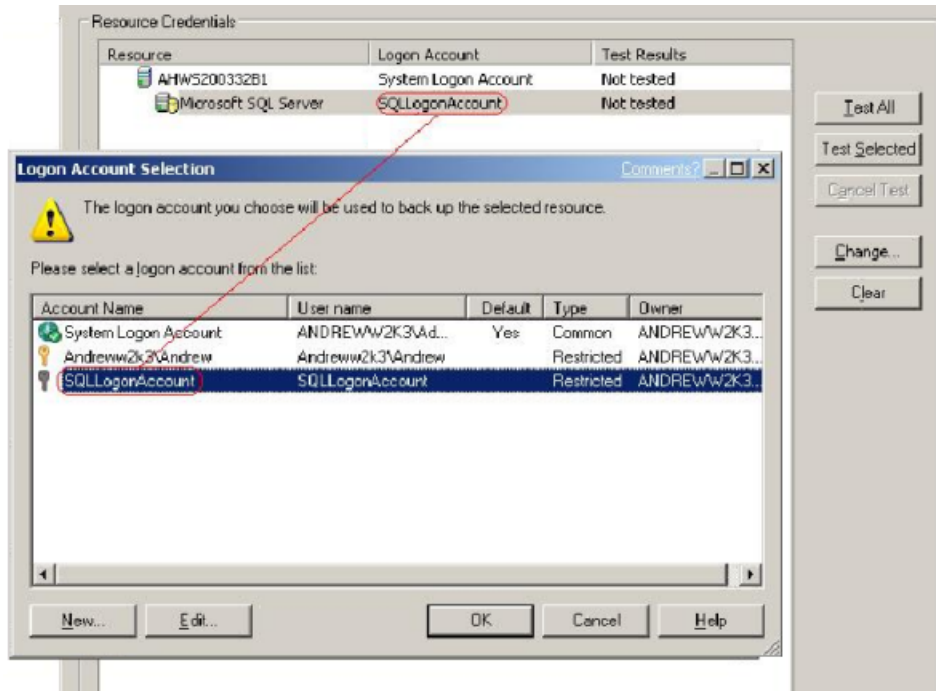
In the backup selections list or in the resource credentials list, apply that logon account to the Windows server that SQL is installed on, not to the actual SQL instance.

Figure J-1 Applying Windows user account credentials

If you are using SQL Server Authentication, then add a Backup Exec logon account that stores the credentials of the SQL user account.

In the backup selections list, apply the Backup Exec logon account for the Windows user account to the Windows server that SQL is installed on, and then apply the logon account for the SQL user account to the SQL instance.

Figure J-2 Applying SQL user account credentials



If you use a Backup Exec logon account that does not have the proper rights, you will receive an error message stating that the username and password are invalid.

See [“About selection lists”](#) on page 295.

See [“Creating a new Backup Exec System Logon Account”](#) on page 193.

About backup strategies for SQL

Backup Exec incorporates online, nondisruptive SQL database protection as part of everyday backup routines, which increases the chance of data recovery and minimizes data loss without inhibiting daily database activity. Using database, differential, and log backups provides a good balance between backup windows and minimizes the amount of time that will be spent recovering a database if the need arises.

To decide which backup methods to use for the best data protection, consider the following for typical environments:

- In small environments, consider running a daily full database backup every evening and daily transaction log backups.

- In mid-sized environments, consider running a weekly full database backup and daily transaction log backups along with daily differential backups except on the day when the full backup is run.
- In large environments, consider running daily differential database backups, weekly full database backups, and transaction log backups as necessary. Many shops run full backups on a weekly basis, preferring to run differential backups throughout the week to keep backup run time to a minimum. Extremely large environments may need to run filegroup backups in order to split the full backup over several days. Log backups are required to be able to recover a system from a filegroup backup.

The trade-off with running fewer full backups and running more differential backups occurs at recovery time when you must recover using the full database backup as well as the last differential database backup, and all log backups made after the last differential database backup.

What will work best for you will be based on the size of your environment, the number of transactions processed each day, and the expectations of your users when a recovery is required.

SQL backup strategy recommendations

When you develop a SQL backup strategy, consider the following:

Table J-1 Recommendations for backing up SQL

SQL Server backup strategies	Description
Protect the entire SQL Server	To make sure SQL is completely protected, back up the following on a regular basis: <ul style="list-style-type: none">■ The system drive that SQL is on.■ The Windows registry and System State.■ SQL databases or filegroups. You do not need to back up both.■ Transaction logs.
When you upgrade, run new full database backups.	If you upgrade SQL, run new full database backups. You may not be able to restore backups from one version or service pack level of SQL to other versions.

Table J-1 Recommendations for backing up SQL (*continued*)

SQL Server backup strategies	Description
<p>Run consistency checks after backups.</p>	<p>Symantec recommends that you run a consistency check after a backup. If a database, transaction log, or filegroup contains errors when it is backed up, the backup will still contain the errors when it is restored, if it is restorable at all.</p> <p>These consistency checks include the following:</p> <ul style="list-style-type: none"> ■ A full consistency check, including indexes. This check will have significant impacts on SQL performance; therefore, it should be performed in off-peak hours. ■ A full consistency check with no index check. While not as thorough as a full consistency check that includes indexes, this check is faster and can be done during peak hours with little impact on system performance. ■ A physical check only. Another low-overhead check, this method checks only the integrity of the physical structure of the page and record headers, and the consistency between the pages' object ID and index ID and the allocation structures.
<p>Back up the master database whenever data is changed in the master database.</p>	<p>Back up the master database whenever procedures are run that change information in the database, especially after the following:</p> <ul style="list-style-type: none"> ■ New databases are created. ■ Files are added to an existing database. ■ Usernames or passwords are added or changed. <p>If changes are not backed up before the master database must be restored, the changes are lost.</p>
<p>Run one backup at a time.</p>	<p>Do not schedule more than one backup to occur simultaneously against a database or its transaction log, or a filegroup.</p>

Table J-1 Recommendations for backing up SQL (*continued*)

SQL Server backup strategies	Description
Back up both system and user databases and transaction logs regularly.	<p>Copies of the master and model databases are automatically created by Backup Exec whenever you back up the master and model databases. If these databases become corrupted or are missing, and SQL cannot be started, you can replace them with the copies of the master and model databases, and then start SQL. After SQL is running again, you can restore the latest copy of the master database using Backup Exec's Automate master database restore option, and then restore any other databases, if needed.</p> <p>If you use the Intelligent Disaster Recovery (IDR) option, then during an IDR recovery of the C: drive, it will automatically replace the damaged databases with the copies of the master and model databases that you made.</p>
If you have filegroups, back them up instead of databases. Do not back up filegroups and databases.	When databases grow too large to be backed up all at once, filegroups can provide an alternative backup method. Different filegroups can be backed up at different times and frequencies. A combination of filegroup and log backups provides complete database protection.

About consistency checks for SQL

If you back up a database, transaction log, or filegroup that contains errors, these errors will still exist when the backup is restored. In some cases, this can prevent a successful restore. Backup Exec enables you to check the logical and physical consistency of the data before and after a backup. SQL reports any consistency check failures in the Backup Exec job log. Symantec strongly recommends that you always run a consistency check either before or after the backup.

Backup Exec's consistency check uses the following SQL consistency check utilities:

- CHECKDB
- CHECKCATALOG
- CHECKFILEGROUP
- PHYSICAL_ONLY

CHECKDB, CHECKCATALOG, and PHYSICAL_ONLY are performed for database-related operations.

CHECKFILEGROUP is performed for filegroup-related operations.

For more information concerning these utilities, see your MS SQL documentation.

See [“Setting backup options for SQL”](#) on page 1273.

How to use snapshot technology with the SQL Agent

The SQL Agent supports snapshot technology for SQL through the use of the Symantec Backup Exec - Advanced Open File Option (AOFO) and the Advanced Disk-based Backup Option (ADBO). ADBO can only be installed on Windows Server 2003/2008. The use of ADBO and AOFO can reduce both restore time and backup impact on the server.

Note: The SQL Agent also supports SQL 2005 or later database snapshot technology.

See [“About SQL 2005 or later database snapshots”](#) on page 1286.

Before you use snapshot technology with the SQL Agent, review the following information:

- With snapshot technology, a point in time view of the SQL database is "snapped" and then backed up, leaving the actual SQL database open and available for users.
- Symantec recommends that SQL backup jobs be run separately from AOFO or ADBO backup jobs because SQL backups made with snapshot technology are considerably bigger than regular SQL backups.
- Performing consistency checks before backup is highly recommended. See [“About consistency checks for SQL”](#) on page 1261.
- The SQL Agent only supports full snapshot backups; filegroup snapshots, log snapshots, and differential snapshots are not supported.
- If a filegroup, differential, or transaction log backup method is selected, AOFO or ADBO Backup Job Properties are ignored and a traditional differential or transaction log backup will be performed.
- With the SQL Agent, snapshot and traditional backups are interoperable when restoring SQL data.
- For the Intelligent Disaster Recovery Option to work with SQL backups, copies are made of the master and model databases. Copies are only made when non-snapshot backups of master and model are run. If you are using AOFO or ADBO for SQL backups, make at least one backup of the master and model databases without using AOFO or ADBO.

- If SQL is upgraded, refresh the copies with another non-snapshot backup.
- SQL backups made using AOFO or ADBO will fail if multiple databases are selected for backup and SQL Service Pack 2 is not installed. If SQL 2000 Service Pack 2 or later is installed, you can select multiple databases at the same time for backup.
- Snapshot backups of the master database cannot be redirected.
- Performing database consistency checks both before and after backups impacts the time required for the backup jobs.

See [“Setting backup options for SQL”](#) on page 1273.

How to use AOFO with the SQL Agent

When using the SQL Agent together with AOFO, depending on which operating system you are running, you can select to use the VERITAS Storage Foundation for Windows FlashSnap Option or the Microsoft Volume Shadow Copy Service. What happens when the Automatically select open file technology option is selected, also depends on the operating system being used.

Note: When used with the SQL Agent, AOFO snapshot backups are limited to full backups of the Microsoft SQL Server databases.

See [“How to use AOFO and protect SQL on Windows 2003”](#) on page 1263.

See [“How to use AOFO and protect SQL on Windows 2000 ”](#) on page 1264.

How to use AOFO and protect SQL on Windows 2003

The SQL Agent supports Microsoft’s Volume Shadow Copy Service (VSS), a snapshot provider service only available on Windows 2003 or later, and the VERITAS Storage Foundation for Windows FlashSnap Option.

In order for the SQL Agent to use VSS, the SQL Agent must be installed and running on the SQL server. VSS can use different providers, including Default, System, Hardware, and Software.

In order for the SQL Agent to use the FlashSnap Option, the SQL Agent, VERITAS Storage Foundation for Windows, and the Advanced Open File Option (AOFO) must be installed and running on the SQL server.

See [“About the Advanced Open File Option”](#) on page 937.

How to use AOFO and protect SQL on Windows 2000

When you protect Windows 2000, the SQL Agent only supports the VERITAS Storage Foundation for Window FlashSnap Option.

In order to protect the SQL server using the FlashSnap Option, the SQL Agent, VERITAS Storage Foundation for Windows, and the Advanced Open File Option (AOFO) must be installed and running on the SQL server.

You can select the FlashSnap option through the Advanced Open File options dialog box. If Automatically select open file technology is selected on the Advanced Open File options dialog box when running Windows 2000, Backup Exec attempts to use the FlashSnap Option to perform the backup. If the FlashSnap option is not available, the job will fail. If you select the Symantec Volume Snapshot Provider or Microsoft Volume Shadow Copy Service options, a non-snapshot backup is performed.

Also, when using AOFO with FlashSnap note the following:

- The SQL user data and log files must exist on the mirrored volume (plex) in order for the snapshot to occur.
- To protect system databases, the system database files must reside on the mirrored volume (plex).

See [“About the Advanced Open File Option”](#) on page 937.

See [“Setting default options for the Advanced Open File Option”](#) on page 943.

See [“About SQL 2005 or later database snapshots”](#) on page 1286.

How to use ADBO with the SQL Agent

In order to protect the SQL server using ADBO, both the SQL Agent and ADBO must be installed. ADBO, which can only be installed on Windows 2003 or later, can use different snapshot providers, including hardware and software. You can select the provider to be used through the Advanced Disk-based Backup options when creating the backup job.

In order to use the Software - Use VERITAS Storage Foundation for Windows option, VERITAS Storage Foundation for Windows (VSWF) must be installed.

If Automatic - Use hardware if available; otherwise use software is selected, the first available hardware provider is used. If a hardware provider is not available, then the first software provider is used. If neither a hardware nor a software provider are available, the job status will depend the job disposition option set on the Advanced Disk-based Backup dialog box.

Also, when using ADBO note the following:

- The SQL user data and log files must exist on the mirrored volume (plex) in order for the snapshot to occur.
- To protect system databases, the system database files must reside on the mirrored volume (plex).

See [“About selecting data to back up”](#) on page 280.

See [“About selection lists”](#) on page 295.

See [“About consistency checks for SQL”](#) on page 1261.

Setting default backup and restore options for SQL

You can use the defaults set by Backup Exec during installation for all SQL backup and restore jobs, or you can choose your own defaults. You can also change the defaults for any specific backup or restore job.

See [“Setting backup options for SQL”](#) on page 1273.

See [“About restoring SQL databases and file groups”](#) on page 1293.

To set default backup and restore options for SQL

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Microsoft SQL**.
- 3 Select the appropriate options.

See [“Microsoft SQL default options”](#) on page 1265.

Microsoft SQL default options

You can set the following default options for all backup and restore jobs for Microsoft SQL.

See [“Setting default backup and restore options for SQL”](#) on page 1265.

Table J-2 Microsoft SQL default options

Item	Description
Backup method	

Table J-2 Microsoft SQL default options (*continued*)

Item	Description
	<p>Specifies one of the following backup methods:</p> <ul style="list-style-type: none"> <li data-bbox="677 369 1225 487"> <p>■ Full - Back up entire database or filegroup. This option backs up the entire database or filegroup. This option is selected by default. See “About backing up SQL databases” on page 1281.</p> <li data-bbox="677 496 1241 913"> <p>■ Log - Back up transaction log. This option backs up only the data contained in the transaction log; it does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated). If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery mode before a log backup is run. Alternatively, you can run full backups only, or run full and differential backups of the SQL databases. See “Backing up SQL transaction logs” on page 1285.</p> <li data-bbox="677 921 1241 1303"> <p>■ Log No Truncate - Back up transaction log - no truncate. This option backs up the database when it is corrupt or database files are missing. Since the Log No Truncate method does not access the database, you can still back up transactions that you may not be able to access when the database is in this state. You can then use this transaction log backup along with the database backup and any previous transaction log backups to restore the database to the point at which it failed; however, any uncommitted transactions are rolled back. The Log No Truncate method does not remove committed transactions after the log is backed up. See “Backing up SQL transaction logs” on page 1285.</p> <li data-bbox="677 1312 1241 1555"> <p>■ Differential - Back up database or filegroup changes only. This option backs up only the changes made to the database or filegroup since the last full backup. Because differential backups allow the restore of a system only to the point that the differential backup was created, you should also create multiple log backups between the differential backups. See “About backing up SQL databases” on page 1281.</p> <li data-bbox="677 1564 1241 1581"> <p>■ Database Snapshot (SQL 2005 Enterprise Edition or later)</p>

Table J-2 Microsoft SQL default options (*continued*)

Item	Description
	<p>- Read-only, point-in-time copy of another database This option creates a read only, point-in-time copy of another database. See “About SQL 2005 or later database snapshots” on page 1286.</p> <ul style="list-style-type: none"> ■ Full Copy-only (SQL 2005 or later) - Back up entire database or filegroup. Unlike the Full backup method, the Full Copy-only backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup. After making a full backup, you can use the Full Copy-only backup method to make a copy of a SQL database without affecting the baseline backup set required to run future differential backups.
<p>Consistency check before backup</p>	<p>Specifies one of the following consistency checks to run before a backup.</p> <ul style="list-style-type: none"> ■ None. This option does not run a consistency check before a backup. Symantec strongly recommends that you always run a consistency check either before or after the backup. This option is selected by default. ■ Full check, excluding indexes. This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the nonclustered index pages is not checked. ■ Full check, including indexes. This option includes indexes in the consistency check. Any errors are logged. ■ Physical check only. This option performs a low overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page and record headers, and the consistency between the pages' object ID and index ID and the allocation structures.

Table J-2 Microsoft SQL default options (*continued*)

Item	Description
Continue with backup if consistency check fails	Continues with the backup operation even if the consistency check fails. You may want to continue with the backup when the consistency check fails if you think that a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database with only a small problem in a table.
Consistency check after backup	<p>Specifies the consistency check to run after a backup. Because database transactions can occur during or after the consistency check, but before the backup runs, consider running a consistency check after the backup to ensure the data was consistent at the time of the backup.</p> <p>The following checks are available:</p> <ul style="list-style-type: none">■ None. This option does not run a consistency check after a backup. Symantec strongly recommends that you always run a consistency check either before or after the backup.■ Full check, excluding indexes. This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the nonclustered index pages is not checked.■ Full check, including indexes. This option includes indexes in the consistency check. Any errors are logged.■ Physical check only. This option performs a low overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page and record headers, and the consistency between the pages' object ID and index ID and the allocation structures. This option is selected by default.
Display filegroups when creating new backup jobs	Displays filegroups that you want to select for backup. If this checkbox is not selected, filegroups are not displayed as backup selections.

Table J-2 Microsoft SQL default options (*continued*)

Item	Description
Use checksums on backup (SQL 2005 or later)	Adds checksums to the SQL database data being backed up by Backup Exec. Adding checksums to the data being backed up is required if you want to use the option Run verify only; do not restore data. Using this option, along with Run verify only; do not restore data, ensures that during a restore of the SQL database, you are restoring from a verified SQL backup.
Database snapshots to keep	(SQL Server 2005 or later) Displays the number of database snapshots to keep on disk. As the threshold is met, older database snapshots are deleted, which are then replaced with new snapshots. Because database snapshots continue to grow as the SQL Server database is updated, limiting the number of snapshots enables you to minimize both the disk space and SQL Server processing time that is required when the snapshots are updated. See “About SQL 2005 or later database snapshots” on page 1286.
Create on-disk copies of SQL backups to be placed on the SQL server where the database is located	Creates an on-disk copy of the SQL database being backed up. This option lets you simultaneously back up a SQL database to storage media while also writing a copy of the database to a disk path you specify in the Save to path box. This option gives IT administrators the ability to back up SQL databases while also providing database administrators with copies of the database on disk, which can be used for such things as tests and restores. This option is not compatible with Advanced Open File Option backups or with database snapshot backups.
Save to path	Displays a path in which to save on-disk copies of SQL backups.

Table J-2 Microsoft SQL default options (*continued*)

Item	Description
SQL Server (2008 or later) software compression	<p>Specifies the following compression setting you want to use for this backup job:</p> <ul style="list-style-type: none">■ None. Do not use compression.■ Compress. Use SQL Server 2008 or later compression. <p>SQL compresses the data on the computer on which SQL Server 2008 Enterprise Edition or later is installed. Therefore, faster SQL 2008 or later backups should occur if you use SQL compression.</p> <p>If you back up remote SQL 2008 or later computers and you use SQL 2008 or later software compression, you must use the latest version of the Remote Agent.</p> <p>You can find a list of compatible operating systems, platforms, and applications at the following URL: http://entsupport.symantec.com/umi/V-269-1</p> <p>Symantec recommends that you do not use SQL 2008 or later software compression in a backup job that uses Backup Exec-initiated software compression. Minimal additional SQL 2008 or later compression benefits are gained when you enable Backup Exec compression. In fact, in jobs where both compression schemes are used, backup times may increase.</p> <p>SQL 2008 or later software compression is not used if a backup job that includes SQL 2008 or later data uses the Advanced Open File Option.</p>
Leave database ready to use. Additional transaction logs cannot be restored	<p>Lets the restore operation roll back all uncompleted transactions when you restore the last database, differential, or log backup. After the recovery operation completes, the database is ready for use. If Leave the database ready to use is not performed, the database is left in an intermediate state and is not usable.</p> <p>If you select the option when an intermediate backup is applied, you cannot continue to restore backups. You must restart the restore operation from the beginning.</p> <p>This option is selected by default.</p>

Table J-2 Microsoft SQL default options (*continued*)

Item	Description
<p>Leave the database nonoperational. Additional transaction logs or differential backups can be restored</p>	<p>Indicates that you have additional differential or transaction log backups to be restored in another restore job.</p>
<p>Leave the database in read-only mode</p>	<p>Creates and maintains a standby database during transaction log and database restores. See your SQL documentation for information on standby databases.</p>
<p>Consistency check after restore</p>	<p>Specifies one of the following consistency checks:</p> <ul style="list-style-type: none"> ■ None. This option is for sequential restores. Do not run a consistency check after a restore until all sequential restores have been done. If a consistency check is selected during a restore, the restore will complete but the consistency check will not be done. Check the job log for this information. If you selected the option Leave the database ready to use, select one of the following consistency checks: ■ Full check, excluding indexes. This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the nonclustered index pages is not checked. ■ Full check, including indexes. This option includes indexes in the consistency check. Any errors are logged. This option is selected by default. ■ Physical check only. This option performs a low overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page and record headers, and the consistency between the pages' object ID and index ID and the allocation structures.

Table J-2 Microsoft SQL default options (*continued*)

Item	Description
Overwrite the existing database	Replaces a database or file group, even if another database or file group with the same name already exists on the server. If Overwrite the existing database is not specified for a restore, SQL performs a safety check to ensure that a different database or file group is not accidentally overwritten. Refer to your SQL documentation for more information about the safety check that occurs when this option is not selected.

Setting backup options for SQL

This procedure details how to select backup job properties.

See [“About selecting data to back up”](#) on page 280.

See [“About selection lists”](#) on page 295.

See [“Creating a backup job by setting job properties”](#) on page 332.

To set backup job options for SQL

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Settings**, click **Microsoft SQL**.
- 4 Select the appropriate options.

See [“SQL backup options”](#) on page 1273.

SQL backup options

You can set the following options when you create a backup job for SQL.

See [“Setting backup options for SQL”](#) on page 1273.

See [“Backing up SQL databases”](#) on page 1283.

Table J-3 SQL backup options

Item	Description
Backup method	

Table J-3 SQL backup options (*continued*)

Item	Description
	<p>Specifies one of the following backup methods:</p> <ul style="list-style-type: none"> <li data-bbox="581 366 1243 487"> <p>■ Full - Back up entire database or filegroup. This option backs up the entire database or filegroup. This option is selected by default. See “About backing up SQL databases” on page 1281.</p> <li data-bbox="581 496 1243 913"> <p>■ Log - Back up transaction log. This option backs up only the data contained in the transaction log; it does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated). If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery mode before a log backup is run. Alternatively, you can run full backups only, or run full and differential backups of the SQL databases. Select this option to be able to select No recover - Place database in loading state or Standby - Place database in standby state under Enable advanced log backup options. See “Backing up SQL transaction logs” on page 1285.</p> <li data-bbox="581 921 1243 1269"> <p>■ Log No Truncate - Back up transaction log - no truncate. This option backs up the database when it is corrupt or database files are missing. Since the Log No Truncate method does not access the database, you can still back up transactions that you may not be able to access otherwise when the database is in this state. You can then use this transaction log backup along with the database backup and any previous transaction log backups to restore the database to the point at which it failed; however, any uncommitted transactions are rolled back. The Log No Truncate method does not remove committed transactions after the log is backed up. See “Backing up SQL transaction logs” on page 1285.</p> <li data-bbox="581 1277 1243 1459"> <p>■ Differential - Back up database or filegroup changes only. This option backs up only the changes made to the database or filegroup since the last full backup. Because differential backups allow the restore of a system only to the point in time that the differential backup was created, you should also create multiple log backups between the differential backups.</p> <li data-bbox="581 1468 1243 1581"> <p>■ Database Snapshot (SQL 2005 Enterprise Edition or later) - Read-only, point-in-time copy of another database. This option creates a read only, point-in-time copy of another database.</p>

Table J-3 SQL backup options (*continued*)

Item	Description
	<p>See “About SQL 2005 or later database snapshots” on page 1286.</p> <ul style="list-style-type: none"> ■ Full Copy-only (SQL 2005 or later) - Back up entire database or filegroup. <p>This option backs up the entire database or filegroup without affecting future differential or log backups.</p> <p>Unlike the Full backup method, the Full Copy-only backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.</p> <p>After making a full backup, you can use the Full Copy-only backup method to make a copy of a SQL database without affecting the baseline backup set required to run future differential backups.</p>
<p>Database snapshots to keep</p>	<p>(SQL 2005 or later) Displays the number of database snapshots to keep on disk. As the threshold is met, older database snapshots are deleted, which are then replaced with new snapshots. Because database snapshots continue to grow as the SQL database is updated, limiting the number of snapshots enables you to minimize both the disk space and SQL Server processing time that is required when the snapshots are updated.</p> <p>See “About SQL 2005 or later database snapshots” on page 1286.</p>

Table J-3 SQL backup options (*continued*)

Item	Description
Consistency check before backup	<p>Specifies one of the following consistency checks to run before a backup:</p> <ul style="list-style-type: none">■ None. This option does not run a consistency check before a backup. Symantec strongly recommends that you always run a consistency check either before or after the backup.■ Full check, excluding indexes. This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the nonclustered index pages is not checked.■ Full check, including indexes. This option includes indexes in the consistency check. Any errors are logged.■ Physical check only. This option performs a low overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page and record headers, and the consistency between the pages' object ID and index ID and the allocation structures. This option is selected by default. <p>See “About consistency checks for SQL” on page 1261.</p>
Continue with backup if consistency check fails	<p>Continues with the backup operation even if the consistency check fails. You may want to continue with the backup when the consistency check fails if you think that a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database with only a small problem in a table.</p>

Table J-3 SQL backup options (*continued*)

Item	Description
Consistency check after backup	<p>Specifies the consistency check to run after a backup. Because database transactions can occur during or after the consistency check, but before the backup runs, consider running a consistency check after the backup to ensure the data was consistent at the time of the backup.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> ■ None. This option does not run a consistency check after a backup. Symantec strongly recommends that you always run a consistency check either before or after the backup. This option is selected by default. ■ Full check, excluding indexes. This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the nonclustered index pages is not checked. ■ Full check, including indexes. This option includes indexes in the consistency check. Any errors are logged ■ Physical check only . This option performs a low overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page and record headers, and the consistency between the pages' object ID and index ID and the allocation structures.
Enable advanced log backup options	<p>Enables either the No Recover - Place database in loading state option or the Standby - place database in standby state option to apply to the backup.</p> <p>This option is only available after you select the backup method Log - Back up transaction log.</p>
No recover - Place database in loading state	<p>Places the database in a loading state when the log file backup completes. Users cannot connect to or query the database while it is in a loading state.</p> <p>This option is only available after you select Enable advanced log backup options.</p>

Table J-3 SQL backup options (*continued*)

Item	Description
Standby - Place database in standby state	<p>Places the database in standby mode when the log file backup completes. Users can connect to and query the database when it is in standby mode, but cannot update it.</p> <p>You can convert a standby database to a live database by restoring the latest transaction log. Ensure that you select the following recovery completion state Leave the database ready to use; additional transaction logs or differential backup cannot be restored.</p> <p>This option is only available if Enable advanced log backup options has been selected.</p>
Use checksums on backup (SQL 2005 or later)	<p>Adds checksums to the SQL database data being backed up by Backup Exec. Adding checksums to the data being backed up is required if you want to use the option Run verify only; do not restore data. Using this option, along with Run verify only; do not restore data, ensures that during a restore of the SQL database, you are restoring from a verified SQL backup.</p>
Create on-disk copies of SQL backups to be placed on the SQL server where the database is located	<p>Creates an on-disk copy of the SQL database being backed up. This option lets you simultaneously back up a SQL database to storage media while also writing a copy of the database to a disk path you specify in the Save to path box.</p> <p>This option gives IT administrators the ability to back up SQL databases while also providing database administrators with copies of the database on disk, which can be used for such things as tests and restores.</p>
Save to path	<p>Displays a path in which to save on-disk copies of SQL backups.</p>

Table J-3 SQL backup options (*continued*)

Item	Description
<p>SQL Server (2008 or later) software compression</p>	<p>Specifies the following compression setting you want to use for this backup job:</p> <ul style="list-style-type: none"> ■ None. Do not use compression. ■ Compress. Use SQL Server 2008 or later compression. <p>SQL compresses the data on the computer on which SQL Server 2008 Enterprise Edition or later is installed. Therefore, faster SQL 2008 or later backups should occur if you use SQL compression.</p> <p>If you back up remote SQL 2008 or later computers and you use SQL 2008 or later software compression, you must use the latest version of the Remote Agent.</p> <p>You can find a list of compatible operating systems, platforms, and applications at the following URL:</p> <p>http://entsupport.symantec.com/umi/V-269-1</p> <p>Symantec recommends that you do not use SQL 2008 or later software compression in a backup job that uses Backup Exec-initiated software compression. Minimal additional SQL 2008 or later compression benefits are gained when you enable Backup Exec compression. In fact, in jobs where both compression schemes are used, backup times may increase.</p> <p>SQL 2008 or later software compression is not used if a backup job that includes SQL 2008 or later data uses the Advanced Open File Option.</p>
<p>Guide Me</p>	<p>Starts a wizard that helps you select backup job properties for SQL.</p>

About automatic exclusion of SQL data during volume level backup

If you select a volume that contains SQL data for backup, the SQL Agent determines which SQL data should not be included in a volume level backup. For example, .MDF and .IDF files should not be part of the backup because they are opened for exclusive use by the SQL system. These files will be automatically excluded for backup by a feature called Active File Exclusion. If this exclusion did not happen during a non-snapshot backup, these files would appear as in use - skipped. If this exclusion did not happen during a snapshot backup, the files would be backed up in a possible inconsistent state, which could create restore issues.

While it is not recommended, if you want to include SQL data in a volume level backup, you must first dismount the database you want backed up. Then, run the backup job.

See [“About backing up SQL databases”](#) on page 1281.

See [“About backing up SQL filegroups”](#) on page 1282.

See [“Displaying SQL filegroups on the backup selections pane”](#) on page 1284.

About backing up SQL databases

Backup Exec includes three methods for backing up databases: Full, Differential, and for SQL 2005 or later, Full Copy-only. The full method backs up the entire database including all system tables and filegroups. The differential method backs up only the changes made to the database since the last full backup. The copy method works in the same manner as the full method, except that it does not affect future differential or log backups.

A differential backup is smaller and faster than a full backup, so differential backups can be run more often than full backups. Because differential backups allow the restore of a system only to the point that the differential backup was created, you should also create multiple log backups between the differential backups. Using transaction log backups allows you to recover the database to the exact point of failure.

Consider using differential backups when only a relatively small amount of data changes between full backups, or if the same data changes often. Differential backups may also work well in your environment if you are using the simple recovery model and need backups more often, but cannot spare the time to do frequent full backups. If you are using the full or bulk-logged recovery models, you can use differential backups to decrease the time it takes to roll forward log backups when restoring a database.

If you want to run database backups only, instead of a mix of database and log backups, use the simple recovery model for the database so that the transaction log is automatically truncated when a checkpoint occurs in the database. This helps prevent transaction logs from becoming full since with other recovery models the logs are not cleared after a database backup.

With the simple recovery model, copies of the transactions are not stored in the log file, which prevents transaction log backups from being run.

If you do not run transaction log backups, you can recover the database to the point of the last backup, but you cannot restore the database to the point of failure or to a specific point in time.

The master database can only be backed up with the full method; you cannot use the log or differential methods to back up the master database.

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux Servers is installed.

The SQL Agent supports a mirrored SQL database configuration, although Microsoft places limitations on the mirroring of SQL databases.

These limitations include the following:

- You cannot back up or restore a mirrored SQL database. If you attempt to back up or restore a mirrored database, the backup job or restore job fails.
- You cannot restore the primary SQL database while it is configured in a mirrored configuration. To restore the primary SQL database, you must stop database mirroring of the primary database.
- You can back up a primary SQL database and its transaction logs only if the backup job does not leave the database in a non-recovered state.

See [“Backing up SQL databases”](#) on page 1283.

About backing up SQL filegroups

When databases grow too large to be backed up all at once, filegroups can provide an alternative backup method. Filegroups can be backed up at different times and frequencies. Filegroups that change often can be backed up more frequently than filegroups that remain more static. In certain situations, filegroup backups can greatly reduce restore time. For example, if a nonprimary filegroup is destroyed or corrupted, only that filegroup has to be restored.

Backup Exec includes two methods for backing up filegroups: full and differential. The full method backs up the entire filegroup. The differential method backs up only the changes made to the filegroup since the last full backup. A differential backup is smaller and faster than a full backup, so differential backups can be run more often than full backups. Consider using differential backups when only a relatively small amount of data changes between full filegroup backups, or if the same data changes often.

Because differential backups allow the restore of a system only to the point that the differential backup was created, you should also create multiple log backups between the differential backups. Using transaction log backups enables you to recover the filegroup to the exact point of failure. A combination of full and differential filegroup backups and transaction log backups provides complete database protection. Log backups are required to be able to recover a system from a filegroup backup.

See [“Displaying SQL filegroups on the backup selections pane”](#) on page 1284.

See [“Backing up SQL filegroups”](#) on page 1284.

Backing up SQL databases

The following procedure provides details on how to back up SQL databases.

See [“About backing up SQL databases”](#) on page 1281.

See [“Backing up SQL transaction logs”](#) on page 1285.

To back up SQL databases

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**, and then select the data you want to back up.

See [“How to use Backup Exec logon accounts for SQL resources”](#) on page 1256.

- 4 To select SQL data from local or remote selections, click the domain name icon or icons that contain the SQL installations, and then click the actual Windows computer icon that contains the SQL installation. If you are using a cluster, make backup selections from the virtual server.

A list of shared network directories appears, along with an icon that represents the SQL installation.

To select all databases in SQL, click the check box preceding the SQL icon, or you can select specific databases by clicking the SQL icon, and then selecting individual databases.

Whether you make SQL database selections using the Windows domain, Active Directory, DNS names, or IP addresses, you must use the same method when making full, differential, and incremental backups of your SQL databases. For example, do not make full backup selections of your SQL databases using the Windows domain, and then make incremental or differential selections using an IP address.

- 5 To select the SQL backup job properties, on the **Properties** pane, under **Settings**, click **Microsoft SQL**.
- 6 Select options for the backup job.
See [“SQL backup options”](#) on page 1273.
- 7 Start the backup job or select other backup options from the **Properties** pane, and then start the backup job.

See [“Creating a backup job by setting job properties”](#) on page 332.

Displaying SQL filegroups on the backup selections pane

Filegroups do not appear on the backup selections pane by default.

To display filegroups on the backup selections pane

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Microsoft SQL**.
- 3 Select **Display filegroups when creating new backup jobs**.

See [“Backing up SQL filegroups”](#) on page 1284.

See [“Setting backup options for SQL”](#) on page 1273.

Backing up SQL filegroups

You can back up specific filegroups.

See [“About backing up SQL filegroups”](#) on page 1282.

To back up SQL filegroups

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**, and then select the data you want to back up.

See [“How to use Backup Exec logon accounts for SQL resources”](#) on page 1256.

- 4 To select SQL data, click the domain name icon or icons that contain the SQL installations, and then click the actual Windows computer icon that contains the SQL installation. If you are using a cluster, make backup selections from the virtual server.

- 5 Click a SQL container and select specific filegroups.

- 6 To select SQL backup job properties, on the **Properties** pane, under **Settings**, click **Microsoft SQL**.

- 7 Select options for the backup job.

See [“SQL backup options”](#) on page 1273.

- 8 Start the backup job or select other backup options from the **Properties** pane, and then start the backup job.

See [“Creating a backup job by setting job properties”](#) on page 332.

How to back up SQL transaction logs

Backup Exec includes two methods for backing up transaction logs: Log and Log No Truncate.

When running log backups, it is recommended that you use Backup Exec exclusively to perform log truncations if you decide to truncate the logs. After a transaction log has been truncated by something other than a log backup, you must run a full or differential backup before you run another log backup.

Use the Log No Truncate method only when the database is corrupted or database files are missing. This method backs up transactions that you may not be able to access otherwise when the database is in this state. You can then use this transaction log backup along with the last database backup and any previous transaction log backups to restore the database to the point at which it failed; however, any uncommitted transactions are rolled back. The Log No Truncate method does not remove committed transactions after the log is backed up.

To use the Log No Truncate backup to restore a database, you should also have a database backup that was created before the Log No Truncate backup. The transaction log contains only the log files used in the restore process, which alone are not sufficient to restore a complete database. You must have at least one database backup or a full set of filegroup backups and a log backup of the database to restore a database.

Caution: Do not run a log backup using either method if the SQL database is using the simple recovery model. With the simple recovery model, you can recover data only up to the most recent full or differential backup. If you run a log backup on a database using the simple recovery completion state, the backup will fail.

To check the database properties, from the Database management tools on the SQL Server, right-click the database, click Properties, click the Options tab, and then view the configuration settings.

See [“Backing up SQL transaction logs”](#) on page 1285.

Backing up SQL transaction logs

You can use Backup Exec to back up SQL transaction logs.

See [“How to back up SQL transaction logs”](#) on page 1285.

See [“About restoring SQL databases and file groups”](#) on page 1293.

See [“Restoring from SQL transaction logs up to a point in time ”](#) on page 1296.

See [“Restoring from SQL transaction logs up to a named transaction ”](#) on page 1297.

To back up SQL transaction logs

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**, and then select the data you want to back up.
See [“How to use Backup Exec logon accounts for SQL resources”](#) on page 1256.
- 4 To select SQL data, click the domain name icon or icons that contain the SQL installations, and then click the actual Windows computer icon that contains the SQL installation. If you are using a cluster, make backup selections from the virtual server.
A list of shared network directories appears, along with an icon that represents the SQL installation.
- 5 To select SQL backup job properties, on the **Properties** pane, under **Settings**, click **Microsoft SQL**.
- 6 Select options for the backup job.
See [“SQL backup options”](#) on page 1273.
- 7 Start the backup job or select other backup options from the **Properties** pane, and then start the backup job.
See [“Creating a backup job by setting job properties”](#) on page 332.

About SQL 2005 or later database snapshots

SQL database snapshots enable you to quickly revert a database back to the state it was in when the database snapshot was created. When you use a database snapshot, a full restore of the host database is not required to revert the database. However, that changes made to the host between the time a database snapshot is created and the point at which it is reverted, are lost.

The Backup Exec SQL Agent works with the SQL database to create database snapshots, which are read-only, point-in-time copies of an existing host database. When Backup Exec runs a SQL backup job using the Database Snapshot (SQL 2005 or later) backup method, a request is sent to the host database instructing it to create a database snapshot.

Note: The snapshot backup method for SQL databases is only supported by SQL Server Enterprise Edition (versions 2005 or later).

Database snapshots cannot be backed up to storage media. Rather, they are written to a SQL snapshot file on disk. After running the database snapshot job, Backup Exec creates history and job log information to indicate the job's status.

Because database snapshots cannot be backed up, all database snapshots will be lost if the disk where the host database is installed fails. Therefore database snapshots should not be used as your sole database protection strategy. They should be used in conjunction with an overall Backup Exec database protection strategy that includes full, differential, and transaction log backups of the SQL database.

For more information, see your Microsoft SQL documentation.

Note: SQL database snapshots are not the same as Microsoft Virtual Shadow Copy Service (VSS) snapshots. Whereas VSS snapshots enable you to create point-in-time snapshots of disk volumes and shares, database snapshots enable you to create point-in-time copies of SQL databases. You cannot use the VSS option in Backup Exec's Advanced Open File Option to create SQL database snapshots.

Note: SQL database snapshot catalog information that refers to deleted database snapshots is periodically removed from the catalogs. If backup media is re-cataloged, the database snapshot catalog information will be periodically removed again.

See [“About the database snapshot \(SQL 2005 or later\) backup method”](#) on page 1287.

About the database snapshot (SQL 2005 or later) backup method

The backup method, Database Snapshot (SQL 2005 or later), enables you to do the following:

- Produce SQL database snapshots.
- Set the number of SQL database snapshots to keep on disk.

After selecting the Database Snapshot (SQL 2005 or later) backup method, an option appears called Database snapshots to keep. This option enables you to set the number of database snapshots to be kept on disk for each database. As the threshold is met, older database snapshots are deleted, which are then replaced with new snapshots. Because database snapshots continue to grow as the SQL database is updated, limiting the number of snapshots enables you to minimize both the disk space and SQL Server processing time that is required when the snapshots are updated.

By limiting the number of database snapshots that are kept, you can configure a database protection strategy that minimizes data loss in the event of a host database problem. For example, you can create a strategy that protects the SQL database from inadvertent table deletions. This strategy consists of a Backup Exec database snapshot job that you schedule to run once every hour during a 24 hour period. As part of the strategy, you also configure the job to keep four database snapshots. Because the job is scheduled to run every hour, a new database snapshot is created every hour. Beginning with the fifth hour and going forward, the oldest database snapshot is automatically deleted before a new one is created. Throughout the 24 hour period, there are no more than four database snapshots on disk. If a user deletes a database table, you first determine the time the table was deleted and then run a Backup Exec database snapshot restore job that enables you to revert the host to one of four previous points in time going back four hours. Remember, however, that any changes made to the host between the time a database snapshot is created and the point at which it is reverted, are lost.

See [“Creating SQL database snapshots”](#) on page 1288.

See [“About reverting SQL 2005 or later databases using database snapshots ”](#) on page 1309.

Creating SQL database snapshots

The Backup Exec SQL Agent works with the SQL 2005 or later database to create database snapshots, which are read-only, point-in-time copies of an existing host database.

See [“About SQL 2005 or later database snapshots”](#) on page 1286.

To create database snapshots

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 Select the SQL database for which you want to create a database snapshot.
- 5 On the **Properties** pane, under **Settings**, click **Microsoft SQL**.
- 6 In the Backup method field, select **Database Snapshot (SQL 2005 or later) - Read only, point-in-time copy of another database**.
- 7 Set the number of database snapshots per database to keep, or accept the default of four.
- 8 Start the database snapshot job or select other backup options from the **Properties** pane, and then start the backup job.

Setting restore options for SQL

This procedure details how to select restore job properties for SQL, and provides definitions for SQL-specific restore options. For details on how to create a restore job, and for definitions of all other restore options:

See [“Restoring data by setting job properties”](#) on page 605.

See [“About restoring SQL databases and file groups”](#) on page 1293.

See [“Restoring from SQL transaction logs up to a point in time ”](#) on page 1296.

See [“Restoring from SQL transaction logs up to a named transaction ”](#) on page 1297.

See [“About restoring from SQL filegroup backups ”](#) on page 1298.

To set restore options for SQL

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Settings**, click **Microsoft SQL**.
- 4 Select the appropriate options.

See [“SQL restore options”](#) on page 1289.

SQL restore options

You can set the following options when you create a restore job for SQL.

See [“SQL restore options”](#) on page 1289.

See [“Restoring from SQL database backups”](#) on page 1295.

See [“Restoring from SQL transaction logs up to a point in time ”](#) on page 1296.

Table J-4 SQL restore options

Item	Description
Leave the database ready to use; additional transaction logs or differential backups cannot be restored	Lets the restore operation roll back all uncompleted transactions when you restore the last database, differential, or log backup. After the recovery operation, the database is ready for use. If you do not select this option, the database is left in an intermediate state and is not usable. If you select this option, you cannot continue to restore backups. You must restart the restore operation from the beginning.

Table J-4 SQL restore options (*continued*)

Item	Description
Leave the database nonoperational; additional transaction logs or differential backups can be restored	Indicates that you have additional differential or transaction log backups to be restored in another restore job.
Leave the database in read-only mode	Creates and maintains a standby database during a transaction log and database restore. See your SQL documentation for information on standby databases.
Take existing destination database offline	Lets Backup Exec automatically take the database offline before the restore job runs. If this option is not selected and there are active connections to the SQL database, the restore job will fail.
Overwrite the existing database	Replaces a database or filegroup, even if another database or filegroup with the same name already exists on the server. If Overwrite the existing database is not specified for a restore, SQL performs a safety check to ensure that a different database or filegroup is not accidentally overwritten. Refer to your SQL documentation for more information about the safety check that occurs when this option is not selected.
Automate master database restore	<p>Lets Backup Exec stop SQL so that the master database can be restored. All existing users are logged off, and SQL Server is put into single-user mode.</p> <p>When this option is selected, only the master database can be restored; if this option is selected for any other database, those jobs will fail.</p> <p>If Backup Exec does not have access to the SQL registry keys, HKEY_LOCAL_MACHINE\Software\Microsoft\Microsoft SQL Server, and HKEY_LOCAL_MACHINE\Software\Microsoft\MSSQLServer, then a restore to the default directory may not work, and the Automate master database restore option on the restore job properties for SQL will not work. To ensure that Backup Exec has access rights, verify that the logon account used has administrator rights to the Windows server that the SQL instance is installed on.</p>

Table J-4 SQL restore options (*continued*)

Item	Description
Continue restoring if an error occurs during the restore (SQL 2005 or later)	Lets Backup Exec restore as much of the SQL database as possible if SQL detects database corruption errors during the database restore.
Run verify only; do not restore data	<p>Lets SQL verify your SQL backup jobs. This option returns the entire Backup Exec SQL data stream directly to SQL for verification. Although SQL processes the data stream for errors, existing SQL databases are not affected; all verification processes are handled within SQL itself, and nothing is ever written to the disk.</p> <p>As SQL processes the data streams, a slight performance impact on overall database performance occurs until the verification process finishes.</p> <p>Although supported in SQL 2000, this option's best performance occurs with the Backup Exec SQL backup option, Use checksum on backups (SQL 2005 or later). See "Setting backup options for SQL" on page 1273.</p>

Table J-4 SQL restore options (*continued*)

Item	Description
<p>Consistency check after restore</p>	<p>Specifies one of the following options:</p> <ul style="list-style-type: none"> ■ Full check, excluding indexes. Excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the nonclustered index pages is not checked. ■ Full check, including indexes. Includes indexes in the consistency check. Any errors are logged. This option is selected by default. ■ Physical check only. Performs a low overhead check of the physical consistency of the SQL 2000 database. This option only checks the integrity of the physical structure of the page and record headers, and the consistency between the pages' object ID and index ID and the allocation structures. ■ None. This option is for sequential restores. Do not run a consistency check after a restore until all sequential restores have been done. If a consistency check is selected during a restore, the restore will complete but the consistency check will not be done. Check the job log for this information. <p>If you need to recover the database after restores are complete, select one of the consistency checks mentioned above.</p>
<p>Recover the entire log</p>	<p>Recovers all of the transactions in the transaction logs you select for restore.</p>
<p>Point in time log restore</p>	<p>Restores transactions from a transaction log up to and including a point in time in the transaction log. After the point in time, recovery from the transaction log is stopped.</p> <p>In the Date box, select the part of the date you want to change, and then enter a new date or click the arrow to display a calendar from which you can select a date.</p> <p>In the Time box, select the part of the time you want to change, and then enter a new time or click the arrows to select a new time.</p>

Table J-4 SQL restore options (*continued*)

Item	Description
Restore log up to named transaction	<p>Restores transactions from a transaction log up to a named transaction (or named mark) in the transaction log; after that, recovery from the transaction log is stopped. The named transactions are case-sensitive.</p> <p>Check your client application event log to find dates and times of named transactions.</p>
Include the named transaction	<p>Includes the named transaction in the restore; otherwise the restore will stop immediately before the named transaction is restored.</p> <p>This option is only available if you select the Restore log up to named transaction option.</p>
Found after	<p>Specifies a date and time after which the restore operation is to search for the named transaction. For example, if you specify a restore from a log up to the named transaction AfternoonBreak, found after 6/02/2000, 12:01 p.m., then the restore operation will not search for AfternoonBreak until after that time.</p> <p>This option is only available if you selected the Restore log up to named transaction option.</p>
Check selections	<p>Lets Backup Exec verify or complete the selections required to successfully restore SQL databases. After making your database restore selections, use this feature to verify the database selections are valid. If there are selection issues, Backup Exec notifies you of the error or errors and then corrects them for you.</p>
Guide Me	<p>Starts a wizard that helps you select restore job properties for SQL.</p>

About restoring SQL databases and file groups

You can restore a database by using one job or using multiple jobs to restore all of the backup sets. The number of jobs you decide on depends on the types of backup jobs that protect the database or the file group. If you use one job to restore a database, select all the backup sets that you want to apply. Include the full backup, any differential backups, and any log backups. Also select the Leave the database ready to use option. Additional transaction logs cannot be restored.

Single-job restores and multiple-job restores can both be used in redirected restore operations.

Some restore operations must be completed using separate restore jobs to recover data.

These operations include the following:

- Restoring a database or a primary filegroup from a filegroup backup. Separate restore jobs must be used to restore the primary filegroup, to restore the rest of the filegroup backup sets, and to restore the transaction logs.
- Restoring a nonprimary filegroup. After running a Log No Truncate backup, separate restore jobs must be used to restore the missing filegroup from full and differential backups of the filegroups, and to restore the transaction logs.

If you use multiple jobs to restore a database, ensure that you specify the recovery completion state `Leave the database nonoperational`. Additional transaction logs can be restored for all the jobs except the last one. For the last job, you should specify the recovery completion state `Leave the database ready to use`. If you use this recovery state, additional transaction logs cannot be restored.

SQL database files contain unused space so that the disk file does not have to be grown every time a small amount of data is added to the database. SQL fills the unused space with zeros. When SQL databases are restored, it is not known how much of the file will actually be used by the restored data, so SQL creates the required database files on disk and then fills them with zeros.

With very large databases this process can take several hours to complete. During this time Backup Exec reports that no data is being transferred, and the Byte count field in the Job Monitor view is not updated. When SQL has completed filling the files with zeros, the restore job continues. This occurs for all database restores but is noticeable only on very large databases.

In a mirrored configuration, the primary SQL database cannot be restored. To restore the primary SQL database, you must stop database mirroring of the primary database.

See [“Restoring data by setting job properties”](#) on page 605.

See [“Restoring from SQL transaction logs up to a point in time ”](#) on page 1296.

See [“Restoring from SQL transaction logs up to a named transaction ”](#) on page 1297.

See [“About restoring from SQL filegroup backups ”](#) on page 1298.

See [“Redirecting restores for SQL”](#) on page 1305.

About restoring encrypted SQL databases

SQL 2008 supports Transparent Database Encryption (TDE), which lets you encrypt SQL 2008 databases at the backup set level.

When you back up a database that uses TDE, Microsoft recommends that you back up the certificate keys and encryption keys with the database. If you do not include the certificate keys and encryption keys, you must perform all backup and restore operations within the selected SQL instance.

Note: Backup Exec can redirect the restore of the database data that used TDE only if the certificate keys and encryption keys are applied to the destination instance. If the certificate keys and encryption keys are not applied to the destination instance, an error appears stating that the certificate thumbprint cannot be found.

See your Microsoft SQL 2008 documentation.

Restoring from SQL database backups

If the database is using the simple recovery model, there are no transaction log backups to restore. You only need to restore the most recent full database backup and if you were running differential database backups, restore the most recent differential database backup.

See [“Restoring data by setting job properties”](#) on page 605.

Note: Restoring a full SQL 2005 or later database backup over an existing SQL 2005 or later database with active database snapshots will eliminate all existing database snapshots for the SQL 2005 or later database being restored.

See [“About reverting SQL 2005 or later databases using database snapshots”](#) on page 1309.

See [“About restoring SQL databases and file groups”](#) on page 1293.

See [“About SQL 2005 or later database snapshots”](#) on page 1286.

See [“About reverting SQL 2005 or later databases using database snapshots”](#) on page 1309.

To restore from SQL database backups

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.

- 4 In the restore selections list, select the most recent full database backup set, and the most recent differential database backup set, if any, to restore.
- 5 On the **Properties** pane, under **Settings**, click **Microsoft SQL**.
- 6 On the **Restore Job Properties** dialog box, click **Leave the database ready to use; additional transaction logs or differential backups cannot be restored**.
See [“SQL restore options”](#) on page 1289.
- 7 Start the restore job or select other restore options from the **Properties** pane.

How to restore from SQL transaction logs up to a point in time

You can restore transactions from a transaction log up to and including a point in time in the transaction log. After the point in time is reached, recovery from the transaction log is stopped. To find dates and times of transactions, check your client application event log.

If the specified point in time is later than the time contained in the most recent transaction log being restored, then the restore operation succeeds, but a warning is generated and the database remains in an intermediate state. If the specified point in time is before the time contained in the transaction log or logs being restored, no transactions are restored.

See [“Restoring from SQL transaction logs up to a point in time ”](#) on page 1296.

Restoring from SQL transaction logs up to a point in time

The following procedure provides details on how to restore transactions from a transaction log up to and including a point in time in the transaction log.

See [“How to restore from SQL transaction logs up to a point in time”](#) on page 1296.

See [“Restoring data by setting job properties”](#) on page 605.

See [“About restoring SQL databases and file groups”](#) on page 1293.

See [“Restoring from SQL transaction logs up to a named transaction ”](#) on page 1297.

To restore from SQL transaction logs up to a point in time

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 In the restore selections list, select the most recent full database backup set, and the most recent differential database backup set, if any, and all the log backup sets you want to restore.
- 5 On the **Properties** pane, under **Settings**, click **Microsoft SQL**.

- 6 On the **Restore Job Properties** dialog box, click **Leave the database ready to use; additional transaction logs or differential backups cannot be restored**.
- 7 Select **Point in time log restore**, and then select a date and time.
See [“SQL restore options”](#) on page 1289.
- 8 Start the restore job or select other restore options from the **Properties** pane.
See [“Restoring data by setting job properties”](#) on page 605.

How to restore from SQL transaction logs up to a named transaction

You can restore transactions from a transaction log up to and including a named transaction (or mark). After the named transaction is reached, recovery from the transaction log is stopped.

Since named transactions do not necessarily have unique names, you can also specify a date and time after which the restore operation is to search for the named transaction. For example, if you specify a restore from a log up to the named transaction `AfternoonBreak`, found after 6/02/2000, 12:01 p.m., then the restore operation will not search for `AfternoonBreak` until after that time. To find dates and times of named transactions, check your client application event log.

If the named transaction is not found, then the restore operation succeeds, but a warning is generated and the database remains in an intermediate state.

The names of transactions are case-sensitive. Make sure you enter the correct upper- and lower-case characters when specifying a named transaction.

See [“How to restore from SQL transaction logs up to a named transaction”](#) on page 1297.

Restoring from SQL transaction logs up to a named transaction

The following procedure provides details on how to restore transactions from a transaction log up to and including a named transaction (or mark).

See [“How to restore from SQL transaction logs up to a named transaction”](#) on page 1297.

See [“About restoring SQL databases and file groups”](#) on page 1293.

To restore from SQL transaction logs up to a named transaction

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.

- 4 In the restore selections list, select the most recent full database backup set, and the most recent differential database backup set, if any, and all the log backup sets you want to restore.
- 5 On the **Properties** pane, under **Settings**, click **Microsoft SQL**.
- 6 On **Restore Job Properties** dialog box, click **Leave the database ready to use; additional transaction logs or differential backups cannot be restored**.
- 7 Select **Restore log up to named transaction**, and then enter the name of the transaction.

The names are case-sensitive. Make sure you enter the correct upper- and lower-case characters.
- 8 To include the named transaction in the restore, select **Include the named transaction**.
- 9 To specify a particular named transaction in the log, select **Found after** and then select a date and time.

If a date and time are not entered, recovery from the transaction log is stopped at the first transaction with the specified name.

See [“Redirecting restores for SQL”](#) on page 1305.
- 10 Start the restore job or select other restore options from the **Properties** pane.

See [“Restoring data by setting job properties”](#) on page 605.

About restoring from SQL filegroup backups

With filegroup backups, you can restore the entire database, a primary filegroup, a filegroup containing a deleted or changed table, and a nonprimary filegroup.

The following are conditions for filegroup restores:

- All filegroups must be restored to the same point in time. For example, if a table is deleted from a filegroup, you cannot restore that filegroup to a point in time before the table was deleted and then leave it at that time; you must continue restoring the filegroup to the same point in time shared by all existing filegroups.

To be able to restore a filegroup to the same point in time as the other filegroups, run one of the following log backups:

- If the database is intact, run a Log backup.
- If any files or filegroups are missing, run a Log - No Truncate backup.

Note: If the primary filegroup is missing, the log backup methods are unavailable. You can restore the database only up to the last log backup.

- Filegroup restores can be redirected to a different server, but the database file paths cannot be changed. For example, if the filegroup was backed up from G:\SQLDATA then the filegroup must be restored to G:\SQLDATA, regardless of the server the restore is redirected to.
The options Restore all databases to default drive and Restore all database files to the target instance's data location on the Restore Job Properties for SQL dialog box do not apply to filegroup restores. Filegroups must be restored to the same drive letter and path that they were backed up from.
- When restoring from filegroup backups, separate restore jobs are required.
- Previous versions of Backup Exec cannot restore filegroup backups made with this release of Backup Exec.

See [“Restoring an entire SQL database, a missing primary filegroup, or a filegroup containing a deleted or changed table”](#) on page 1299.

See [“Restoring a missing or corrupted SQL nonprimary filegroup”](#) on page 1300.

Restoring an entire SQL database, a missing primary filegroup, or a filegroup containing a deleted or changed table

Use the following steps to restore an entire database, a missing primary filegroup or a filegroup containing a deleted or changed table.

Use separate restore jobs to restore the primary filegroup, the rest of the filegroup backup sets, and the transaction logs.

See [“Restoring a missing or corrupted SQL nonprimary filegroup”](#) on page 1300.

See [“About restoring SQL databases and file groups”](#) on page 1293.

To restore the entire database, a missing primary filegroup, or a filegroup containing a deleted or changed table

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 In the restore selections list, select the backup set containing the primary filegroup.
- 5 On the **Properties** pane, under **Settings**, click **Microsoft SQL**.

- 6 On the **Restore Job Properties** dialog box, click **Leave the database nonoperational; additional transaction logs or differential backups can be restored**.
See [“Redirecting restores for SQL”](#) on page 1305.
- 7 Start the restore job.
- 8 After the primary filegroup is restored, select the rest of the filegroup backup sets containing the latest full and differential backups.
- 9 On the **Restore Job Properties** dialog box, click **Leave the database nonoperational; additional transaction logs or differential backups can be restored**, and then start the restore job.
- 10 When the full and differential backups are restored, select the backup set containing the transaction logs.
- 11 On the **Restore Job Properties** dialog box, click **Leave the database nonoperational; additional transaction logs or differential backups can be restored**.
This option restores all of the transaction logs. You also can select Point in time log restore or Restore log up to named transaction.
- 12 Start the restore job or select other restore options from the **Properties** pane.
See [“Restoring data by setting job properties”](#) on page 605.

Restoring a missing or corrupted SQL nonprimary filegroup

Use the following steps to restore a missing or corrupted nonprimary filegroup.

See [“Restoring an entire SQL database, a missing primary filegroup, or a filegroup containing a deleted or changed table”](#) on page 1299.

See [“About restoring SQL databases and file groups”](#) on page 1293.

To restore a missing or corrupted nonprimary filegroup

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 In the backup selections list, select the database.
- 5 On the **Properties** pane, under **Settings**, click **Microsoft SQL**.
- 6 Select the backup method **Log No Truncate**, select **None** for a consistency check, and then start the backup job.

- 7 After the Log No Truncate backup is complete, restore the missing or corrupted filegroup by selecting the filegroup backup sets containing the latest full and differential backups, and the transaction log backups.
- 8 On the **Restore Job Properties** dialog box, click **Leave the database ready to use; additional transaction logs or differential backups cannot be restored.**
- 9 Start the restore job or select other restore options from the **Properties** pane, and then start the restore job.

See [“Restoring data by setting job properties”](#) on page 605.

About restoring the SQL master database

If the master database is damaged, symptoms may include the following:

- An inability to start SQL.
- Segmentation faults or input/output errors.
- A report generated by SQL Database Consistency Checker utility (DBCC).

If you can still start SQL, you can restore the latest copy of the master database backup using the Automate master database restore option in Backup Exec’s Restore Job Properties for SQL dialog box and then restore any other databases, if needed.

If the master database is critically damaged and SQL cannot be started, rather than running the Rebuild Master utility, or reinstalling SQL to be able to restart SQL, you can replace the corrupted or missing databases with the copies of the master and model databases that Backup Exec automatically creates and updates whenever backups of those databases are run. After SQL is running again, you can restore the latest copy of the master database using Backup Exec’s Automate master database restore option, and then restore any other databases, if needed.

If copies of the master and model databases were not made, then you must use Microsoft’s rebuildm.exe utility to rebuild the master database and start SQL.

Because all changes made to the master database after the last backup was created are lost when the backup is restored, the changes must be reapplied. If any user databases were created after the master database was backed up, those databases cannot be accessed until the databases are restored from backups or reattached to SQL.

See [“Restarting SQL using database copies”](#) on page 1302.

See [“Restoring the master database”](#) on page 1304.

Restarting SQL using database copies

You can restart SQL manually using copies of the database from previous backups and then restore the master database.

See [“Restoring the master database”](#) on page 1304.

Table J-5 Restarting SQL using database copies

Step	Action
Step 1	Ensure that the SQL services are not running. Refer to the SQL Server documentation for details.
Step 2	Verify that the database copies are present. See “SQL database copy locations” on page 1303. If necessary, restore the master and model database copies from a backup set to the same directory that the original master and model databases are in.
Step 3	Using the Windows Explorer, browse to the default data directory and delete the following files: <ul style="list-style-type: none">■ master.mdf■ mastlog.ldf■ model.mdf■ modellog.ldf.
Step 4	Rename the copies of the databases back to their original names. See “SQL database names” on page 1303. Do not use read-only files. The SQL services will not start with read-only files.
Step 5	Use the SQL Service Control Manager to start SQL Server.

Table J-5 Restarting SQL using database copies (*continued*)

Step	Action
Step 6	Restore the latest changes to the master database. See “Restoring the master database” on page 1304.

SQL database copy locations

The database copies are named master\$4idr, mastlog\$4idr, model\$4idr, and modellog\$4idr.

See [“Restarting SQL using database copies”](#) on page 1302.

Table J-6 SQL database copy locations

SQL database copy	Location
A default installation of SQL 2000	C:\Program Files\Microsoft SQL Server\MSSQL\Data*.*
A named instance of SQL 2000	C:\Program Files\Microsoft SQL Server\MSSQL\$Instance_Name\Data*.*
An initial installation of SQL 2005 or later	C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Data*.*
A second installed instance of SQL 2005 or later	C:\Program Files\Microsoft SQL Server\MSSQL.2\MSSQL\Data*.*
A default installation of SQL 2008	C:\Program Files\Microsoft SQL Server\MSSQL10.<instance name>\MSSQL\Data

SQL database names

The following table lists the copied database name and the original database name.

See [“Restarting SQL using database copies”](#) on page 1302.

Table J-7 SQL database names

Copied database name	Original database name
master\$4idr	master.mdf
master\$4idr	mastlog.ldf

Table J-7 SQL database names (*continued*)

Copied database name	Original database name
model\$4idr	model.mdf
modellog\$4idr	modellog.ldf

Restoring the master database

You can restore the master database after you restart SQL using database copies.

See [“About restoring the SQL master database”](#) on page 1301.

See [“Restarting SQL using database copies”](#) on page 1302.

To restore the master database

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 On the restore selections list, select the backup set containing the last master database backup.
- 5 On the **Properties** pane, under **Settings**, click **Microsoft SQL**.
- 6 On the **Restore Job Properties for SQL** dialog box, select **Automate master database restore**.

All existing users are logged off, and SQL Server is put into single-user mode.

When this option is selected, only the master database can be restored; if this option is selected for any other database, those jobs will fail.

If Backup Exec does not have access to the SQL registry keys HKEY_LOCAL_MACHINE\Software\Microsoft\Microsoft SQL Server and HKEY_LOCAL_MACHINE\Software\Microsoft\MSSQLServer, then a restore to the default directory may not work, and the option Automate master database restore on the restore job properties for SQL will not work. To ensure that Backup Exec has access rights, verify that the account that Backup Exec uses has administrator rights to the computer that is running SQL.

- 7 Select a consistency check to be run after the restore.
- 8 Start the restore job.

After the restore, SQL restarts in multi-user mode.

See [“Restoring data by setting job properties”](#) on page 605.

- 9 Restore the remaining SQL databases.

About redirecting restores for SQL

You can redirect the following:

- A database backup to a different server, database, or instance.
- Differential and log backups to wherever the associated database is restored.
- One or more filegroups in a backup to a different server or instance. Filegroups can be redirected to a different server, but the database file paths cannot be changed. For example, if the filegroup was backed up from G:\SQLDATA, then it must be restored to G:\SQLDATA, even if it is redirected to another server. Filegroups must be restored to the same drive letter and path that they were backed up from.
- A database from a 32-bit or 64-bit platform to any other platform.

Single-job restores and multiple-job restores can both be used in redirected restore operations.

See [“Redirecting restores for SQL”](#) on page 1305.

Redirecting restores for SQL

The following procedure provides details on how to redirect restores for SQL.

See [“About redirecting restores for SQL”](#) on page 1305.

See [“About restoring SQL databases and file groups”](#) on page 1293.

To redirect a restore

- 1 Start a restore job.
 - See [“Restoring from SQL database backups”](#) on page 1295.
 - See [“Restoring from SQL transaction logs up to a point in time”](#) on page 1296.
 - See [“Restoring from SQL transaction logs up to a named transaction”](#) on page 1297.
 - See [“About restoring from SQL filegroup backups”](#) on page 1298.
- 2 After selecting options on the **Restore Job Properties** dialog box, on the **Properties** pane, under **Destination**, click **Microsoft SQL Redirection**.
- 3 Select the appropriate options.
 - See [“Microsoft SQL Redirection options”](#) on page 1306.
- 4 Start the redirected restore job or select other restore options from the **Properties** pane.
 - See [“Restoring data by setting job properties”](#) on page 605.

Microsoft SQL Redirection options

You can set the following options when you restore SQL backup sets to a different server.

See [“Redirecting restores for SQL”](#) on page 1305.

Table J-8 Microsoft SQL Redirection options

Item	Description
Redirect Microsoft SQL Server sets	Lets you enable redirection of SQL backup sets.
Server	<p>Redirects the restore to a different server. After you check the check box, type the target server name.</p> <p>You can redirect a full database backup to a different server or database.</p> <p>If the drive configuration changes after the database backup was created, you must select either of the following options:</p> <ul style="list-style-type: none"> ■ Default drive for restoring database files. ■ Restore all database files to the target instance’s data location. <p>See “Setting restore options for SQL” on page 1289.</p>
Instance	Redirects this restore to a named instance. After checking the checkbox, type the instance name. If you are restoring to the default instance, leave the field empty.
Database	<p>Redirects the restore to a different database on the destination server. After you check the check box, type the destination database name.</p> <p>You can redirect a full database backup to a different server or database.</p> <p>If the drive configuration changes after the database backup was created, you must select either of the following options:</p> <ul style="list-style-type: none"> ■ Default drive for restoring database files. ■ Restore all database files to the target instance’s data location. <p>If you restore a differential or log backup, and the associated database backup was restored to a different server, type the new database name.</p>

Table J-8 Microsoft SQL Redirection options (*continued*)

Item	Description
Use alternate drive	<p>Specifies a default drive to which SQL database files can be restored.</p> <p>When a SQL database is backed up, the physical file names (which include the directory path) of the files that make up the database are stored in the backup set by SQL. For example, for the logical file pubs, the physical file name is stored as E:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\DATA\test.mdf. If the database must later be restored, SQL uses these same physical file names to target the restore to. During a restore, Backup Exec automatically creates any necessary subdirectories that do not exist.</p> <p>However, if the drive where one or more of the database files previously resided no longer exists, Backup Exec moves those files to their original directory path, but on the default drive specified. Using the same example, if the default drive C is specified, then the file with the original directory path of E:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\DATA\test.mdf is restored to C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\DATA\test.mdf.</p> <p>If no default drive is specified in this situation, the job will fail.</p>
Only when original drive does not exist	<p>Lets you use the alternate drive selected in Use alternate drive when the drive from which the database was originally backed up does not exist.</p>
Even when original drive does exist	<p>Restores all database files to their original directory path on the alternate drive selected in Use alternate drive, even if the drive where they originally resided exists.</p> <p>Do not select this option when restoring filegroups. Filegroups must be restored to the same drive letter and path that they were backed up from.</p>

Table J-8 Microsoft SQL Redirection options (*continued*)

Item	Description
Use destination instance's default data directory	<p>Restores files to the default data and log directories of the destination instance. For example, if you are restoring a database to a different instance of SQL, you would select this option to move the database files to the correct location for the new instance.</p> <p>If this option is not selected, then the files are restored to the directory that the master database is in.</p> <p>Do not select this option when restoring filegroups. Filegroups must be restored to the same drive letter and path that they were backed up from.</p>
Use this path	<p>Restores the database to a specific location on disk. To use this option, enter a drive letter and its corresponding path. For example, C:\temp. You can also click the ellipsis button and browse to a disk location. All paths entered are maintained in the Use this path drop-down list, which can be used for future redirected database restore jobs.</p>
Server logon account	<p>Displays a Backup Exec logon account that stores the credentials of a Windows user account when you restore to a server. The Windows user account must have been granted the System Administrator role on the SQL instance. The default logon account is displayed. To use another logon account, click Change.</p> <p>See "How to use Backup Exec logon accounts for SQL resources" on page 1256.</p>
SQL logon account	<p>Displays a Backup Exec logon account that stores the credentials of the SQL user account if you are using SQL Server Authentication. Apply the Backup Exec logon account for the Windows user account to the Windows server that SQL is installed on, and then apply the logon account for the SQL user account to the SQL instance.</p> <p>To use another logon account, click Change. To remove the SQL logon account displayed in this field, click Clear.</p> <p>See "How to use Backup Exec logon accounts for SQL resources" on page 1256.</p>
Retain replication information	<p>Retains the default settings of the database during a redirected database restore job. By default, Backup Exec retains default database settings during database restore jobs, except when database restore jobs are redirected.</p>

Table J-8 Microsoft SQL Redirection options (*continued*)

Item	Description
Check selections	Lets Backup Exec verify your SQL database restore selections. If selection errors are found, Backup Exec notifies you of the error or errors and then attempts to correct them for you.

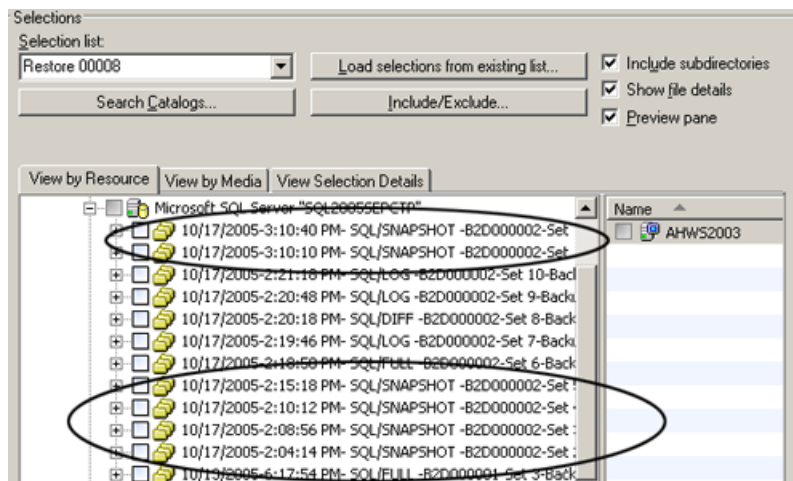
About reverting SQL 2005 or later databases using database snapshots

SQL 2005 or later database snapshots created with Backup Exec can be used to revert a SQL 2005 or later database back to the state it was in at a previous point in time, without having to run a full database restore job.

When viewed by resource in the Restore Job Properties pane, SQL database snapshots appear as backup sets, in chronological order with the most recent snapshot appearing first.

The word SNAPSHOT appears in the backup set description.

Figure J-3 How SQL database snapshots appear in the Restore Job Properties pane



The following caveats apply when reverting a database:

- You cannot undo a SQL 2005 or later database that has been reverted.
- Before reverting a database, Backup Exec deletes all existing database snapshots, including those created with SQL 2005 or later, with the exception

of the snapshot used for the revert. After being deleted, the database snapshots cannot be recovered.

- You cannot redirect a database snapshot restore job.

Reverting a SQL 2005 or later databases using database snapshots

You can revert a SQL 2005 or later database using database snapshots without having to run a full database restore jobs.

See [“Restoring data by setting job properties”](#) on page 605.

See [“About restoring SQL databases and file groups”](#) on page 1293.

To revert a SQL 2005 or later database using database snapshots

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 Select a database snapshot.

After you click Run Now, all previous SQL 2005 or later database snapshots (including those created by SQL 2005 or later) are deleted, and the database is reverted. After the revert completes, the SQL database cannot be returned to its previous state.

- 5 Click **Run Now**.

About disaster recovery of a SQL Server

Backup Exec provides a quicker method for restoring SQL rather than running the Rebuild Master utility or reinstalling SQL to restart SQL. Using Backup Exec, you can replace the corrupted or missing databases with copies of the master and model databases that Backup Exec automatically creates and updates whenever backups of those databases are run. After SQL is running again, you can restore the latest copy of the master database using Backup Exec’s Automate master database restore option, and then restore any other databases, if needed.

If you use the Intelligent Disaster Recovery (IDR) option, then during an IDR recovery of drive C, it will automatically replace the damaged databases with the copies of the master and model databases. You can then restart SQL, and restore the latest master database backup and any other databases that are necessary.

How to prepare for disaster recovery of SQL

To prepare for disaster recovery if you are using SQL do the following:

- Back up both system and user databases and transaction logs regularly. Copies of the master and model databases are automatically created by Backup Exec whenever you back up the master and model databases. Backup Exec places these copies in the same directory that the databases are in, where they must remain in order to be updated.

The following table includes information about MS SQL database locations:

The copies of the master and model databases are named:

- Master\$4idr
 - Mastlog\$4idr
 - Model\$4idr
 - Modellog\$4idr
- Back up the system drives that contain SQL instances. Whenever you back up the system drive that contains a SQL instance, copies of the master and model databases are backed up. Backing up the system drive that SQL is on also backs up all the executables and registry settings needed for SQL to run.
 - Back up the master database whenever any changes are made to SQL.
 - Keep records of any service packs that have been installed.
 - Make sure you are prepared to recover the entire server, not just SQL.

See [“Returning to the last known good configuration”](#) on page 777.

Requirements for SQL disaster recovery

To perform a recovery, you will need the following items:

- The latest backup of the SQL directory (\Program Files\Microsoft SQL Server\MSSQL), and the Windows registry/System State.
- The SQL database or filegroup backups, and differential and log backups.
- An Administrator logon account (or an Administrator equivalent) during the recovery.

Disaster recovery of SQL

You can restore either the entire server, including the SQL databases, from full system backups, or restore only the SQL databases to a newly installed or other available SQL server.

Restoring the entire server, including the SQL databases has the added benefit of recovering other applications and data that may have resided on the server at the time of failure, and can be accomplished using one of the following methods:

- Manual recovery of the Windows server, and then manual recovery of the SQL databases. This method involves manually restoring the Windows server from full system backups, and then recovering the SQL databases.
- The Intelligent Disaster Recovery Option. This option provides an automated method of restoring the Windows server as well as the SQL databases from full system backups.

See [“Microsoft SQL Server recovery notes”](#) on page 1844.

To restore only the SQL databases, review the following:

- To restore only the SQL databases to a newly-installed or other available server, the server must be running on the same hardware platform (cross-platform restores are not supported), and the same version of SQL with the same service pack level as the original server.
- To restore SQL databases to an existing installation of SQL with other active databases, you should redirect the restore.

See [“Redirecting restores for SQL”](#) on page 1305.

See [“About recovering SQL manually”](#) on page 1312.

About recovering SQL manually

When you recover SQL manually, you must first restore the Windows server from full system backups. After recovery of the Windows computer is complete, or after the new server installation is available, you can recover the SQL databases.

See [“About manual disaster recovery of Windows computers”](#) on page 780.

In order to restore SQL databases, SQL must be running; however, SQL cannot be started unless the master and model databases are present.

You can restore the master and model databases and start SQL using one of the following methods:

- Rename the files created by Backup Exec that replace the master and model databases. After the master and model databases are present on SQL, you must start SQL, restore the master database with the Automate master database restore option, and then restore all other databases.

See [“Restarting SQL using database copies”](#) on page 1302.

- Run the Rebuild Master utility (`\Program Files\Microsoft SQL Server\80\Tools\Binn\rebuildm.exe` for SQL 2000).

Note: The Rebuild Master utility is not supported in SQL 2005 or later; see your MS SQL 2005 or later documentation for setup options.

- Reinstall SQL.

This topic only details how to restart SQL by using the copies of the master and model databases made by Backup Exec. For more information on the Rebuild Master utility, or on reinstalling SQL, refer to your MS SQL documentation.

If you are restoring to a new SQL installation, start with the restore of the master database.

See [“To restore the master database”](#) on page 1304.

Symantec Backup Exec Agent for Oracle on Windows or Linux Servers

This appendix includes the following topics:

- [About the Backup Exec Oracle Agent](#)
- [About installing the Oracle Agent](#)
- [Upgrading the Backup Exec Oracle Agent](#)
- [Configuring the Oracle Agent on Windows computers and Linux servers](#)
- [About authentication credentials on the media server](#)
- [About Oracle instance information changes](#)
- [Setting application defaults for Oracle](#)
- [About backing up Oracle resources](#)
- [About restoring and recovering Oracle resources](#)
- [Troubleshooting the Oracle Agent](#)

About the Backup Exec Oracle Agent

The Symantec Backup Exec Agent for Oracle on Windows or Linux Servers (Oracle Agent) uses Oracle's Recovery Manager (RMAN) to protect Oracle databases. RMAN is a tool that manages the backup and restore and recovery of Oracle databases.

The following features are available with the Oracle Agent:

- The ability to initiate backup and restore operations from Backup Exec or from the RMAN console as a Database Administrator (DBA).
Operations that the DBA performs on the RMAN console are referred to as DBA-initiated operations. You should refer to your Oracle documentation for information about RMAN.
- Multiple data stream support for increased performance during backup and restore operations.
- RMAN recovery catalog support to manage the backup, restore, and recovery of Oracle databases.
- Oracle Real Application Cluster (RAC) support.

The following are not supported:

- Tivoli Storage Manager (TSM) devices as storage devices for Oracle backup jobs.
- The Oracle Management Server.
- Oracle backup and restore jobs running on the IPv6 protocol.

See [“About installing the Oracle Agent”](#) on page 1316.

See [“Upgrading the Backup Exec Oracle Agent”](#) on page 1317.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1318.

About installing the Oracle Agent

The Oracle Agent is installed as a separate, add-on component of Backup Exec.

To protect local or remote Oracle instances, you must install the following Backup Exec options:

- Backup Exec Remote Agent for Windows Systems on remote Windows computers.

Note: If you upgrade a previous version of the Remote Agent on an Oracle server, you must restart the Oracle server after the upgrade. Backup Exec jobs cannot complete successfully until you restart the Oracle server.

See [“About installing the Remote Agent for Windows Systems”](#) on page 139.

- Backup Exec Remote Agent for Linux and UNIX Servers on remote Linux computers.

See [“About installing the Remote Agent for Linux or UNIX Servers”](#) on page 1877.

- Backup Exec Oracle Agent on the media server.
See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

Upgrading the Backup Exec Oracle Agent

The Backup Exec Oracle RMAN Agent replaces the legacy GRFS Oracle Agent. All existing Oracle jobs are upgraded for use with the new agent. When you upgrade to the Backup Exec Oracle RMAN Agent, backup jobs for Oracle instances that were created with the legacy Oracle Agent are placed on hold. You must do the following:

Table K-1 Upgrading the Backup Exec Oracle Agent

Step	Action
Step 1	Verify that the logon account for each resource that was backed up by the previous Oracle Agent is valid for the new Oracle Agent. See “Changing and testing resource credentials for restore jobs” on page 630.
Step 2	Use the Remote Agent Utility to configure information about the Oracle instances for the Oracle Agent, and to enable media server access to the Oracle databases. See “Configuring the Oracle Agent on Windows computers and Linux servers” on page 1318.
Step 3	Add the fully qualified domain name and logon account of the Oracle server to the media server’s list of authentication credentials. See “Setting authentication credentials on the media server for Oracle operations” on page 1330.
Step 4	Take the associated jobs off hold. See “Removing the hold on the job queue” on page 570.

After upgrading, the database control file resource no longer appears in the backup selections tree under the Oracle server node. Backup Exec automatically backs up the database control file whenever a tablespace or other resource on the Oracle server is backed up.

Note: If you have a backup selections list that contains the database control file as its single resource, create another selection list that contains other resources. The database control file is then automatically backed up whenever the other resources in the selection list are backed up. This applies only to selection lists created with the legacy GRFS Oracle Agent.

Configuring the Oracle Agent on Windows computers and Linux servers

Before you can back up or restore Oracle databases, you must do the following:

Table K-2 Configuring the Oracle Agent on Windows computers and Linux servers

Step	Action
Step 1	Configure information about the Oracle instances for the Oracle Agent. See “Configuring an Oracle instance on Windows computers” on page 1319. See “Configuring an Oracle instance on Linux servers” on page 1324.

Table K-2 Configuring the Oracle Agent on Windows computers and Linux servers (*continued*)

Step	Action
Step 2	<p>Enable database access for the media server.</p> <p>Whenever Oracle instance information changes or a new configuration is added, you must update the Remote Agent Utility. If credential information is not updated, is incorrect, or the server is down, the error "Unable to attach to a resource..." may appear when you run a backup job. If this message appears, you must bring the server online and configure the information.</p> <p>For Oracle RAC, run the Remote Agent Utility on each node and add information about the instances. When Oracle RAC nodes are added or removed, you must enter information about any changes to instances in the Remote Agent Utility.</p> <p>Note: When you use the Remote Agent Utility, the user account with which you are logged on should be a member of the Oracle DBA group.</p> <p>You must have administrator privileges to run the Remote Agent Utility.</p> <p>See "Enabling database access for Oracle operations on Windows computers" on page 1323.</p> <p>See "Enabling database access for Oracle operations on Linux servers" on page 1327.</p>
Step 3	<p>Set authentication credentials for Oracle.</p> <p>See "Setting authentication credentials on the media server for Oracle operations" on page 1330.</p>

Configuring an Oracle instance on Windows computers

You can use the Remote Agent Utility to configure information about the Oracle instances for the Oracle Agent on Windows computers.

To configure an Oracle instance on Windows computers

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

When the Remote Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 On the **Oracle** tab, click **New**.

Any instances that currently exist on the computer appear on the tab.

- 3 Complete the appropriate options.

See [“Oracle Agent Configuration options”](#) on page 1320.

- 4 Click **OK**.

Oracle Agent Configuration options

You can set the following Oracle Agent Configuration options.

See [“Configuring an Oracle instance on Windows computers”](#) on page 1319.

See [“Editing an Oracle instance on Windows computers”](#) on page 1322.

Table K-3 Oracle Agent Configuration options

Item	Description
Local instance name	<p>Displays the name of the Oracle instance. If you edit an instance, you cannot change the instance name.</p> <p>For Oracle RAC nodes, enter the fully qualified domain name of each physical node.</p> <p>The fully qualified domain name of the node appears on the media server in the Backup Selections tree, under Oracle Real Application Clusters.</p> <p>The name is in the format RAC-<dbname>-<dbid>, where dbname is the database name, and dbid is the database ID.</p>
User name	<p>Displays the user name for the Oracle instance.</p> <p>If the credentials for the Oracle instance change, you must enter a user with SYSDBA rights to the Oracle instance.</p> <p>For Oracle RAC nodes, enter the same set of credentials for all of the nodes.</p>
Password	<p>Displays the password for the Oracle instance user name.</p>

Table K-3 Oracle Agent Configuration options (*continued*)

Item	Description
Confirm password	Displays the password again to confirm it.
Use recovery catalog	Indicates that you plan to use the Oracle recovery catalog. The Oracle Agent supports the use of the RMAN recovery catalog to manage the backup, restore, and recovery of Oracle databases. If you choose not to use a recovery catalog, RMAN uses the target database control file as the sole repository of metadata.
TNS name	Displays the Oracle Net Service name.
User name	Displays the user name for the Oracle recovery catalog.
Password	Displays the password for the Oracle recovery catalog.
Confirm password	Displays the password for the recovery catalog again to confirm it.
Media server name or IP address	Displays the name or IP address of the Backup Exec media server where you want to send the DBA-initiated backup jobs. You must use the same form of name resolution for all operations. For example, if you use the IP address of this computer for backup operations, you must also use the IP address for restore operations. If you use the full computer name for backup operations, you must also use the full computer name for restore operations.
Job template name	Displays the name of the Backup Exec job template that you want the DBA-initiated job to use for backup and restore operations. You create the job template on the DBA-initiated Job Settings dialog box on the Backup Exec media server. If you do not specify a job template, the default job template is used. See “Creating a template for DBA-initiated jobs” on page 424.

Viewing an Oracle instance on Windows computers

You can use the Remote Agent Utility to view information about the Oracle instances for the Oracle Agent on Windows servers.

To view an Oracle instance on Windows computers

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.
- 2 On the **Oracle** tab, view the instances that currently exist on the computer. See [“Oracle options for the Remote Agent Utility”](#) on page 1322.
- 3 Click **OK**.

Oracle options for the Remote Agent Utility

You can set the following Oracle options for the Remote Agent Utility.

See [“Viewing an Oracle instance on Windows computers”](#) on page 1321.

Table K-4 Oracle options for the Remote Agent Utility

Item	Description
Instance	Displays the name of the Oracle instance.
User Name	Displays the user name for the Oracle instance.
Recovery Catalog	Displays the name of the recovery catalog.
Media Server	Displays the name or IP address of the Backup Exec media server where you want to send the DBA-initiated backup jobs.
Job Template	Displays the name of the DBA-initiated template. See “About performing a DBA-initiated backup job for Oracle” on page 1340.
New	Lets you add an Oracle instance.
Edit	Lets you revise an Oracle instance.
Delete	Lets you remove an Oracle instance.

Editing an Oracle instance on Windows computers

You can use the Remote Agent Utility to revise information about the Oracle instances for the Oracle Agent on Windows computers.

To edit an Oracle instance on Windows computers

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

When the Remote Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 On the **Oracle** tab, click **Edit**.

Any instances that currently exist on the computer appear on the tab.

- 3 Edit the appropriate options.

See [“Oracle Agent Configuration options”](#) on page 1320.

- 4 Click **OK**.

Deleting an Oracle instance on Windows computers

You can use the Remote Agent Utility to remove an Oracle instance for the Oracle Agent on Windows computers.

To delete an Oracle instance on Windows computers

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

When the Remote Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 On the **Oracle** tab, click **Delete**.

Any instances that currently exist on the computer appear on the tab.

- 3 Click **OK**.

Enabling database access for Oracle operations on Windows computers

You can use the Remote Agent Utility to enable database access for the Windows media server after you configure an Oracle instance.

See [“About backing up Oracle resources”](#) on page 1335.

See [“About backing up Oracle RAC resources”](#) on page 1337.

See [“Creating a template for DBA-initiated jobs”](#) on page 424.

See [“Setting application defaults for Oracle”](#) on page 1334.

To enable database access for Oracle operations on Windows computers

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

When the Remote Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 On the **Database Access** tab, complete the appropriate options.
See “[Database access options for the Remote Agent Utility](#)” on page 1956.
- 3 Click **OK**.
- 4 For Oracle RAC installations, type the fully qualified domain name that you want to publish to.

The media server that you publish to lists the RAC databases in its backup selection tree, under the node named Oracle Real Application Clusters.

If you do not enter a fully qualified domain name to publish to, the RAC databases are not listed in the media server’s backup selection tree.

See “[About publishing the Remote Agent for Windows Systems to media servers](#)” on page 1951.

- 5 On the media server, add the name of the Oracle server and the user name that you enabled for database access to the media server’s list of authentication credentials.

See “[About authentication credentials on the media server](#)” on page 1329.

Configuring an Oracle instance on Linux servers

You can use the Remote Agent Utility to configure information about the Oracle instances for the Oracle Agent on Linux servers.

To configure an Oracle instance on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.
- 2 Change to the following directory:
cd /opt/VRTSralus/bin
- 3 Start the Remote Agent Utility:
./AgentConfig
- 4 Type **2** to select Configure Oracle instance information, and then press **Enter**.
- 5 Type **1** to select the Add a new Oracle Instance option, and then press **Enter**.

- 6** Enter the name of the Oracle instance in upper case characters.
For example, ORACLENAME.
- 7** Enter the user name for the Oracle instance.
If the credentials for the Oracle instance are changed, you must update the credentials in this field. For Oracle RAC nodes, enter the same set of credentials for all of the nodes.

When you use the Remote Agent Utility to enter the Oracle credentials for an instance, the credentials cannot be verified if the user account with which you are logged on is a member of the Oracle DBA group. If the credentials are incorrect, the error "Unable to attach to a resource..." may appear when you run a backup job.
- 8** To display the Oracle database in a media server's backup selection list under Favorite Resources, type the media server name or IP address to which you want the remote computer to publish to.

The media server lists the Oracle database under Favorite Resources, under Linux/Unix, under <computer name> <root>.

Oracle RAC databases are listed in the media server's backup selection list, under Oracle Real Application Clusters. They are not listed under Favorite Resources.
- 9** When prompted, specify if you want to use a recovery catalog.

The Oracle Agent supports the use of the RMAN recovery catalog to manage the backup, restore, and recovery of Oracle databases. If you choose not to use a recovery catalog, RMAN uses the target database control file as the sole repository of metadata.

If you specify a recovery catalog, any database that you want to back up must be registered in the recovery catalog before you can run backup jobs from the media server.
- 10** To use a recovery catalog, type the recovery catalog name and a user name and password for the recovery catalog.

- 11 To use a customized DBA-initiated job settings template, type the name of the template.

See [“Creating a template for DBA-initiated jobs”](#) on page 424.

- 12 Do one of the following:

To commit the new entry to the configuration file Type **Y**, and then press **Enter**.

To cancel this entry Type **N**, and then press **Enter**.

Viewing an Oracle instance on Linux servers

You can use the Remote Agent Utility to view information about the Oracle instances for the Oracle Agent on Linux servers.

The following information is listed:

- Name of the instance
- Logon name for the instance
- IP address of the default media server name for DBA-initiated operations
- Name of the DBA-initiated job template

To view an Oracle instance on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.
- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```
- 3 Start the Remote Agent Utility:

```
./AgentConfig
```
- 4 Type 4.

Editing an Oracle instance on Linux servers

You can use the Remote Agent Utility to revise information about the Oracle instances for the Oracle Agent on Linux servers.

To edit an Oracle instance on Linux computers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.
- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```
- 3 Start the Remote Agent Utility:

```
./AgentConfig
```
- 4 Type **2** to select Configure Oracle Instance Information, and then press **Enter**. Any instances that currently exist on the computer are discovered.
- 5 Type **2**.
- 6 Follow the prompts.

Deleting an Oracle instance on Linux servers

You can use the Remote Agent Utility to remove an Oracle instance for the Oracle Agent on Linux servers.

To delete an Oracle instance for the Oracle Agent on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.
- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```
- 3 Start the Remote Agent Utility:

```
./AgentConfig
```
- 4 Type **2** to select Configure Oracle Instance Information, and then press **Enter**. Any instances that currently exist on the computer are discovered.
- 5 Type **3**.
- 6 Follow the prompts.

Enabling database access for Oracle operations on Linux servers

You can use the Remote Agent Utility to enable database access for the Linux server after you configure an Oracle instance.

See [“Setting authentication credentials on the media server for Oracle operations”](#) on page 1330.

See [“About backing up Oracle resources”](#) on page 1335.

See [“About backing up Oracle RAC resources”](#) on page 1337.

See [“Creating a template for DBA-initiated jobs”](#) on page 424.

See [“Setting application defaults for Oracle”](#) on page 1334.

To enable database access for Oracle operations on Linux servers

1 On the Linux server on which the Oracle instances are installed, open a Terminal window.

2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```

3 Start the Remote Agent Utility:

```
./AgentConfig
```

4 Type **1** to select Configure database access, and then press **Enter**.

5 Type the user name that is in the beoper group on the Linux system.

See [“About the Backup Exec operators group for the Remote Agent for Linux or UNIX Servers”](#) on page 1880.

If the authentication fails when the Oracle resources are backed up, the backup job fails. If the authentication fails when you browse the backup sets for a restore job, then the backup sets become unavailable, and you must run a DBA-initiated restore job to restore data.

6 Type the password for this logon account, and then confirm it.

The logon credentials are not stored on this computer.

- 7 When prompted, specify if you want to use a custom port to connect to the media server communications between this computer and the media server during Oracle operations.

Port 5633 is used by default. If you change the port number on this computer, you must also change it on the media server, and then restart the Backup Exec Job Engine Service on the media server. If a Windows firewall is enabled, you must add this port as an exception.

See [“Setting default backup network and security options”](#) on page 404.

- 8 Do one of the following:

To commit the Oracle operation settings to the configuration file Type **Y**, and then press **Enter**.

To cancel this entry Type **N**, and then press **Enter**.

About authentication credentials on the media server

You must add the Oracle fully qualified domain name and the logon account name to the media server’s list of Oracle servers and authentication credentials. The media server has database access for operations on Oracle instances that are included in the authentication list. Before you start any backup or restore operations, on the computer on which the Oracle instances are installed, make sure that you use the Remote Agent Utility to configure instance information and database access.

The logon account name must have administrative rights or backup operator rights to the Oracle server. If the user name is incorrect or is not provided, or if it does not have the appropriate rights, then you cannot perform Oracle backup or restore operations to that computer.

Note: For Oracle RAC nodes, enter the fully qualified domain name for the logon account name. You can view the fully qualified domain name of the node in the backup selections list. It is in the form RAC-`<database name>`-`<database ID>`.

See [“Setting authentication credentials on the media server for Oracle operations”](#) on page 1330.

See [“Editing authentication credentials on the media server for Oracle operations”](#) on page 1332.

See [“Deleting an Oracle server from the media server’s list of authentication credentials”](#) on page 1333.

Setting authentication credentials on the media server for Oracle operations

You must add the Oracle server to the list so that the media server has database access for operations.

See [“About authentication credentials on the media server”](#) on page 1329.

See [“Authentication Credentials options”](#) on page 1331.

See [“About Oracle instance information changes ”](#) on page 1333.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1318.

See [“Clustering Backup Exec using Veritas Cluster Server”](#) on page 845.

To set authentication credentials on the media server for Oracle operations

- 1 On the media server, on **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Oracle**.
- 3 Click the **Modify list** button.

On the **Authentication Credentials for Oracle and DB2 Servers** dialog box, you can add, edit, or delete a fully qualified domain name and a logon account.

- 4 Click **New**.
- 5 Enter the name of the Oracle server on which the instance is installed.

The name of the Oracle server should match the name of the server that lists the Oracle resource. Symantec recommends that you enter the fully qualified domain name. For example, Servername.domain.com is the fully qualified domain name and Servername is the NETBIOS name. For Oracle RAC nodes, enter the fully qualified domain name for the each node in the list.

6 To add the logon account name, do one of the following:

- Click the arrow Select the logon account name that you want to add.
- Click New On the Logon Account Selection dialog box, click **New**.
See [“Creating a Backup Exec logon account”](#) on page 187.

Use the same logon account format that you use when you enter the logon account name on the Database Access tab in the Remote Agent Utility. For example, if you entered Domainname\Username on the Remote Agent Utility, use that same format on the list of authentication credentials.

7 Click **OK**.**8** On the **Authentication Credentials for Oracle and DB2 Servers** dialog box, click **OK**.

Authentication Credentials options

You can set the following Authentication Credentials for a server.

See [“About authentication credentials on the media server”](#) on page 1329.

Table K-5 Authentication credentials for Oracle and DB2 servers options

Item	Description
Server	Displays the name of the Oracle and DB2 media server.
Logon Account	Displays the name of the logon account that has rights to the Oracle or DB2 server.
New	Lets you add the fully qualified domain name and logon account credentials to the list.
Edit	Lets you revise the fully qualified domain name and logon account credentials.
Delete	Lets you remove the fully qualified domain name and logon account credentials.

Add or edit server options

The following options are available when you add or edit a computer name and logon account name.

See [“About authentication credentials on the media server”](#) on page 1329.

Table K-6 Add or edit server options

Item	Description
Oracle or DB2 Server	Specifies the name of the Oracle or DB2 media server.
Logon Account	Specifies the name of the logon account for the Oracle or DB2 server.
New	Lets you add a logon account that has rights to the Oracle or DB2 server.

Editing authentication credentials on the media server for Oracle operations

If the Oracle server name or the logon account name for the Oracle server changes, you must update the media server's list of Oracle servers and authentication credentials. Make the same changes on the Oracle server by using the Remote Agent Utility to configure instance information and database access.

The logon account name must have administrative rights or backup operator rights to the Oracle server. If the user name is incorrect or is not provided, or if it does not have the appropriate rights, then you cannot perform Oracle backup or restore operations to that computer.

See [“About authentication credentials on the media server”](#) on page 1329.

See [“Authentication Credentials options”](#) on page 1331.

See [“Setting authentication credentials on the media server for Oracle operations”](#) on page 1330.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1318.

See [“Deleting an Oracle server from the media server's list of authentication credentials”](#) on page 1333.

To edit authentication credentials on the media server for Oracle operations

- 1 On the media server, on **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Oracle**.
- 3 Click **Modify list**.

On the **Authentication Credentials for Oracle and DB2 Servers** dialog box, you can add, edit, or delete a server name and a logon account.

- 4 Select the item that contains the server name or logon account that you want to edit.
- 5 Click **Edit**.
- 6 Change the server name or change the logon account name.
See [“Editing a Backup Exec logon account”](#) on page 190.
- 7 Click **OK**.
- 8 On the **Authentication Credentials for Oracle and DB2 Servers** dialog box, click **OK**.

Deleting an Oracle server from the media server’s list of authentication credentials

You can delete an Oracle server name or logon account from a media server’s list of authentication credentials.

To delete an Oracle server from the media server’s list of authentication credentials

- 1 On the media server, on Tools menu, click **Options**.
- 2 On the properties pane, under Job Defaults, click **Oracle**.
- 3 Click the **Modify list** button.
On the Authentication Credentials for Oracle and DB2 Servers dialog box, you can add, edit, or delete a server name and a logon account.
- 4 Select the item that contains the server name or logon account that you want to delete.
- 5 Click **Delete**.
See [“Deleting a Backup Exec logon account”](#) on page 192.
- 6 Click **OK**.

About Oracle instance information changes

Whenever information about the Oracle instance changes, such as the instance user name and password, you must update the Remote Agent Utility.

When Oracle RAC nodes are added or removed, you must enter information about any changes to instances in the Remote Agent Utility. After these changes are entered, the Backup Exec media server discovers them.

If the changes are not updated in the Remote Agent Utility, the error "Unable to attach to a resource..." may appear when you run a backup job.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1318.

Setting application defaults for Oracle

You can use the defaults that Backup Exec sets during installation for all Oracle backup jobs, or you can choose your own defaults.

See [“About Oracle instance information changes”](#) on page 1333.

See [“About backing up Oracle resources”](#) on page 1335.

See [“Troubleshooting the Oracle Agent”](#) on page 1353.

To set application defaults for Oracle

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Oracle**.
- 3 Complete the appropriate options.
See [“Oracle default options”](#) on page 1334.
- 4 Click **OK**.

Oracle default options

You can set options for Oracle backup jobs and the list of servers and authentication credentials.

See [“Setting application defaults for Oracle”](#) on page 1334.

Table K-7 Oracle default options

Item	Description
Backup Method	<p>Specifies one of the following backup methods:</p> <ul style="list-style-type: none"> ■ Full - Back up selections. This method is the equivalent of the Oracle RMAN Incremental: Level 0 backup. Select this method to perform a full backup of Oracle selections. ■ Differential - Back up changes since last full. This method is the equivalent of the Oracle RMAN Cumulative Incremental: Level 1 backup. Select this method to back up all database changes since the last full backup. A full backup is performed of all archived redo logs since log files are never partially backed up. ■ Incremental - Back up changes since last full or incremental. This method is the equivalent of the Oracle RMAN Incremental: Level 1 backup. Select this method to back up all database changes since the last full or incremental backup. A full backup is performed of all archived redo logs since log files are never partially backed up.
Delete backed up archive log files	Enables Backup Exec to delete the archived log files automatically after the backup job.
Do not back up archived log files that have already been backed up	Enables Backup Exec to skip any archived log files that have been backed up previously.
Perform the backup offline	Lets you take the database offline before you start the backup job. Backup Exec brings the database online after the backup job is complete.
Modify list	<p>Lets you add, edit, or delete the Oracle computer name and the logon account name to the media server's list of authentication credentials for Oracle servers.</p> <p>See "Setting authentication credentials on the media server for Oracle operations" on page 1330.</p>

About backing up Oracle resources

Before you back up Oracle resources, review the following:

- You must run the Remote Agent Utility on the Oracle server and add information about the instances before you can perform any backup or restore operations.

When Oracle instance information changes, you must update the Remote Agent Utility. After these changes are entered, the Backup Exec media server discovers them.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1318.

- During a backup operation, the amount of data that is backed up may not equal the size of the Oracle files that are on the disk. This behavior is normal. Backup Exec backs up the selected data files as well as a copy of the control file.
- In a Central Admin Server Option environment, all backup jobs for a specific Oracle instance must be delegated to the same managed media server. If you do not restrict the backup job to the same managed media server, then before you can restore data, you must move the physical media that contains the backup sets to a single managed media server.
See [“Restricting the backup of a selection list to specific devices in CASO”](#) on page 1546.

- If the Oracle database resides on volumes that are configured with Oracle Automatic Storage Management (ASM), you cannot select these volumes as part of a file system backup.

The following message appears when you attempt to select the volumes:

```
An error was encountered while attempting to browse the  
contents of <drive>. A device-specific error occurred.
```

- The database must be in a mounted or open state before you can make backup selections.
- The database must be in ARCHIVELOG mode before an archive log can be displayed in the backup selections list.
- The Backup Exec option to display progress indicators for backup jobs is not available for backup jobs when Oracle resources are included in the backup selection list.

See [“Backing up Oracle resources”](#) on page 1337.

See [“About performing a DBA-initiated backup job for Oracle”](#) on page 1340.

See [“Creating a template for DBA-initiated jobs”](#) on page 424.

About backing up Oracle RAC resources

Oracle Real Application Cluster (RAC) is an active-active cluster with shared storage, in which multiple instances share a single physical database. Since all of the participating nodes can access the database, you can initiate backup, restore, or recovery from any node. Oracle RAC databases appear in the media server's backup selection list, under the node Oracle Real Application Clusters.

Requirements for backing up Oracle RAC resources include the following items:

- You must run the Remote Agent Utility on each node and add information about the instances before you can perform any backup or restore operations. When RAC nodes are added or removed, you must update the Remote Agent Utility with information about the affected instances. After these changes are entered, the Backup Exec media server discovers them.
See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1318.
- You must select the RAC fully qualified domain name when making backup selections.
Each node in the cluster uses the same fully qualified domain name. The fully qualified domain name of the node appears under the Oracle Real Application Clusters resource in the media server's backup selections list. It is in the form RAC-<database name>-<database ID>.

Backing up Oracle RAC is similar to backing up standard Oracle databases.

You should be aware of the following differences:

- By default, each node in an Oracle RAC stores its archive logs locally. To have a meaningful backup of the archive logs, back up each archive log. Alternatively, you can move the archive logs to a shared device for backup.
- Each node that is part of the cluster is assigned a priority. For database backups, Backup Exec connects to the node that has the highest priority. Backup Exec uses the fully qualified domain name to connect to the node.

See [“About performing a DBA-initiated backup job for Oracle”](#) on page 1340.

See [“Creating a template for DBA-initiated jobs”](#) on page 424.

See [“Setting application defaults for Oracle”](#) on page 1334.

Backing up Oracle resources

Before you back up Oracle resources, make sure that you have completed all of the installation and configuration requirements.

Note: The password for the credentials that you use to connect to the Oracle resource cannot contain special characters.

See [“About installing the Oracle Agent”](#) on page 1316.

See [“About backing up Oracle resources”](#) on page 1335.

See [“Setting application defaults for Oracle”](#) on page 1334.

See [“Troubleshooting the Oracle Agent”](#) on page 1353.

See [“Creating a template for DBA-initiated jobs”](#) on page 424.

To back up Oracle resources

- 1** On the navigation bar, click the arrow next to Backup.
- 2** Click **New Backup Job**.
- 3** In the backup selections list, do one of the following:

For Oracle RAC Expand the RAC fully qualified domain name under the Oracle Real Application Clusters node.

Each node in the cluster uses the same fully qualified domain name. It is in the form RAC-*database name*-*database ID*.

For Oracle resources Expand the appropriate node under Favorite Resources.

The current state of the database appears in brackets next to the database name. If the database is down, you cannot select it for backup.

- 4** To select data for backup, select the check box next to the item that you want to back up.

If you select a container item for backup, you cannot exclude individual items in that container. You must clear the check box for the container item, and then select the individual items in the container that you want to include in the backup.

- 5** On the **Properties** pane, under **Settings**, click **Oracle**.

- 6** Complete the appropriate options.

See [“Oracle backup options”](#) on page 1339.

- 7** To configure multiple data streams for backup, under Destination, click **Device and Media**.

- 8 Complete the appropriate options.
See [“Oracle device and media options for backup jobs”](#) on page 1340.
- 9 Complete the remaining backup job properties as necessary.
See [“How to back up data”](#) on page 329.

Oracle backup options

You can set the following options when you create a backup job for Oracle.
See [“Backing up Oracle resources”](#) on page 1337.

Table K-8 Oracle backup options

Item	Description
Backup method	<p>Specifies one of the following backup methods:</p> <ul style="list-style-type: none"> ■ Full - Back up selections. This method is the equivalent of the Oracle RMAN Incremental: Level 0 backup. Select this method to perform a full backup of Oracle selections. ■ Differential - Back up changes since last full. This method is the equivalent of the Oracle RMAN Cumulative Incremental: Level 1 backup. Select this method to back up all database changes since the last full backup. A full backup is performed of all archived redo logs since log files are never partially backed up. ■ Incremental - Back up changes since last full or incremental. This method is the equivalent of the Oracle RMAN Incremental: Level 1 backup. Select this method to back up all database changes since the last full or incremental backup. A full backup is performed of all archived redo logs since log files are never partially backed up.
Delete backed up archive log files	Lets you delete the archived log files automatically after the backup.
Do not back up archived logfiles that have already been backed up	Enables Backup Exec to skip any archived logfiles that have been backed up previously.
Perform the backup offline	Enables Backup Exec to take the database offline before you start the backup job. Backup Exec brings the database online after the backup job is complete.

Oracle device and media options for backup jobs

You can set the following device and media options when you create a backup job for Oracle.

See [“Backing up Oracle resources”](#) on page 1337.

Table K-9 Oracle device and media options

Item	Description
Maximum number of devices to use for resources that support multiple data streams	<p>Specifies the maximum number of devices that the backup job can use.</p> <p>If you specify more than one device, you must choose one of the following items as a destination device for the backup job:</p> <ul style="list-style-type: none"> ■ A device pool. ■ A backup-to-disk folder that has at least two concurrent operations enabled. <p>If there is only one device for the backup job to use, then the data streams from RMAN are backed up serially to the media.</p> <p>See “Creating a backup-to-disk folder by setting properties” on page 499.</p> <p>This feature is not available for DBA-initiated jobs.</p>
Minimum number of devices, terminate job if fewer devices are available	<p>Specifies the minimum number of devices that the job can use.</p> <p>If the job cannot acquire the minimum number of devices, the job fails.</p> <p>This feature is not available for DBA-initiated jobs.</p>

About performing a DBA-initiated backup job for Oracle

A Database Administrator (DBA) can initiate a backup or restore operation for Oracle from the RMAN console. Example scripts for backup and restore operations that you can run from the RMAN console are installed to the following location:

```
\Program Files\Symantec\Backup Exec\scripts\Oracle
```

Refer to your Oracle documentation for more information on using the RMAN console.

Review the following notes before initiating backup jobs for Oracle from the RMAN console:

- Make sure that you have completed all of the preparations for configuring the Oracle Agent.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1318.

- The channel is not released if the RMAN console is not exited, or if a new manual channel is not allocated on that console.
See [“Troubleshooting the Oracle Agent”](#) on page 1353.
- The SKIP INACCESSIBLE option is available in RMAN to skip corrupt data and log files. Jobs that include this option may complete successfully, but it is likely that if this data is restored, the database will be in an inoperable state. The SKIP INACCESSIBLE option is not available for media server operations. If a backup job encounters corrupt data or log files, the job fails. Symantec recommends that you do not use this option.
- In a CASO environment, the destination device that you select in the DBA-initiated job template must be locally attached to the central administration server.
If the destination device includes a device pool, all devices in the pool must be locally attached to the central administration server.

See [“About Oracle instance information changes”](#) on page 1333.

See [“About backing up Oracle resources”](#) on page 1335.

See [“About configuring DBA-initiated job settings”](#) on page 423.

About restoring and recovering Oracle resources

The restore selections that you choose in Backup Exec are converted to a script. RMAN uses the script to determine what to restore from the Backup Exec media. After the data has been restored to the Oracle server, RMAN completes any requested recovery and restore operations. These recovery and restore operations are determined by the options that you select.

Note: Symantec no longer supports the Symantec Backup Exec - Agent for Oracle Server and its use of GRFS technology. Files that are backed up using that agent can be restored as a file system restore.

See [“Restoring from a legacy GRFS Oracle Agent database backup”](#) on page 1348.

Some recovery operations may not require media from the media server. For example, the redo logs may still be on the Oracle server. During a restore operation, the amount of data that is restored may not be equal to the amount of data that is backed up. In some cases, the amount of data that is restored is listed as 0 bytes. This behavior is normal because Oracle might skip data files that are already up-to-date on the disk.

If you perform a complete recovery on the whole database, or on a tablespace or datafile, you must restore a backup of the database or files that you want to recover. Then you must apply online or archived redo logs, or both. For jobs that are initiated both from the media server and from a DBA, RMAN determines the specific data that it requires from Backup Exec to complete the restore and the recovery that you request.

Note: Backup Exec does not support Oracle tablespace point-in-time restore (TSPITR) through server-initiated operations.

You can only choose Oracle restore selections from the View by Resource tab on the Restore Job Properties dialog box. The View by Media tab displays backup sets, but you cannot browse or select the contents.

On the View by Resource tab, you can make restore selections from the online database or from control files.

Table K-10 Restore selections for Oracle resources

View restore data in	Description
Online database	<p>Provides a view of the live database (if available). You can select an entire database or select individual tablespaces and datafiles.</p> <p>Note: For Oracle RAC, the Oracle database is listed under its fully qualified domain name. It is in the form RAC-<i><database name></i>-<i><database ID></i>.</p>
Control files	<p>Provides a list of all backed up control files. Each control file lists the date it was backed up and the control file's piece ID.</p> <p>You cannot select individual tablespaces or datafiles for restore.</p> <p>Caution: When you recover to a point in time by using a control file, make sure that the date of the control file backup is before the specified recovery point in time. There should not be any database structure changes between the two times. Additionally, when you restore a control file, the entire database reverts to the point in time of the restored control file.</p>

See [“Restoring Oracle data”](#) on page 1343.

See [“Redirecting a restore of Oracle data”](#) on page 1347.

About DBA-initiated restore and recovery for Oracle

DBAs can initiate restore jobs directly from the RMAN console. For example, you can specify the resources you want restored, and the number of channels to allocate for the restore job. Refer to your Oracle documentation for more information on using the RMAN console.

All DBA-initiated restore jobs are deleted after the jobs have completed.

Note: If you attempt to use a DBA-initiated restore job to restore a datafile, a tablespace, or a database that is online, a message appears on the RMAN console. The message indicates that the restore cannot be performed because Oracle does not allow the restore of these items if they are online. However, this message is not reported to Backup Exec. Therefore, the DBA-initiated restore job is reported in Backup Exec as completing successfully.

Restoring Oracle data

Before you restore Oracle resources, make sure that you have completed all of the configuration requirements.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1318.

See [“About restoring and recovering Oracle resources”](#) on page 1341.

Note: In a CASO environment, you can delegate an Oracle restore job to a managed media server. However, if the restore job uses encrypted Oracle backup sets from which to restore, the restore job may fail. An error message may appear that indicates the managed media server does not have the required encryption keys necessary to complete the job. You must create the encryption keys on the managed media server that runs the restore job.

See [“Creating an encryption key”](#) on page 420.

See [“Creating a template for DBA-initiated jobs”](#) on page 424.

See [“About DBA-initiated restore and recovery for Oracle”](#) on page 1343.

See [“Troubleshooting the Oracle Agent”](#) on page 1353.

To restore Oracle data

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.

- 4 Click the **View by Resource** tab.
- 5 Expand the **All Resources** icon.
- 6 Expand the system resource that contains the database instance that you want to restore.
- 7 Expand the database instance that you want to restore.

The current state of the database is listed to the right of the database name. The database must be in a Mounted, Nomounted, or Open state. You cannot select databases for restore jobs if they have a status of Down.

- 8 Expand either the **Current Database** or **Control Files** resource.

If you restore a control file, the entire database reverts to the point in time of the selected control file backup. You cannot use this option to restore an individual tablespace or datafile.

To restore an individual tablespace or datafile, make your selection from the online database view instead.

- 9 Select the appropriate items to restore.
- 10 On the **Properties** pane, under **Settings**, click **Oracle**.
- 11 Complete the appropriate options.

See [“Oracle restore options”](#) on page 1344.

- 12 Select other restore options from the **Properties** pane as appropriate, and then start the restore job.

See [“Restoring data by setting job properties”](#) on page 605.

- 13 Run a full backup of the restored database.

Oracle restore options

You can set the following options when you create a restore job for Oracle.

See [“Restoring Oracle data”](#) on page 1343.

Table K-11 Oracle restore options

Item	Description
Restore from full and/or incremental backups	<p>Specifies a restore method.</p> <p>Select an option to restore data from one of the following:</p> <ul style="list-style-type: none"> ■ The most recent available backups ■ To a point in time ■ To a specific system control number (SCN) <p>RMAN determines which backup objects are necessary for this restore job, and then Backup Exec restores those objects.</p> <p>You may also need to check a recovery option for recovering from redo logs to complete the restore.</p>
To the most recent available	Restores the Oracle database to the most recent full and incremental backups that are available.
To a point in time	<p>Restores data up to and including a point in time. After the point in time, recovery stops.</p> <p>In the Date box, select the part of the date that you want to change, and then enter a new date or click the arrow to display a calendar from which you can select a date.</p> <p>In the Time box, select the part of the time that you want to change, and then enter a new time or click the arrows to select a new time.</p>
To an SCN	Restores up to and including a specific system control number (SCN). Type the SCN in the field provided.
Restore read-only files if they are not current	Enables RMAN to examine the headers of all read-only data files and restore any that are not current.
Validate only; do not restore data	<p>Mounts all required media and reads it as necessary. RMAN selects the backup sets that are necessary to perform the operation, and scans them all to ensure that they are available and not corrupted. No data is written or restored to the database server. Validation of the control file is not supported.</p> <p>Symantec recommends that you select this option to ensure that all required media is available before you attempt to restore to the database.</p>
Restore/recover data if validation completes without errors	<p>Lets you run the restore job immediately if the validation was successful.</p> <p>All options that you have selected for restore and recovery are performed.</p>

Table K-11 Oracle restore options (*continued*)

Item	Description
Restore the control file only	Recovers the control file for the Oracle database, but does not include the tablespaces or the associated data files.
Recover using redo logs	<p>Recovers committed transactions from online and archived redo logs. Select an option to recover transactions up to the most recent available, or to a point in time, or to a specific system control number (SCN).</p> <p>RMAN determines which backup objects are necessary for this restore job, and then Backup Exec restores those objects.</p>
To the most recent available	Recovers up to the last committed transaction that is available from the online and archived redo logs.
To a point in time	<p>Recovers committed transactions from the online and archived redo logs up to and including a point in time. After the point in time, recovery stops.</p> <p>In the Date box, select the part of the date that you want to change, and then enter a new date or click the arrow to display a calendar from which you can select a date.</p> <p>In the Time box, select the part of the time that you want to change, and then enter a new time or click the arrows to select a new time.</p> <p>This option is only available when the database is in ARCHIVELOG mode.</p> <p>Caution: When you recover to a point in time by using a control file, make sure that the backup time of the control file is before the specified recovery point in time. The database structure should not have changed between the two times.</p>
To an SCN	Recovers committed transactions from the online and archived redo logs to a specific system control number (SCN). After the SCN is recovered, recovery stops.
After recovery, delete archived redo logs that are no longer needed	Deletes older archived redo log files and free space on the hard disk.
Open database after recovery	Ensures that the database is opened as soon as the recovery is finished. Check this option if you want the database to be online after the recovery.

About redirecting a restore of Oracle data

In Backup Exec, you can redirect an Oracle instance or its files by redirecting the following:

- An Oracle instance to another Oracle server.

Note: If you redirect the instance to a different Oracle server, ensure that an instance with the same name and database ID (DBID) is set up on that server. The database status should be Nomount. Refer to your Oracle documentation for details on creating an instance with the same name and database ID.

- An Oracle instance to another Oracle server and specifying alternate paths for the Oracle files.
- Tablespaces, datafiles, and archive logs to an alternate location on the original server.

Symantec recommends that you select only one instance for each redirected restore operation.

Redirecting a restore of Oracle data

You can use Backup Exec to redirect an Oracle instance or Oracle files.

To redirect a restore of Oracle data

- 1 Create a restore job for Oracle data.
See [“Restoring Oracle data”](#) on page 1343.
- 2 After you select options on the **Restore Job Properties** dialog box for Oracle, on the **Properties** pane, under **Destination**, click **Oracle Redirection**.
- 3 Select the appropriate options.
See [“Oracle redirection options ”](#) on page 1347.
- 4 Start the redirected restore job or select other restore options from the **Properties** pane.

After the restore job is complete, Symantec recommends that you run a full backup of the restored data.

See [“About backing up Oracle resources”](#) on page 1335.

Oracle redirection options

You can set the following options when you restore an Oracle instance to a different server.

See “[Redirecting a restore of Oracle data](#)” on page 1347.

Table K-12 Oracle Redirection options

Item	Description
Restore Oracle instance to server	Lets you redirect the restore of the Oracle instance to a server other than the source server.
Server	Indicates the name of the server to which you want to redirect the restore job.
Server logon account	Displays a logon account that has rights to restore data to the server to which you want to redirect the restore job.
Instance logon account	Displays a logon account for the Oracle instance that you want to restore.
Restore datafiles to the following path	<p>Lets you specify the path to which you want to restore the datafiles. You must enter a valid path or the restore job fails.</p> <p>If you checked Restore Oracle instance to server, use this option to specify a path other than the default on that server.</p> <p>If you do not want to redirect the instance, check this option to specify an alternate local path for the file on the Oracle server.</p>
Restore archived log files to the following path	<p>Lets you specify the path to which you want to restore the archived log files. You must enter a valid path or the restore job fails.</p> <p>If you checked Restore Oracle instance to server, check this option to specify a path other than the default on that server.</p> <p>If you do not want to redirect the instance, check this option to specify an alternate local path for the file on the Oracle server</p>

Restoring from a legacy GRFS Oracle Agent database backup

To restore legacy GRFS Oracle Agent datafile backups, you must use a redirected file system restore job. After you restore the datafiles, you must use your Oracle database administration tools to add the datafile to the intended Oracle instance.

See your Oracle documentation for assistance.

To restore from a legacy GRFS Oracle Agent database backup

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 In the **Properties** pane, under **Source**, click **Selections**.
- 4 On the **View by Resource** tab, expand **All Resources**.

- 5 Browse to the GRFS Oracle resource that you want to restore.
- 6 On the results pane, check the check box of the Oracle datafile that you want to restore.
- 7 In the **Properties** pane, under **Destination**, click **File Redirection**.
Do not click **Oracle Redirection**. This is a file system restore process only.
- 8 Check the Redirect file sets checkbox.
- 9 Type a drive letter in **Restore to drive**.
- 10 Type a path in **Restore to path**.
- 11 Do one of the following:

To run the job now Click **Run Now**.

To schedule the job to run later Do the following in the order listed:

- In the Properties pane, under Frequency, click **Schedule**.
- Set the scheduling options.
- Click **Submit**.

See [“Scheduling jobs”](#) on page 356.

- 12 Rename the restored datafile to match the name of the original Oracle datafile.
- 13 Add the datafile to the intended Oracle instance.
See your Oracle documentation for assistance.

Requirements for recovering the complete Oracle instance and database using the original Oracle server

If you experience a complete loss, deletion, or destruction of the Oracle instance or database, you can use the same Oracle server for the recovery. You can also use these instructions when you configure a new physical server that uses the same server name and SID name.

To successfully complete the recovery using this scenario, you must have the following items:

Table K-13 Requirements when you recover using the original Oracle server

Item	Description
DBID	If you do not know the DBID, you can find it in the Backup Exec job log or in RMAN after you login.

Table K-13 Requirements when you recover using the original Oracle server
(continued)

Item	Description
ControlFile piece ID	You can identify the ControlFile piece ID in the Backup Exec restore view in the Control Files subnode under the Oracle node.
A full system Oracle backup	The full system Oracle backup must include the following: <ul style="list-style-type: none"> ■ controlfile ■ datafiles ■ archive logs
The original Oracle server	To successfully recover the Oracle system using disaster recovery scenario 1, you must restore to the original Oracle server.

Recovering the complete Oracle instance and database using the original Oracle server

You can use the same Oracle server for a recovery if you experience a complete loss, deletion, or destruction of the Oracle instance or database.

See [“Requirements for recovering the complete Oracle instance and database using the original Oracle server”](#) on page 1349.

To recover the complete Oracle instance or database using the original Oracle server

- 1 Recreate the Oracle database using the same name you used for the original database that was lost.
- 2 Find and rename the `pwd<SID>.ora` file.
- 3 Do the following in the order listed to create a new `pwd<SID>.ora` file:
 - Open a command prompt.
 - Type the following command:

```
orapwd file=path\pwsid.ora password=<password>
```
- 4 Type the following commands in the order listed:
 - `RMAN`
 - `CONNECT TARGET <sys/password@sid>;`
 - `SHUTDOWN ABORT;`
 - `STARTUP NOMOUNT;`
 - `SET DBID<dbid ID>;`

- 5 Move to the Backup Exec media server.
- 6 On the navigation bar, click the arrow next to **Restore** and click **New Restore Job**.
- 7 On the **Properties** pane, under **Source**, click **Selections**.
- 8 Select the appropriate ControlFile for restore.
- 9 Click **Run Now**.
The restore job will fail because the recovery portion encounters inconsistent archive logs. This is a normal occurrence during a disaster recovery.
- 10 After the restore job finishes, exit Backup Exec.
- 11 At the Oracle server command prompt, type:
Alter database open resetlogs;
- 12 Close the command prompt.

Requirements for recovering the complete Oracle instance or database to a computer other than the original Oracle server

If you experience a complete loss, deletion, or destruction of the Oracle instance or database, you can restore the instance and database to a computer other than the original Oracle server.

See [“Recovering the complete Oracle instance or database to a computer other than the original Oracle server”](#) on page 1352.

To successfully complete the recovery, you must have the following items:

Table K-14 Requirements when you recover using a new or alternate Oracle server

Item	Description
DBID	If you do not know the DBID, you can find it in the Backup Exec job log or in RMAN after you login.
ControlFile piece ID	You can identify the ControlFile piece ID in the Backup Exec restore view in the Control Files subnode under the Oracle node.
A full system Oracle backup	The full system Oracle backup must include the following: <ul style="list-style-type: none"> ■ controlfile ■ datafiles ■ archive logs

Recovering the complete Oracle instance or database to a computer other than the original Oracle server

You can restore an Oracle instance or database to a computer other than the original Oracle server.

See [“Requirements for recovering the complete Oracle instance or database to a computer other than the original Oracle server”](#) on page 1351.

To recover the complete Oracle instance and database to a computer other than the original Oracle Server

- 1 Recreate the Oracle instance using the same name you used for the original instance that was lost.
- 2 Find and rename the `pwd<SID>.ora` file.
- 3 Do the following in the order listed to create a new `pwd<SID>.ora` file:
 - Open a command prompt.
 - Type the following command:
`orapwd file=path\pwdsid.ora password=<password>`
- 4 Type the following commands in the order listed:
 - **RMAN**
 - **CONNECT TARGET <sys/password@sid>;**
 - **SHUTDOWN ABORT;**
 - **STARTUP NOMOUNT;**
 - **SET DBID<dbid ID>;**
- 5 Move to the Backup Exec media server.
- 6 On the navigation bar, click the arrow next to **Restore** and click **New Restore Job**.
- 7 On the **Properties** pane, under **Source**, click **Selections**.
- 8 Select the appropriate ControlFile to restore.
- 9 On the **Restore job properties** pane, under **Destination**, click **Oracle Redirection**.
- 10 Check the check box for the option, **Restore Oracle instance to server**.
- 11 Enter account credentials to access the new or alternate Oracle server.
- 12 Check the check box for the option, **Restore datafiles to the following path**
- 13 Type a path to the new database.

14 Check the check box for the option, **Restore archived log files to the following path**

15 Click **Run Now**.

The restore job will fail because the recovery portion encounters inconsistent archive logs. This is a normal occurrence during a disaster recovery.

16 Move to the Oracle server.

17 Type **Alter database open resetlogs;**

18 Do one of the following:

If an error is encountered while Oracle tries to open the database Note the online redo log path and then update the path. See [“Updating the online redo log file path”](#) on page 1358.

If an error does not occur Do nothing. The disaster recovery is complete.

Troubleshooting the Oracle Agent

If you have a problem with the Oracle Agent, the following questions and answers may help you solve the problem.

Table K-15 Questions and answers about the Oracle Agent

Question	Answer
<p>What should I do if I get a message that an attempt by Backup Exec to change the state of the Oracle database timed out?</p>	<p>For media server operations, the Oracle database may take some time to change states, such as from open to shut down, from shut down to mount, and so on. A SQLplus script in Backup Exec allows a default time-out of 10 minutes to handle the changing database state. For Oracle Real Application Cluster (RAC), a srvctl script is used.</p> <p>The time-out for database state change is named <code>SqlplusTimeout</code>.</p> <p>You may need to change the length of the default time-out if the following error message appears:</p> <p>An attempt by Backup Exec to change the state of the database timed out. For details, refer to the Database Script output section in the job log. Contact your database administrator to change the state of the database.</p> <p>Try shutting down the database. If you succeed, then the SQLplus time-out is too short. Change the default time-out appropriately, based on how long it took to shut down the database. If you cannot shut down the database, contact your DBA to troubleshoot the database.</p> <p>If the time-out is too short, then restore jobs and offline backups may fail with a time-out error. If the time-out is too long, and the database does not respond to the state change request, the job takes longer to fail.</p> <p>See “Changing the SqlplusTimeout for Oracle instances on Windows computers” on page 1356.</p> <p>See “Changing the SqlplusTimeout for Oracle instances on Linux computers” on page 1357.</p>

Table K-15 Questions and answers about the Oracle Agent (*continued*)

Question	Answer
What should I do if a job continues to run on the media server even after it ends on the Oracle RMAN console?	<p data-bbox="821 326 1239 852">When a backup or restore operation is run on an automatically allocated channel, and if the channel is not released, the job continues to run on the media server even after the operation ends on the RMAN console. The channel is not released if the RMAN console is not exited, or if a new manual channel is not allocated on that console. The job ends on the media server when either the automatic channel is released, or after a time-out period elapses without any activity on that channel, whichever occurs first. If a new backup or restore operation is started within the time-out period on the same automatic channel, a new job is not created. Instead, the existing job performs the operation at the media server.</p> <p data-bbox="821 869 1239 1100">The channel time-out has a default value of 10 minutes, which is recommended for most purposes. If the time-out is too short, then multiple jobs are created for successive operations on a channel. If the time-out is too long, the job runs for a long time on the media server unnecessarily, after the operation has ended.</p> <p data-bbox="821 1117 1239 1204">See “Changing the time-out for an automatic RMAN channel for Oracle instances on Windows computers” on page 1357.</p> <p data-bbox="821 1222 1239 1308">See “Changing the time-out for an automatic RMAN channel for Oracle instances on Linux computers” on page 1358.</p>

Table K-15 Questions and answers about the Oracle Agent (*continued*)

Question	Answer
<p>The error "Unable to attach to a resource..." is displayed when Oracle instance information changes</p>	<p>Whenever Oracle instance information changes, you must update the Remote Agent Utility. If credential information is not updated or is incorrect, the error "Unable to attach to a resource..." may be displayed when you run a backup job. If this message appears, you must bring the server online and configure the information.</p> <p>See "Configuring the Oracle Agent on Windows computers and Linux servers" on page 1318.</p>
<p>What should I do if the error ORA-12546: TNS: Permission denied appears on the Linux computer where Oracle is installed?</p>	<p>If a Backup Exec operation fails on the Linux computer on which the Oracle instances are installed, and the error in the RMAN output section is ORA-12546: TNS: Permission denied, then you must change the machine-level resource credentials in the job. The resource credentials must be an account that is a member of the dba and beoper groups on the Linux computer. Retry the operation.</p> <p>See "Setting authentication credentials on the media server for Oracle operations" on page 1330.</p> <p>See "About the Backup Exec operators group for the Remote Agent for Linux or UNIX Servers" on page 1880.</p>

Changing the SqlplusTimeout for Oracle instances on Windows computers

You can change the length of time Backup Exec handles a change in the state of the Oracle database. Backup Exec allows a default time-out of 10 minutes to handle the changing database state.

See ["Troubleshooting the Oracle Agent"](#) on page 1353.

To change the SqlplusTimeout for Oracle instances on Windows computers

- 1 Create a registry entry of the type DWORD in:
`Software\Symantec\Backup Exec\Engine\Agents\XBSA\Oracle RMAN Agent`
- 2 Name the entry SqlplusTimeout.
- 3 Set the time-out value in seconds.
For example, a time-out of 5 minutes is set as 300 seconds.

Changing the SqlplusTimeout for Oracle instances on Linux computers

You can change the length of time Backup Exec handles a change in the state of the Oracle database. Backup Exec allows a default time-out of 10 minutes to handle the changing database state.

See [“Troubleshooting the Oracle Agent”](#) on page 1353.

To change the SqlplusTimeout for Oracle instances on Linux computers

- 1 In a command prompt, type the following:
`vi etc/VRTSralus/ralus.cfg`
- 2 Create the following entry:
`Software\Symantec\Backup Exec\Engine\Agents\XBSA\Oracle RMAN Agent\SqlplusTimeout`
- 3 Set the time-out value in seconds.
For example, a time-out of 5 minutes is set as 300 seconds.

Changing the time-out for an automatic RMAN channel for Oracle instances on Windows computers

You can change the default channel time-out of 10 minutes for an automatic RMAN channel.

See [“Troubleshooting the Oracle Agent”](#) on page 1353.

To change the time-out for an automatic RMAN channel for Oracle instances on Windows computers

- 1 Create a registry entry of the type DWORD in:

```
HKLM\Software\Symantec\Backup Exec\Engine\Agents\XBSA\Oracle RMAN Agent
```

- 2 Name the entry ChannelTime.
- 3 Set the time-out value in minutes.

Changing the time-out for an automatic RMAN channel for Oracle instances on Linux computers

You can change the default channel time-out of 10 minutes for an automatic RMAN channel.

See [“Troubleshooting the Oracle Agent”](#) on page 1353.

To change the time-out for an automatic RMAN channel for Oracle instances on Linux computers

- 1 In a command prompt, type the following:

```
vi etc/VRTSralus/ralus.cfg
```

- 2 Create the following entry:

```
HKLM\Software\Symantec\Backup Exec\Engine\Agents\XBSA\Oracle RMAN Agent <time-out>
```

- 3 Set the time-out value in minutes.

Updating the online redo log file path

You may have to update the online redo log file path during the recovery of a complete Oracle instance or database.

See [“Recovering the complete Oracle instance or database to a computer other than the original Oracle server”](#) on page 1352.

To update the online redo log file path

- 1 At the Oracle server, open a command prompt.
- 2 Type the following commands in the order listed:

```
■ SQLPLUS /nolog  
■ connect<sys/password@SID>;
```

- 3 Type the following SQLPlus command:

```
SQLPLUS ALTER DATABASE RENAME FILE <old path from backup to any  
redolog file name> to <path to expected restored redolog file  
name>;
```

For example,

```
ALTER DATABASE RENAME FILE  
'D:\ORACLE\ORADATA\JACOB\REDO01.LOG' to  
'C:\ORACLE\ORADATA\JACOB\REDO01.LOG';
```

- 4 In the command prompt, type **RMAN**.
- 5 Type the following command at the RMAN prompt:

```
Alter database open resetlogs;
```
- 6 Close the command prompt.

Symantec Backup Exec Agent for SAP Applications

This appendix includes the following topics:

- [About the SAP Agent](#)
- [Requirements for using the SAP Agent](#)
- [About installing the SAP Agent](#)
- [About SAP Agent security and privileges](#)
- [Before backing up SAP data](#)
- [About system level SAP backup jobs](#)
- [About backing up and restoring with the SAP Agent](#)
- [Backing up SAP data with RMAN](#)
- [Restoring SAP data with RMAN](#)
- [Migrating the SAP Agent catalog from _backint.mdb to _backint.xml](#)
- [About backing up a clustered SAP database on Microsoft Cluster Server](#)
- [About backing up MaxDB databases by using the SAP Agent](#)
- [Restoring MaxDB databases by using the SAP Agent](#)
- [About performing disaster recovery using the SAP Agent](#)

About the SAP Agent

The Backup Exec Agent for SAP Applications (SAP Agent) is a separate, add-on component of Backup Exec. It supports multiple media servers that run on the same network. The SAP Agent lets you back up and restore individual files, entire databases, or individual tablespaces in online or offline mode. You can also back up and restore offline redo logs.

The SAP Agent lets you back up and restore SAP® for Oracle and MaxDB databases using one of the following:

- BACKINT, the backup and restore interface developed by SAP
- Oracle's Recovery Manager (RMAN)

To use RMAN to back up SAP for Oracle databases, you must have both the SAP Agent and the Symantec Backup Exec Agent for Oracle on Windows or Linux Servers (Oracle Agent).

To back up Oracle databases that are not managed by SAP, you can purchase the Oracle Agent. To back up SAP applications that are installed on Microsoft SQL Server, you can purchase the Symantec Backup Exec Agent for Microsoft SQL Servers.

The following features are available with the SAP Agent:

- Ability to name jobs.
- Data encryption.
- Data compression at the client side.
- Recipient notification.
- Dedicated network path for SAP jobs.
- Data integrity verification following a backup job.
- Increased protection of the SAP Agent catalog.

See [“Configuring DBA-initiated job settings for SAP”](#) on page 1369.

See [“About encrypting SAP data”](#) on page 1366.

See [“General options for backup jobs and templates”](#) on page 342.

See [“About preserving the integrity of the SAP Agent catalog”](#) on page 1366.

How the SAP Agent works

The SAP Agent acts as a client for Backup Exec. Through the biparam.ini file, you set Backup Exec parameters for the jobs that you submit from SAP interfaces. For

example, you can set a job name, specify a server or device to use for the job, or specify the backup compression mode.

The SAP Agent connects to your Backup Exec server. The requests are processed through Backup Exec. Jobs that you submit from SAP interfaces, through the SAP Agent, are treated as Run Now jobs. If all drives are busy, jobs can be queued in a Backup Exec job queue. From the queue, the Backup Exec administrator can edit or cancel the job.

When jobs are complete, Backup Exec writes a standard job log as it does for any submitted job. You can view the job log from the Backup Exec administration console. The SAP Agent sends the results of its jobs to the SAP tools. When a job is initiated by a SAP interface, a job log is produced with an eight-character name that represents the job. The SAP system stores the job log in the following directory:

```
<x>:\Oracle\<SID>\sapbackup
```

where <x> is the database installation drive and <SID> is the system ID of the Oracle instance. The file is a plain ASCII text file that you can view with any text editor.

The SAP Agent displays errors and their details on the console, so some problems can be solved without looking at the log file.

You can type either of the following commands to view command line help:

- backin/?
- backint/h

The `_backint.xml` file, which stores Backup Exec catalog information about the job, is placed in the folder local to BRTOOLS, a SAP utility program. Generally the path is as follows:

```
Usr\sap\<SID>\sys\exe\run
```

In the event of a disaster, you must restore this file to restore data to the SAP server.

See [“About system level SAP backup jobs”](#) on page 1371.

See [“About performing disaster recovery using the SAP Agent”](#) on page 1380.

About using the SAP Agent with RMAN

Backup Exec requires the Oracle Agent when integrating with RMAN for backing up and restoring data files.

When you back up the database through RMAN the following occurs:

- The data files are backed up by RMAN using the Oracle Agent.
- Control and log files are backed up by BACKINT using the SAP Agent.

You must run the Remote Agent Utility to configure some settings before you run a backup or restore job through RMAN.

See “[About the Backup Exec Oracle Agent](#)” on page 1315.

See “[About using the SAP Agent with RMAN](#)” on page 1363.

See “[About the Remote Agent Utility for Windows Systems](#)” on page 1948.

Requirements for using the SAP Agent

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Other requirements are as follows:

- The SAP Agent license key must be installed on the media server.
- The SAP Agent must be installed on the SAP server.
- The Backup Exec Remote Agent for Windows Systems must be installed on the SAP server.

Note: The Backup Exec Remote Agent for Windows Systems is automatically installed on the SAP server as part of the SAP Agent installation.

- All SAP databases that you want to back up must be placed in ARCHIVE_LOG mode. Also, Automatic Archival must be enabled using the database administration tools.
See the *SAP User Manual* or your *SAP Database Administrator's Guide*.
- All Oracle databases that you want to back up must be managed by the SAP system.
- BRTOOLS 6.40 must be installed on the Oracle server that is being backed up if you use Oracle 9.i. BRTOOLS 7 or 7.10 must be installed on the Oracle server that is being backed up if you use Oracle 10g.
- The SAP Agent must be configured.
See “[Configuring biparam.ini for the SAP Agent](#)” on page 1368.
- The backup operator must be a member of the ORA_DBA group.
- You must create the ORA_<SID>_OPER group and add the backup user.

- For RMAN backup and restore jobs, the Backup Exec Oracle Agent must be installed on the SAP server.

The SAP Agent follows the BC-BRI BACKINT Interface for ORACLE Databases specification, version 3.0.

Note: The SAP Agent does not support the data that is stored on RAW partition types.

See [“Installing a custom installation of Backup Exec”](#) on page 119.

See [“System requirements”](#) on page 114.

About installing the SAP Agent

Before you install the SAP Agent, do the following:

- Ensure that the backup operator is a member of ORA_DBA. (If the backup operator is Administrator, add Administrator to ORA_DBA group).
- Create group ORA_<SID>_OPER and add the current user to this group.
- Verify that the SAP system environment variables are set for the server that you plan to back up. The SAP Agent default directory is in the SAP database's home as follows:

```
Usr\sap\<SID>\sys\exe\run
```

You can install the SAP Agent to your media server when you install Backup Exec. Or, if you have already installed Backup Exec, you can install only the SAP Agent.

See [“Installing a custom installation of Backup Exec”](#) on page 119.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

About SAP Agent security and privileges

Since BACKINT functions as a client to the Backup Exec server, all Microsoft Windows security limitations that apply to the Backup Exec administration console also apply to BACKINT.

You must have the appropriate privileges on both the SAP and Backup Exec media servers to be able to back up and restore data.

The Backup Exec service account must have the following:

- Access to selections in the jobs that are submitted by the BACKINT interface.
- Rights to the volumes on which the selections are contained.

See [“About encrypting SAP data”](#) on page 1366.

See [“About generating SAP Agent alerts”](#) on page 1366.

See [“About preserving the integrity of the SAP Agent catalog”](#) on page 1366.

See [“About performing disaster recovery using the SAP Agent”](#) on page 1380.

See [“About changing Windows security”](#) on page 108.

About encrypting SAP data

The SAP Agent lets you encrypt data with encryption keys. This feature is available for backup jobs done using either RMAN or BACKINT.

Create the encryption key through the DBA-initiated Job Settings when you create or edit a job template. You must specify the job template name in the biparam.ini file.

If you run jobs with a particular encryption key and the key is deleted, the data that is backed up using the encryption key cannot be restored.

See [“Configuring DBA-initiated job settings for SAP”](#) on page 1369.

See [“Encryption keys”](#) on page 416.

About generating SAP Agent alerts

Backup Exec generates alerts, which BACKINT handles in one of the following modes:

Table L-1 SAP Agent alert modes

Item	Description
Unattended mode	In Unattended mode, any alert that requires a response causes the job to fail and the alert to appear in the SAP system console. Informational alerts also appear in the SAP system console.
Interactive mode	In Interactive mode, all alerts appear in the SAP system console. You must respond to the alerts for the job to continue.

About preserving the integrity of the SAP Agent catalog

You can preserve the integrity of the SAP Agent catalog (`_backint.xml`) in the following ways:

- By restricting which groups can access `_backint.xml`.
- By backing up the catalog along with normal backup data.

Only users in the Administrators or Backup Operators groups can access the SAP Agent catalog (`_backint.xml`). As an administrator, you can give other users permission to access `_backint.xml`. However, you should not revoke the default permission that is given to the Administrators and Backup Operators groups.

You can include a backup of `_backint.xml` along with your normal backups by setting the parameter Backup Catalog to `on` in `biparam.ini`. Catalog backups are recommended at least monthly, but backing up the catalog more frequently provides additional safeguards if the catalog becomes corrupted.

Note: To restore the SAP Agent catalog, you must create a restore job from the Backup Exec media server.

See [“Configuring `biparam.ini` for the SAP Agent”](#) on page 1368.

Before backing up SAP data

Before you submit a backup operation, you must do the following:

- Place all SAP databases to be backed up in the `ARCHIVE_LOG` mode.
- Enable Automatic Archival using the SAP interfaces.
For more information on how to enable Automatic Archival, consult your SAP or Oracle documentation.

Caution: The SAP Agent does not support concurrent backup or restore operations. If you attempt to back up or restore the same SAP database from more than one media server at the same time, the job fails.

If you use RMAN, you must do the following:

- Configure the SAP Agent by using the Backup Exec Remote Agent Utility.
See [“About the Remote Agent Utility for Windows Systems”](#) on page 1948.
- Add the Oracle Server to the media server’s **Modify** list.
See [“Setting authentication credentials on the media server for Oracle operations”](#) on page 1330.

If you want to enter specific configuration information, you should also edit the the following files:

- The biparam.ini file that is found in your `Usr\sap\<SID>\sys\exe\run` directory.
- The `init<SAP>.sap` file that is found in the `<ORACLE_HOME>\database` folder.

See [“Configuring biparam.ini for the SAP Agent”](#) on page 1368.

See [“Setting application defaults for Oracle”](#) on page 1334.

Configuring biparam.ini for the SAP Agent

The BACKINT interface lets you specify Backup Exec parameters for backup jobs that you submit from the SAP interfaces. The parameters are stored in the backup utility parameter file, `biparam.ini`. A template of this file is installed with the SAP Agent.

SAP requires that all SAP tools be in a common folder. The SAP Agent is installed in the same folder as `BRTOOLS`.

The `biparam.ini` file lets you specify job parameters from a single location for operations through `RMAN` and `BACKINT`.

To configure biparam.ini

- 1 Do one of the following:
 - Ensure that the `biparam.ini` file is local to the `BRTOOLS` and the SAP Agent installation.

- Ensure that the parameter `util_par_file` in the `<ORACLE_HOME>database\init<SID>.sap` file specifies the path of the `biparam.ini` file.
- 2 Edit any of the following options in the `biparam.ini` file to configure the media server:

Server=<server name>	Name of the Backup Exec server that will process this backup job. Restore jobs are automatically directed to the Backup Exec server from where the original backup job was executed (except in the case of restore jobs through RMAN). The default is the local computer.
Job Name=<Job Name>	User-specified Job Name. The default is the media server generated Job Name.
Job Template=<DBA-initiated Job Template on Media Server>	The job template that will be used for this job. The job template includes settings like device, media, and encryption key. The DBA-initiated Job Template has to be created at the media server. See “Configuring DBA-initiated job settings for SAP” on page 1369. The default value is the DEFAULT job template. All job parameters would be considered from the DEFAULT job template in this case. If DEFAULT job template is unavailable in the media server and if the user does not mention any job template name in <code>biparam.ini</code> , then the job fails. The job also fails if you specify a wrong job template name.
Backup Catalog=<On/Off>	On: The backup job includes the client-side catalog file (<code>_backint.xml</code>) Off: The backup job does not include the client-side catalog file (<code>_backint.xml</code>) The default is Off.

- 3 Save the file.

Configuring DBA-initiated job settings for SAP

When you create a DBA-initiated backup operation, you can specify the default job template in Backup Exec, or specify a new job template that you create in Backup Exec. The job template contains the settings that Backup Exec applies to DBA-initiated jobs.

Make sure that the name of the job template that you want to use is configured in the biparam.ini file.

See [“Configuring biparam.ini for the SAP Agent”](#) on page 1368.

Note the following about DBA-initiated jobs:

- DBA-initiated jobs fail when the related job template is deleted. To stop DBA-initiated jobs from running, delete the related DBA-initiated job template.
- All DBA-initiated backup and restore jobs are deleted after the jobs are complete.
- You cannot set minimum device requirements for DBA-initiated jobs.

To configure DBA-initiated job settings for SAP

- 1 On the **Tools** menu, click **Options**.
- 2 Click **DBA-initiated Job Settings**.
- 3 Do any of the following:

To create a new job template Do the following in the order listed:

- Click **New**.
- Proceed to step 4.

To edit a job template Do the following in the order listed:

- Select the job template that you want to edit.
- Click **Edit**.
- Proceed to step 4.

To delete a job template Do the following in the order listed:

- Select the job template that you want to delete.
- Click **Delete**.

- 4 On the DBA-initiated Job Settings pane, under **Destination**, click **Device and Media**, and then complete the options as appropriate.

See [“Device and media options for backup jobs and templates”](#) on page 339.

Some options are not available for the DBA-initiated job settings.

- 5 On the DBA-initiated Job Settings pane, under **Settings**, click **General**, and then complete the options as appropriate.

See [“General options for backup jobs and templates”](#) on page 342.

Some options are not available for the DBA-initiated job settings.

- 6 On the **Properties** pane, under **Settings**, click **Network and Security**, and then complete the options as appropriate.
See [“Encryption keys”](#) on page 416.
Some options are not available for the DBA-initiated job settings.
- 7 If you want Backup Exec to notify someone when the backup job completes, click **Notification**, and then complete the options as appropriate.
See [“Sending a notification when a job completes”](#) on page 682.
- 8 Click **OK**.

About system level SAP backup jobs

The more frequently you back up the SAP database, the less time it takes to recover it in the event of data loss. In addition to regularly scheduled SAP Agent backups, it is recommended that you close the SAP database and run a file-level backup whenever you alter the structure of the database.

Along with backing up the SAP database files, you should do the following:

- Create a backup of the Windows directory on the SAP database server and include the Windows registry.
- Back up the `_backint.xml` file. The `_backint.xml` file is usually located in the following directory:

```
Usr\sap\<<SID>\sys\exe\run
```

See [“About performing disaster recovery using the SAP Agent”](#) on page 1380.

About backing up and restoring with the SAP Agent

You create a backup job for SAP applications in one of the following ways:

- Using the CCMS console, which is a graphical user interface.
- Using BRTOOLS, which is a command line utility.

When you back up the database from the CCMS console, status messages appear on the console. These messages report when the database server starts, or stops. They also report when the backup mode of the tables changes. The detail debugging messages and log messages also appear. After all files are backed up, the entire file list appears in the format that is required by the SAP Agent BACKINT interface specification and reports the success or failure of the submitted job.

When you back up the database from BRTOOLS, you can set the backup_mode by changing the init<ORACLE_SID>.sap parameter file. For example:

```
backup_type = online_file
```

Or, you can specify -d with an appropriate backup type on the BRBACKUP command line. For example, on the command line, type:

```
-d util_file_online
```

This command provides a better online backup of very large files, since only the necessary tablespaces are placed in backup mode. When Backup Exec is ready to process another file, it notifies BRBACKUP.

Requirements for submitting jobs from remote computers by using the SAP Agent

If the SAP Agent and the Backup Exec media server are installed on different computers, the following requirements must be met in order for backup and restore jobs to be successful:

- The computer on which the SAP Agent is installed and the media server must be in the same domain.
- The System Logon Account must exist on both the computer on which the SAP Agent is installed and on the Backup Exec media server.
- The System Logon Account must be a member of either the Administrator or Backup Operators groups on both the computer on which the SAP Agent is installed and on the Backup Exec media server.

Restoring data with BRRESTORE and the SAP Agent

BRRESTORE, the BRTOOL utility for restoring data, submits the BID and filename list to the SAP Agent BACKINT interface. BACKINT checks the date and time when the backup was made and uses Backup Exec to recover the file. BACKINT monitors the progress of the restore job and reports status back to BRRESTORE.

When the job completes, BACKINT saves a copy of the Backup Exec restore logs for auditing purposes. You must restart the database.

To restore data with BRRESTORE and the SAP Agent

- ◆ Do one of the following:

To restore data

Type the following command:

```
BRRESTORE -d util_file -b last -m full
```


To restore the database

Type the following commands:

```
SQL>startup mount  
SQL>recover database  
SQL>alter database open;
```

About redirecting SAP restore jobs

The SAP Agent lets you redirect restore jobs to both local and remote computers. If you are redirecting to a remote computer, you must use a valid full UNC path for the location.

For example, if you want to restore a tablespace that originally existing on ComputerA to D:\RestoreDirectory on ComputerB, type:

```
brrestore -d util_file -b <last | logfile name> -m <tablespace to  
restore>=\\ComputerB\D$\RestoreDirectory
```

Note: The System Logon Account for the media server must be a member of either the Administrators or Backup Operators groups on the computer on which the data is being restored.

Backing up SAP data with RMAN

Backup Exec integrates with RMAN, an Oracle utility that does the following:

- Manages backup operations
- Creates backups of database files

To back up SAP data using RMAN, the Backup Exec Oracle Agent is required, and the Remote Agent Utility must be run first.

See [“About the Remote Agent Utility for Windows Systems”](#) on page 1948.

You must modify the rman_send parameters in the Init<SID>.sap file as follows:

```
rman_send = ( "channel sbt_1 'NBBSA_SAP_AGENT_CONFIG_PATH=<INI file  
Path>'")
```

where <INI file path> is the full path for the biparam.ini. For example:
C:\oracle\ora92\database\biparam.ini.

See [“Configuring biparam.ini for the SAP Agent”](#) on page 1368.

Ensure that the ini file path mentioned in the parameter 'util_par_file' in the init<sid>.sap is the same path specified in the rman_send command.

For backups and restores jobs done using RMAN, Symantec Backup Exec does not honor the configuration parameter (-r option) passed from BRBACKUP or BRRESTORE.

Update init<SID>.ora as follows:

```
control_file_record_keep_time <n>, (say n = 45)
```

This parameter controls the minimum number of days that a reusable record is kept in the control file.

To do an online backup, type the following command:

```
brbackup -d rman_util -t online -m all
```

To do an offline backup, type the following command:

```
brbackup -d rman_util -t offline -m all
```

Note: Before doing restores make sure that the database is in the mount state.

Before running online backup jobs, run the following scripts:

```
$ORACLE_HOME\rdbms\admin\catalog.sql  
$ORACLE_HOME\rdbms\admin\catspace.sql  
$ORACLE_HOME\rdbms\admin\catproc.sql
```

These scripts will configure the database for an online backup. If the database is not configured properly, the job could fail.

If you get the following error:

```
RMAN-00571:  
=====  
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS  
=====  
RMAN-00571:  
=====  
ORA-06550: line 1, column 7:  
  
PLS-00201: identifier 'DBMS_BACKUP_RESTORE.SET_CHARSET' must be  
declared  
  
ORA-06550: line 1, column 7:  
  
PL/SQL: Statement ignored
```

```
RMAN-04015: error setting target database character set to  
WE8MSWIN1252
```

Run the following scripts:

```
$ORACLE_HOME\rdbms\admin\catalog.sql  
$ORACLE_HOME\rdbms\admin\catspace.sql  
$ORACLE_HOME\rdbms\admin\catproc.sql
```

After running the scripts, run the backup job again.

Restoring SAP data with RMAN

Backup Exec integrates with RMAN, an Oracle utility that does the following:

- Manages recovery operations.
- Restores or recovers a database from backups.

To restore SAP data using RMAN, the Backup Exec Oracle Agent is required, and the Remote Agent Utility must be run first.

See [“About the Remote Agent Utility for Windows Systems”](#) on page 1948.

You must modify the `rman_send` parameters in the `Init<SID>.sap` file as follows:

```
rman_send = ( "channel sbt_1 'NBBSA_SAP_AGENT_CONFIG_PATH=<INI file  
Path>'")
```

where `<INI Path >` is the full path for the `biparam.ini`. For example:
`C:\oracle\ora92\database\biparam.ini`.

See [“Configuring biparam.ini for the SAP Agent”](#) on page 1368.

Ensure that the ini file path mentioned in the parameter `'util_par_file'` in the `init<sid>.sap` is the same path specified in the `rman_send` command.

For restores done using RMAN, Symantec Backup Exec does not honor the configuration parameter (`-r` option) passed from `BRBACKUP` or `BRRESTORE`.

Update `init<SID>.ora` as follows:

```
control_file_record_keep_time <n>, (say n = 45)
```

This parameter controls the minimum number of days that a reusable record is kept in the control file.

When `BRRESTORE` is used to run full restores via RMAN, the database should be in mount state.

The command for restoring database files only is as follows:

```
brrestore -d rman_util -b last -m full
```

For restoring the control files, use the -m 0 [,00] option

For example, to restore the .ctl files type the following on the command line:

```
brrestore -d rman_util -b last -m 0
```

To restore the .dbf files, type the following on the command line:

```
brrestore -d rman_util -b last -m 00
```

When restoring with RMAN, make sure that the media server specified in the biparam.ini file is the media server on which the backup job was done.

Migrating the SAP Agent catalog from _backint.mdb to _backint.xml

This release of Backup Exec stores the SAP Agent catalog in an XML file. If you are upgrading from a previous version of Backup Exec, during the install process select the Upgrade option to migrate all the backup metadata from _backint.mdb to _backint.xml.

If you do not select the Upgrade option during the installation, the earlier backup catalog is not be available to this release of Backup Exec.

You must manually migrate the data in the following situations:

- You did not select the Upgrade option during the time of installation.
- You want Backup Exec to use a catalog created by earlier version of the SAP Agent.

Before migrating the data, do the following:

- Make sure that a file named _backint.xml is not already in the path that you plan to use.
- Make sure that _backint.mdb is in the same location were BRTOOLS and _backint.exe are located.

To manually migrate from `backint.mdb` to `backint.xml`

- 1 Copy the `_backint.mdb` file to the location where the `BRTTOOLS` and `_backint.exe` is located (if not already present).

The migration utility (`MdbToXML.exe`) is located in the directory where the SAP Agent is installed.

- 2 Run the Migration Utility (`MdbToXML.exe`) with proper usage parameters.

For example:

```
MdbToXml.exe <Path for the _backint.mdb> <optional Path for Log file >
```

If `_backint.mdb` is in the path `C:\usr\sap\CER\sys\exe\run`, then the command is:

```
MdbToXML.exe C:\usr\sap\CER\sys\exe\run
```

To get online help for this utility, type the following command:

```
MdbToXML.exe /?
```

The Path for Log file is optional. If a path is not specified, then the log file is created local to the `_backint.xml` file.

The `_backint.xml` file is created in the same path as the `_backint.mdb` file.

After the migration is completed, the `_backint.mdb` is renamed to `_backint_migrated.mdb`.

A log file named `MdbToXmlMigrationLog.txt` is created in the path that is specified in the command line. If that parameter is ignored, then it is created local to the `_backint.xml`. If the path for log file is mentioned incorrectly, no log file is created.

But this does not affect the migration process.

About backing up a clustered SAP database on Microsoft Cluster Server

The Backup Exec for SAP Agent supports backup and restore jobs in a clustered environment for Oracle with the help of Microsoft Cluster Server (MSCS) and Oracle Failsafe.

To use the SAP Agent in a clustered environment, do the following:

- Install both MSCS and Oracle along with Oracle Failsafe on both nodes of a two-node cluster environment.

- Install the database that you want to back up on the shared disk to ensure that the database is properly failed over to the other node.
- Ensure that the cluster has a virtual cluster name configured.
- Ensure the Is Alive poll interval of the Oracle database resource is more than the average time needed to back up the full database.

See your Oracle Failsafe documentation for additional information on configuring and installing Oracle Failsafe.

See your Microsoft documentation for additional information on installing and configuring MSCS.

Note: In a cluster environment, if a job created through BACKINT or RMAN is processing and the node fails over, the job operation does not restart from the point when the node went down.

About backing up MaxDB databases by using the SAP Agent

The SAP Agent supports the backup of SAP applications that run on the MaxDB(SAPDB) database.

Use the DBM command line interface or the MaxDB Database Manager graphical user interface to initialize and send a backup job to MaxDB's BACKINT interface. MaxDB's BACKINT then sends the backup job to the Symantec SAP Agent BACKINT interface, which executes the job. The SAP Agent supports backups to both local and remote Backup Exec media servers.

The SAP Agent supports the following SAP DB backup functions:

- Backup of complete data.
- Backup of incremental data.
- Backup of log files.

Note: The SAP Agent must be installed in a directory structure that is only two directories deep. If this directory structure is changed, and then change the SAP.PAR and BSI.ENV files accordingly.

See [“Preparing MaxDB databases for backup”](#) on page 1379.

See [“Backing up MaxDB databases”](#) on page 1379.

Preparing MaxDB databases for backup

Use the following steps to prepare MaxDB databases for backup.

To prepare MaxDB databases for backup

- 1 Ensure that the following files are present:
 - SAP.PAR
 - BSI.ENV
- 2 Ensure that the following parameters are set during the installation of MaxDB:
 - independent program path = C:\sapdb\programs
 - dependant path to C:\sapdb\\db
 - independent data path = C:\sapdb\data
- 3 Ensure that the following paths exist:
 - C:\sapdb\programs
 - C:\sapdb\data
 - C:\sapdb\ - C:\sapdb\
- 4 Ensure that the appropriate versions of the following applications are installed:
 - SQL studio
 - DBMGUI

Backing up MaxDB databases

Use the following steps to back up MaxDB databases.

To back up MaxDB databases

- 1 Copy BSI.ENV into C:\sapdb\data\wrk\- 2 Copy SAP.PAR into C:\sapdb\- 3 From the MaxDB Database Manager, select the Backup Wizard.
- 4 Provide the appropriate inputs for the backup job, including the following:
 - Type of backup

- Pipe used
 - Name of pipe
- 5 Click **Start**.

Restoring MaxDB databases by using the SAP Agent

The SAP Agent supports the restore of SAP applications that run on the MaxDB(SAPDB) database.

Use the DBM command line interface or the MaxDB Database Manager graphical user interface to initialize and send a restore job to MaxDB's BACKINT interface. MaxDB's BACKINT then sends the restore job to the Symantec SAP Agent BACKINT interface, which executes the job. The SAP Agent supports restores to both local and remote Backup Exec media servers.

The SAP Agent supports the following SAP DB restore functions:

- Restore last backup.
- Restore a specified backup from history.
- Restore a medium.
- Restore database to a specific time.

To restore data

- ◆ From the MaxDB Database Manager, run the Recovery Wizard, ensuring the database is in Admin mode.

About performing disaster recovery using the SAP Agent

To recover your SAP database server after a catastrophic failure, you must implement a backup strategy before a failure occurs.

When developing a disaster recovery plan, the following backup strategies should be used:

- Make at least one flat file database backup and make regular offline backups using CCMS. See your *SAP Database Administrator's Guide*.
- Back up the Windows directory on the SAP database server, including the Windows registry.
- If the structure of the database is altered, perform a full offline database backup.

For example, if you create a new tablespace or remove an old one, make a complete offline database backup.

- Always include the `backint.xml` file in your regular flat file backups of the SAP database server.

The `backint.xml` file is usually located in the following directory:

```
Usr\sap\<SID>\sys\exe\run
```

- Schedule full online backups of your SAP database server regularly.

See “[Configuring biparam.ini for the SAP Agent](#)” on page 1368.

SAP disaster recovery prerequisites

The following backups are required to fully recover your SAP database server in the event of a disaster.

- Create a full SAP database server file system backup using Backup Exec. When creating this backup, include both the SAP database directory and the Windows system directory. However, if the database must remain open, do not include the SAP database tablespace data files in this backup.
- Create a second backup containing the SAP database tablespace data files. See your *SAP Database Administrators Guide* for details.

After creating these backups, you can recover your SAP database server as needed.

See “[Creating a backup job by setting job properties](#)” on page 332.

Recovering a remote SAP database server from a disaster

Recovering a remote SAP database involves re-installing the Microsoft Windows operating system and restoring files from a recent backup.

To recover a remote SAP database server from a disaster

- 1 Re-install the Microsoft Windows operating system on the SAP database server.

During the re-install process, install Windows into a temporary directory that you can delete after your SAP database server is back up and running.
- 2 At the media server, and using the storage media containing the FULL flat file SAP database server file system backup, restore the entire contents of the media to the SAP database server using Backup Exec.

3 Reboot your SAP database server.

The computer reboots using its original version of Windows. The system contains the original version of Windows, the SAP Agent, the SAP database minus the tablespaces, and any other files contained on the full backup media.

4 If you have a full offline SAP database backup, restore your last full offline SAP database backup and start your database.

If you do not have a full offline database backup, your database is operational. Proceed to step 5.

5 Restore the backint.xml file from the latest full server backup.

The backint.xml file correlates the SAP catalog with the Backup Exec catalog.

6 To bring your database up-to-date, restore your most recent online or offline SAP database backup.

7 At the media server, run another restore operation. This time, use the SAP Agent to restore the storage media containing all of the SAP tablespace data files.

8 When the restore operation completes, open the CCMS console, and click **Check and Repair Database**.

9 Click **Automatic Recovery**, and then follow the online prompts.

Recovering a combination SAP database server and media server

Recovering a combination SAP database server and media server involves re-installing the Microsoft Windows operating system and restoring files from a recent backup.

To recover a combination SAP database server and media server

1 Re-install the Microsoft Windows operating system on the SAP database server media server.

During the re-install process, install Windows into a temporary directory that you can delete after your SAP database server media server is running.

2 Re-install Backup Exec.

3 Recatalog the media containing the full flat file SAP database server file system backup, and the media containing the SAP database tablespace data files.

- 4** Restore the entire contents of the media containing the full flat file server file system backup.

This restores your original Windows system, along with any services required to run your SAP database.
- 5** Reboot the computer after the restore operation completes.

Because the full system backup was restored, your computer now boots using its original version of Windows. The system now contains the original version of Windows, the SAP Agent, the SAP database minus the tablespaces, and any other files contained on the full backup media.
- 6** Restore the backint.xml file from the latest full server backup.

This file correlates the SAP catalog with the Backup Exec catalog.
- 7** Run another restore operation.

This time, use the SAP Agent to restore the media containing all of the SAP tablespace data files.
- 8** When the restore operation completes, open the CCMS console, and select **Check and Repair Database**.
- 9** Select **Automatic Recovery**, and then follow the online prompts to complete Disaster recovery of the SAP database server.

Symantec Backup Exec Agent for VMware Virtual Infrastructure

This appendix includes the following topics:

- [About the Agent for VMware](#)
- [Requirements for using the Agent for VMware](#)
- [About installing the Agent for VMware](#)
- [Adding VMware vCenter and ESX servers](#)
- [Deleting VMware vCenter and ESX servers](#)
- [About backing up VMware resources](#)
- [Creating a full backup of VMware resources](#)
- [Creating an incremental or a differential backup of VMware resources](#)
- [How Granular Recovery Technology works with the Agent for VMware](#)
- [About protecting databases and applications with the Symantec VSS Provider](#)
- [About restoring VMware resources](#)
- [Redirecting the restore of a VMware virtual machine](#)
- [Setting default backup and restore options for the Agent for VMware](#)

About the Agent for VMware

The Symantec Backup Exec Agent for VMware Virtual Infrastructure (Agent for VMware) lets you back up and restore virtual machines that use the following VMware products:

- ESX Server
- vCenter Server (formerly VirtualCenter)
- vSphere 4.0 and 4.1

Backup Exec performs a single-pass backup to protect all Guest virtual machines and VSS-aware applications that are installed on the Guest virtual machines. Backup Exec's Granular Recovery Technology (GRT) is enabled by default for backup jobs. You can use a GRT-enabled backup to restore individual files and folders from a Windows Guest virtual machine without restoring the entire virtual machine. In addition, you can restore individual items from Microsoft Exchange, SQL, and Active Directory applications that reside on Guest virtual machines.

Additional features of the Agent for VMware let you do the following:

- Redirect the restore of data from a Guest virtual machine to an alternate folder, datastore, host, or network.
- Back up to a disk device or to a tape device.
- Perform policy-based incremental and differential backup jobs (If your virtual machines are configured with hardware version 7).

See [“Requirements for using the Agent for VMware”](#) on page 1386.

See [“How Backup Exec protects Exchange, SQL, and Active Directory data on virtual machines”](#) on page 1397.

See [“About backing up VMware resources”](#) on page 1388.

See [“About restoring VMware resources”](#) on page 1400.

Requirements for using the Agent for VMware

The Agent for VMware uses the following components, which can reside on the same computer or on separate computers:

Table M-1 Agent for VMware components

Item	Description
Backup Exec media server	This component runs the backup and restore jobs. You must enter the Agent for VMware license key on this component.
VMware vCenter Server	This component is optional. It manages the ESX servers. You do not have to install the Remote Agent for Windows Systems (Remote Agent) on this computer. If the Remote Agent is installed, it is used only to publish the vCenter Server to the Backup Exec media server.

To use Backup Exec's Granular Recovery Technology (GRT) with the Agent for VMware, install the Backup Exec Remote Agent for Windows Systems on any virtual machines that run Windows.

See “[How to restore individual items by using Granular Recovery Technology](#)” on page 321.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

About installing the Agent for VMware

The Agent for VMware is installed locally as a separate, add-on component of Backup Exec. You do not have to install the VMware Agent on the ESX host.

See “[Installing additional Backup Exec options to the local media server](#)” on page 123.

Adding VMware vCenter and ESX servers

You can add VMware vCenter and ESX servers to the Backup Exec selection list so that these servers can be selected for backup jobs.

To add VMware vCenter and ESX servers

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 Right-click one of the following:
 - **All Resources**
 - **VMware vCenter and ESX Servers**
- 4 Click **Manage VMware vCenter and ESX Servers**.
- 5 In the **Name** field, type the name of the server that you want to add.
- 6 Click **Add**.
- 7 Click **Close**.

Deleting VMware vCenter and ESX servers

You can delete VMware vCenter and ESX servers from the Backup Exec Database. If a server is not in the Backup Exec Database, then it cannot be selected for backup jobs.

To delete VMware vCenter and ESX servers

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 Right-click one of the following:
 - **All Resources**
 - **VMware VrtualCenter and ESX Servers**
- 4 Click **Manage VMware vCenter and ESX Servers**.
- 5 From the list of servers that appear, right-click the name of the server that you want to delete.
- 6 Click **Delete**.
- 7 Click **Close**.

About backing up VMware resources

When you create a backup job, you can select the following VMware resources:

- An entire vCenter or ESX server, DataCenters, and folders
- Individual virtual machines

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux Servers is installed.

If you select the vCenter or ESX server as a backup resource, all virtual machines are backed up.

If you select to back up the vCenter or ESX server, the backup job does not include the following:

- Configuration files for the vCenter or ESX server
- Physical Raw Disk Mapping (RDM) devices
- Independent disks

Backup Exec can automatically protect new virtual machines and folders that are found when a backup job runs.

See [“How Backup Exec automatically protects new virtual machines during a backup job”](#) on page 1390.

The following backup methods are supported for VMware resources:

Table M-2 Supported backup methods for VMware resources

Backup method	Requirements
Full	<p>This option is available for both VMware vCenter Server and VMware vSphere. It is the only available backup method for VMware resources if you do not use a policy-based backup job.</p> <p>See “Creating a full backup of VMware resources” on page 1390.</p> <p>Backup Exec's Granular Recovery Technology (GRT) lets you use the full image backup to restore individual files for virtual machines that use the Windows operating system. GRT also lets you restore individual items from VSS-aware applications that are installed on the virtual machines.</p> <p>See “How Granular Recovery Technology works with the Agent for VMware” on page 1396.</p>

Table M-2 Supported backup methods for VMware resources (*continued*)

Backup method	Requirements
Incremental or Differential	<p>This option is available only if the virtual machine is configured with hardware version 7. You must use a policy to create a backup job that uses the incremental or differential backup method.</p> <p>See “Creating an incremental or a differential backup of VMware resources” on page 1395.</p>

How Backup Exec automatically protects new virtual machines during a backup job

Backup Exec's dynamic inclusion feature protects new virtual machines and folders that are found when a backup job runs. If new virtual machines are added between the time when the backup job is created and when the backup job runs, Backup Exec automatically backs up the new virtual machines. Because the backup job may include new virtual machines, the job may require more storage space and more time to run than you anticipated. The job history shows the number of virtual machines that were backed up.

In the backup selection list, dynamic inclusion is enabled for the following VMware resources:

- ESX
- vCenter 4
- All nodes that have a folder icon

If you select ESX or vCenter 4, then dynamic inclusion is enabled automatically for all of the nodes below them that have a folder icon. If no virtual machines are located during a backup job, then the job fails.

Creating a full backup of VMware resources

Follow these steps to create a full backup of a VMware vCenter server, an ESX server, or a virtual machine.

If your virtual machines are configured with hardware version 7, you can create an incremental or a differential backup job by using a policy.

See [“Creating an incremental or a differential backup of VMware resources”](#) on page 1395.

To create a full backup of VMware resources

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Properties** pane, under **Source**, click **Selections**.
- 4 Expand **VMware vCenter and ESX Servers**.
- 5 Select one of the following:
 - A VMware vCenter or ESX server.
 - One or more of the virtual machines that appear under the name of a VMware vCenter or ESX server.
- 6 In the **Properties** pane, under **Settings**, click **VMware**.
- 7 Select the appropriate options.
See “[VMware backup options](#)” on page 1391.
- 8 To change the setting for granular recovery for VSS-aware applications that are installed on virtual machines, click **Edit**.
See “[Virtual Machine Application Granular Recovery Technology Settings](#)” on page 1394.
- 9 Start the backup job or select other backup options from the Properties pane.

VMware backup options

The following options are available for VMware backup jobs.

See “[Creating a full backup of VMware resources](#)” on page 1390.

Table M-3 VMware backup options

Item	Description
Backup method	Indicates the backup method to use for the backup job. If your virtual machines are configured with hardware version 7 and you create the backup job from a policy, the Incremental and Differential backup methods are available. If you use VMware vCenter Server, the Full backup method is the only available method. You can use the Incremental and Differential backup methods only if you use a policy, regardless of the version of VMware that you use.

Table M-3 VMware backup options (*continued*)

Item	Description
Use the full backup method for virtual machines that do not support incremental or differential backups	<p>Lets Backup Exec perform a full backup if an incremental backup or a differential backup cannot be performed. If you do not select this option and Backup Exec cannot perform an incremental backup or a differential backup, then the job fails. In addition, if Backup Exec detects a configuration change, then a full backup must be performed. If a configuration change is detected and Backup Exec cannot perform a full backup, then the job fails if this option is not selected. This scenario applies only if a full backup and some incremental or differential backups have already been performed and the next scheduled job is for an incremental or a differential backup.</p> <p>Note: This option appears only when you use a policy.</p>

Table M-3 VMware backup options (*continued*)

Item	Description
Transport mode priority list	<p>Lets you select the method to transport the Virtual Machine Disk Format (VMDK) file from the ESX server. You must select at least one of these options. If you select more than one option, the method is determined by the priority and the resources that are available. You can move the options up or down in the list to change the priority.</p> <p>The following methods are available:</p> <ul style="list-style-type: none"> ■ SAN - Use the SAN to move virtual disk data. If you select this option, the virtual machine must reside on a SAN that the media server can access. With this transport mode, the data is offloaded to the media server, so that the ESX server is not affected. ■ NBD - Do not encrypt the virtual disk data for over-the-network transfers Use this option if you do not use SSL for security and one of the following conditions exist: <ul style="list-style-type: none"> ■ The virtual machine is not located on the SAN. ■ The media server does not have access to the SAN. The data is placed on the ESX server. Then, the data moves across the network. ■ NBDSSL - Encrypt virtual disk data for over-the-network transfers Use this option if you use SSL for security and one of the following conditions exist: <ul style="list-style-type: none"> ■ The virtual machine is not located on the SAN. ■ The media server does not have access to the SAN. ■ Hotadd - Use virtual disk files from the Backup Exec server on the virtual machine Use this option if you want to use the hot add feature for ESX. See your VMware documentation for more information about hot add. <p>The snapshot is placed on the ESX server. Then, the data moves across the network.</p> <p>The VMDK file is not backed up if the virtual hard disk is configured as an Independent disk.</p>
Move Up	Lets you move the selected transport mode to a higher priority in the list.

Table M-3 VMware backup options (*continued*)

Item	Description
Move Down	Lets you move the selected transport mode to a lower priority in the list.
Back up virtual machines that are powered off	Enables Backup Exec to back up virtual machines that are turned off.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	<p>Enables individual files and folders to be restored from the backup. This option is for virtual machines that use a Windows operating system only.</p> <p>The VMDK file is not backed up if the virtual hard disk is configured as an Independent disk.</p> <p>Note: GRT is not meant for system recovery but only for the restore of individual files and folders on Windows computers.</p>
Edit	<p>Lets you change the GRT settings for Microsoft Active Directory, Exchange, and SQL.</p> <p>See “Virtual Machine Application Granular Recovery Technology Settings” on page 1394.</p>
Microsoft Active Directory	Indicates whether GRT is enabled or disabled for Microsoft Active Directory on the virtual machine. It is enabled by default.
Microsoft Exchange	Indicates whether GRT is enabled or disabled for Microsoft Exchange on the virtual machine. It is enabled by default.
Microsoft SQL	Indicates whether GRT is enabled or disabled for Microsoft SQL on the virtual machine. It is enabled by default.
vSphere Port Number	Indicates the port that Backup Exec uses to connect to vCenter Server. The default port is 902.

Virtual Machine Application Granular Recovery Technology Settings

Use the following options to enable or disable granular recovery of individual items from Microsoft Active Directory, Exchange, and SQL.

See [“Creating a full backup of VMware resources”](#) on page 1390.

Note: If you enable or disable Granular Recovery Technology for one of the following applications, the setting applies to both VMware virtual machines and Hyper-V virtual machines. If you do not want to use the same settings, Symantec recommends that you set up separate backup jobs for each type of virtual machine.

Table M-4 Virtual Machine Application Granular Recovery Technology Settings

Item	Description
Enable GRT for Microsoft Active Directory objects on virtual machines	Enables Backup Exec to collect the information that is required to restore individual Active Directory objects on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft Active Directory is installed.
Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines	Enables Backup Exec to collect the information that is required to restore individual Exchange databases and mailbox items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft Exchange is installed.
Enable GRT for Microsoft SQL (database-level only) on virtual machines	Enables Backup Exec to collect the information that is required to restore individual SQL database items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft SQL is installed.

Creating an incremental or a differential backup of VMware resources

If your virtual machines are configured with vSphere 4.0 or 4.1 with hardware version 7, you can create incremental or differential backups of VMware resources by creating policy backup jobs.

See [“Creating a full backup of VMware resources”](#) on page 1390.

Backup Exec includes example policies for VMware incremental and differential backups. These example policies contain the standard settings. You can copy the example policies and then customize them to meet your needs.

See [“Using an example policy”](#) on page 526.

Table M-5 How to create an incremental or a differential backup of VMare resources

Action	For more information
Create a policy.	See “Creating a new policy” on page 522.
Add two backup templates to the policy. One template must use the full backup method and one template must use either the incremental or the differential backup method. Select these methods on the VMware settings.	See “VMware backup options” on page 1391.
Create a new template rule and select the following rule: <Template A> must complete at least once before any other templates will be allowed to start.	See “Adding a backup template to a policy” on page 530.
Create a new job using the policy.	See “Creating new jobs for a policy” on page 544.

How Granular Recovery Technology works with the Agent for VMware

Backup Exec's Granular Recovery Technology (GRT) lets you restore individual drives, files, and folders without having to restore the entire virtual machine. It also lets you restore individual items from VSS-aware applications that are installed on virtual machines.

See [“How Backup Exec protects Exchange, SQL, and Active Directory data on virtual machines”](#) on page 1397.

GRT works only for the virtual machines that use a Windows operating system. GRT does not work for system recovery. You should review the requirements for a GRT-enabled backup before you configure it.

See [“About requirements for jobs that use Granular Recovery Technology”](#) on page 325.

To use GRT, you select the individual files and folders that you want to restore from the list that appears when you expand the network bios or the computer name of the virtual machine. You cannot select individual folders and files from the virtual machines that appear when you expand VMware vCenter and ESX Servers.

See [“About selecting VMware resources for restore”](#) on page 1400.

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

See [“Recommended devices for backups that use Granular Recovery Technology”](#) on page 324.

How Backup Exec protects Exchange, SQL, and Active Directory data on virtual machines

Backup Exec can restore individual items from the following VSS-aware applications that reside on virtual machines:

Table M-6 Types of data that Backup Exec protects for VSS-aware applications on virtual machines

Application	Types of data that Backup Exec protects
Microsoft Exchange	Mailboxes, individual messages, calendar items, tasks, journal entries, and public folder data (disk-backups only)
Microsoft SQL	Databases
Microsoft Active Directory	Individual user accounts, printer objects, sites, and organizational units

When you create a backup job, Backup Exec automatically locates VSS-aware applications on virtual machines. During the backup job, Backup Exec backs up the data from the VSS-aware applications by using Granular Recovery Technology (GRT). By default, Backup Exec enables GRT using the same credentials that were used to connect to the virtual machine. You can disable GRT for any of the VSS-aware application types.

Note: Backup Exec supports the granular recovery of individual Exchange and SQL items only in non-clustered and non-distributed configurations.

See [“VMware backup options”](#) on page 1391.

Note: If you enable or disable GRT for an application, the setting applies to both VMware virtual machines and Hyper-V virtual machines. If you do not want to use the same settings, Symantec recommends that you set up separate backup jobs for each type of virtual machine.

During the backup job, Backup Exec collects metadata from the applications. If Backup Exec is unable to collect the metadata, then you cannot restore individual items for the applications. However, the backup job may otherwise complete successfully.

See [“Requirements for protecting Exchange, SQL, and Active Directory data on virtual machines”](#) on page 1398.

Requirements for protecting Exchange, SQL, and Active Directory data on virtual machines

Backup Exec can back up and restore individual items from VSS-aware applications that are installed on virtual machines.

The following items are required to protect data for Microsoft Exchange, SQL, and Active Directory on virtual machines:

- The virtual machine must be turned on.
- You must enter the appropriate credentials for the virtual machine. Ensure that the credentials for the virtual machine allow access to the VSS-aware applications.
- The media server must be able to connect to the virtual machine using the network name or IP address.
- The Backup Exec Remote Agent for Windows Systems must be installed on the virtual machine.
- The correct number of licenses must be entered for the applications that you want to protect on the virtual machines.
- The operating system on the virtual machine must support VSS.

If you want to use Backup Exec's Granular Recovery Technology (GRT), you must purchase and install the application agents, such as Backup Exec's Agent for Microsoft Exchange, on your virtual machines.

See [“How Backup Exec protects Exchange, SQL, and Active Directory data on virtual machines”](#) on page 1397.

About protecting databases and applications with the Symantec VSS Provider

The Symantec VSS Provider helps Backup Exec protect VSS-aware applications, such as Microsoft Exchange, SQL, and Active Directory. The Symantec VSS

Provider provides an automatic snapshot of the Windows applications and databases for each backup job.

Some of your Guest virtual machines may already have the VMware VSS Provider. However, only one VSS Provider can be used on a Guest virtual machine. Therefore, you must uninstall the VMware VSS Provider.

Warning: Having more than one VSS provider installed on a Guest virtual machine can cause quiesce and other snapshot errors. For example, the virtual machine may time out during a backup job. Verify that your Guest virtual machines have only one VSS provider.

When you install the Remote Agent for Windows Systems on a Guest virtual machine, the Symantec VSS Provider is installed automatically. You can also install it manually from the Backup Exec installation media.

Warning: Do not reinstall the VMware VSS Provider on a Guest virtual machine after you install the Remote Agent for Windows Systems.

See [“Push-installing the Remote Agent and Advanced Open File Option to remote computers”](#) on page 134.

By default, the Symantec VSS Provider takes full backups and truncates database log files. However, you can change the default setting by modifying the script files.

See [“Changing the log truncation setting of the Symantec VSS Provider”](#) on page 1399.

Changing the log truncation setting of the Symantec VSS Provider

By default the Symantec VSS Provider takes full backups and truncates database log files. You can change the settings to enable the Symantec VSS Provider to take copy backups without log truncation.

Note: You must add the -copy flag to the Pre-freeze-script.bat file in both the system root directory and %Programfiles%\Symantec\Backup Exec\RAWS\VSS Provider.

To change the log truncation setting of the Symantec VSS Provider

- 1 Locate the Pre-freeze-script.bat file in both of the following locations:
 - Your system root directory

- %Programfiles%\Symantec\Backup Exec\RAWS\VSS Provider
- 2 Add the -copy flag to the end of each of the three lines that include BeVssRequestor.exe.

For example:

```
"%Programfiles%\Symantec\Backup Exec\BE VSS  
Provider\BeVssRequestor.exe" -pre2 -log -logscreen -copy
```

About restoring VMware resources

You can configure restore jobs to do the following:

- Restore data to the original location or to an alternate location.
- Turn on virtual machines after the restore job completes.
- Restore over an existing virtual machine.
- Restore with a new virtual machine name in vCenter Server.
- Select the preferred network for virtual machines to use after the restore job completes.

Note: Backup Exec does not support the restore of individual Virtual Machine Disks (VMDK files). If you attempt to restore individual VMDK files, the virtual machine is created with all of the disks that were present at backup time and only the selected disk's data is restored. The selection of individual disks is recommended only for redirected restores.

Note: To restore virtual machines that were backed up with Backup Exec 12.5, the VMware Converter (4.01 or later) must be installed on the Backup Exec media server.

See [“About redirecting restore jobs”](#) on page 634.

About selecting VMware resources for restore

You can restore virtual machine data in the following ways:

- You can restore a complete virtual machine, or its Virtual Machine Disk Format (VMDK) file, for disaster recovery purposes.
- You can restore individual files or folders that were backed up from inside the VMDK file, if you selected the Granular Recovery Technology (GRT) option for the backup job.

In the **Restore** view, a virtual machine that was enabled to use GRT appears under its physical network or Netbios name. If you expand the network name, then individual drives, files, and folders appear.

Virtual machines also appear in the **Restore** view under **VMware vCenter and ESX Servers**. Under **VMware vCenter and ESX Servers**, the virtual machines appear by their display name, or the name that you provided for the virtual machine during its creation. If you expand the display name for a virtual machine, its contents appear. If you select the virtual machine by its display name, you can recover the entire virtual machine and its VMDK files.

Restoring VMware resources

By default, Backup Exec restores data to the location from which it was originally backed up. If you want to restore data to a different virtual machine than where the data originally resided, you must create a redirected restore job.

See [“Redirecting the restore of a VMware virtual machine”](#) on page 1404.

Note: To restore virtual machines that were backed up with Backup Exec 12.5, the VMware Converter (4.01 or later) must be installed on the Backup Exec media server.

Note: Granular Recovery Technology (GRT) allows the restore of individual data files and folders. GRT cannot restore system state files such as the active registry.

To restore VMware resources

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Select the data that you want to restore.
See [“About selecting VMware resources for restore”](#) on page 1400.
- 4 In the **Properties** pane, under **Settings**, click **VMware**.
- 5 Select the appropriate options.
See [“VMware restore options”](#) on page 1401.
- 6 Start the restore job or select other restore options from the **Properties** pane.

VMware restore options

The following options are available for VMware restore jobs.

See [“Restoring VMware resources”](#) on page 1401.

Table M-7 VMware restore job options

Item	Description
Delete existing virtual machines prior to restore	Deletes existing virtual machines during the restore job. If you select this option, the virtual machines may be deleted even if the restore job fails. You cannot restore a virtual machine if it already exists on the virtual server unless you select this option.
Power on virtual machine after restore	Turns on a virtual machine after the restore job completes.

Table M-7 VMware restore job options (*continued*)

Item	Description
Transport mode priority list	<p>Lets you select the method to transport the Virtual Machine Disk Format (VMDK) file from the ESX server. You must select at least one of these options. If you select more than one option, the method is determined by the priority and the resources that are available. You can move the options up or down in the list to change the priority.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> ■ SAN-Use the SAN to move virtual disk data If you select this option, the virtual machine must reside on a SAN that the media server can access. With this transport mode, the data is offloaded to the media server so that the ESX server is not affected. ■ NBD-Do not encrypt the virtual disk data for over-the-network transfers Use this option if you do not use SSL for security and one of the following conditions exist: <ul style="list-style-type: none"> ■ The virtual machine is not located on the SAN. ■ The media server does not have access to the SAN. ■ NBDSSL-Encrypt virtual disk data for over-the-network transfers Use this option if you use SSL for security and one of the following conditions exist: <ul style="list-style-type: none"> ■ The virtual machine is not located on the SAN. ■ The media server does not have access to the SAN. ■ Hotadd-Use virtual disk files from the Backup Exec server on the virtual machine Use this option if you want to use the hot add feature for ESX. The hot add feature lets you use a virtual machine as your proxy server. See your VMware documentation for more information about hot add. <p>The snapshot is placed on the ESX server. Then, the data moves across the network.</p> <p>The VMDK file is not backed up if the virtual hard disk is configured as an Independent disk.</p>
Move Up	Lets you move the selected transport mode to a higher priority in the list.

Table M-7 VMware restore job options (*continued*)

Item	Description
Move Down	Lets you move the selected transport mode to a lower priority in the list.
Enter a path to store the temporary files required to restore legacy backup sets	Indicates the location where you want to store temporary files when you restore data from backup sets that were created with a previous version of Backup Exec.
vSphere port number	Indicates the port that Backup Exec uses to connect to vCenter Server. The default port is 902.

Redirecting the restore of a VMware virtual machine

By default, Backup Exec restores data to the location from which it was originally backed up. If you want to restore data to a different virtual machine than where the data originally resided, you must create a redirected restore job.

To redirect the restore of a VMware virtual machine

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Select the data to be restored.
- 4 In the **Properties** pane, under **Destination**, click **VMware Redirection**.
- 5 Select the appropriate options.
See “[VMware Redirection options](#)” on page 1404.
- 6 Start the job or select other options from the **Properties** pane.

VMware Redirection options

The following options are available for VMware restore redirection jobs.
 See “[Redirecting the restore of a VMware virtual machine](#)” on page 1404.

Table M-8 VMware Redirection options

Item	Description
Redirect VMware sets	Lets you set options to redirect data to the vCenter or ESX server.

Table M-8 VMware Redirection options (continued)

Item	Description
Server logon account	Uses the default logon account that appears. You can select another logon account to use for the vCenter or ESX server to which you want to redirect data.
Change	Lets you select a different logon account to use for the vCenter or ESX server to which you want to redirect data.
Redirect to a different vCenter or ESX Server	Lets you set options to redirect VMware data to a different server.
vCenter and ESX Servers	Provides the name of the vCenter or ESX server to which you want to redirect data.
Browse vCenter and ESX servers for destination	Lets you select the virtual server on which you want to redirect data. You can use this option instead of typing the name of the server.
Data Center	Displays the name of the datacenter, or group of ESX servers.
Virtual machine Datastore	Displays the name of the storage location on the ESX server that is used to store the data.
Host or Cluster	Displays the name of the ESX server that will run the virtual machine after the restore job completes.
Virtual machine folder	Indicates the name of the existing vSphere folder to which you want to restore.
Resource pool	Indicates the name of the resource pool to which you want to restore.
New virtual machine name	Indicates the new virtual machine name. You may want to provide a new virtual machine name if a virtual machine with the same name already exists on the server.
Use the original disk datastore selections if available on the selected host	Uses the original datastore selections on the virtual server. If the original datastore selections do not exist, then the datastore selections from the backup data are used.

Table M-8 VMware Redirection options *(continued)*

Item	Description
Choose a network	Indicates the network for the virtual machine to use after the restore job completes.
Restore virtual machine with VMware virtual hardware version 7	Restores the virtual machine with VMware hardware version 7. Selecting this option causes jobs to fail when you restore to VMware ESX Server version 3.5.
Restore with thin provisioning	Restores the virtual machine with thin provisioning. Thin provisioning can help you more efficiently dedicate storage capacity in your VMware ESX Server version 4.0 environment. Selecting this option causes jobs to fail when you restore to VMware ESX Server version 3.5.
Redirect to a folder	Lets you restore data to a folder without restoring it to the ESX server. When the restore job completes, the folder contains all of the .vmdk files for the virtual machine.
Restore to drive	Indicates the drive where the folder is located.
Restore to path	Indicates the path where the folder is located.

Setting default backup and restore options for the Agent for VMware

You can use the defaults that Backup Exec sets during installation for all VMware backup and restore jobs, or you can choose your own defaults. You can also set backup or restore options for individual jobs.

To set default backup and restore options for the Agent for VMware

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Virtual Machines**, click **VMware**.

- 3 Select the appropriate options.
See [“VMware default options”](#) on page 1407.
- 4 Click **OK**.

VMware default options

You can change the following default options for all VMware backup and restore jobs.

See [“Setting default backup and restore options for the Agent for VMware”](#) on page 1406.

Table M-9 VMware default options

Item	Description
Transport mode priority list	<p>Lets you select the method to transport the Virtual Machine Disk Format (VMDK) file from the ESX server. You must select at least one of these options. If you select more than one option, the method is determined by the priority and the resources that are available. You can move the options up or down in the list to change the priority.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> <p>■ SAN - Use the SAN to move virtual disk data</p> <p>If you select this option, the virtual machine must reside on a SAN that the media server can access. With this transport mode, the data is offloaded to the media server, so that the ESX server is not affected.</p> <p>■ NBD - Do not encrypt the virtual disk data for over-the-network transfers</p> <p>Use this option if you do not use SSL for security and one of the following conditions exist:</p> <ul style="list-style-type: none"> ■ The virtual machine is not located on the SAN ■ The media server does not have access to the SAN. <p>The data is placed on the ESX server. Then, the data moves across the network.</p> <p>■ NBDSSL - Encrypt virtual disk data for over-the-network transfers</p> <p>Use this option if you use SSL for security and one of the following conditions exist:</p> <ul style="list-style-type: none"> ■ The virtual machine is not located on the SAN. ■ The media server does not have access to the SAN. <p>■ Hotadd - Use virtual disk files from the Backup Exec server on the virtual machine</p> <p>Use this option if you want to use the hot add feature for ESX. The hot add feature lets you use a virtual machine as your proxy server.</p> <p>See your VMware documentation for more information about hot add.</p> <p>The snapshot is placed on the ESX server. Then, the data moves across the network.</p> <p>The VMDK file is not backed up if the virtual hard disk is configured as an Independent disk.</p>

Table M-9 VMware default options (*continued*)

Item	Description
Move Up	Lets you move the selected transport mode up in the priority list.
Move Down	Lets you move the selected transport mode down in the priority list.
Back up virtual machines that are powered off	Enables Backup Exec to back up virtual machines that are turned off. Note: If a virtual machine is automatically discovered, it is not backed up if it is turned off.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	Enables individual files and folders to be restored from virtual machines. This option is for virtual machines that use a Windows operating system only. GRT is not meant for system recovery but only for the restore of individual files and folders on Windows computers.
Edit	Lets you change the GRT settings for Microsoft Active Directory, Exchange, and SQL.
Microsoft Active Directory	Indicates whether GRT is enabled or disabled for Microsoft Active Directory on the virtual machine. It is enabled by default.
Microsoft Exchange	Indicates whether GRT is enabled or disabled for Microsoft Exchange on the virtual machine. It is enabled by default.
Microsoft SQL	Indicates whether GRT is enabled or disabled for Microsoft SQL on the virtual machine. It is enabled by default.
vSphere Port Number	Indicates the port that Backup Exec uses to connect to vCenter Server. The default port is 902.
Delete existing virtual machines prior to restore	Deletes virtual machines during the restore job. If you select this option, the virtual machines may be deleted even if the restore job fails. You cannot restore a virtual machine if it already exists on the virtual server.
Power on virtual machine after restore	Turns on a virtual machine after the restore job completes.

Table M-9 VMware default options (*continued*)

Item	Description
Transport mode priority list	<p>Lets you select the method to transport the Virtual Machine Disk Format (VMDK) file from the ESX server. You must select at least one of these options. If you select more than one option, the method is determined by the priority and the resources that are available. You can move the options up or down in the list to change the priority.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> ■ SAN - Use the SAN to move virtual disk data If you select this option, the virtual machine must reside on a SAN that the media server can access. With this transport mode, the data is offloaded to the media server, so that the ESX server is not affected. ■ NBD - Do not encrypt the virtual disk data for over-the-network transfers Use this option if you do not use SSL for security and one of the following conditions exist: <ul style="list-style-type: none"> ■ The virtual machine is not located on the SAN. ■ The media server does not have access to the SAN. ■ NBDSSL - Encrypt virtual disk data for over-the-network transfers Use this option if you use SSL for security and one of the following conditions exist: <ul style="list-style-type: none"> ■ The virtual machine is not located on the SAN. ■ The media server does not have access to the SAN. ■ Hotadd - Use virtual disk files from the Backup Exec server on the virtual machine Use this option if you want to use the hot add feature for ESX. The hot add feature lets you use a virtual machine as your proxy server. See your VMware documentation for more information about hot add. <p>The snapshot is placed on the ESX server. Then, the data moves across the network.</p> <p>The VMDK file is not backed up if the virtual hard disk is configured as an Independent disk.</p>
Enter a path to store the temporary files created by staging	Indicates the location where you want to store temporary files when you restore data from backup sets that were created with a previous version of Backup Exec.

Table M-9 VMware default options (*continued*)

Item	Description
vSphere port number	Indicates the port that Backup Exec uses to connect to vCenter Server. The default port is 902.

Symantec Backup Exec Archiving Option

This appendix includes the following topics:

- [About the Archiving Option](#)
- [Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option](#)
- [Installing the Backup Exec Archiving Option](#)
- [How the Archiving Option works](#)
- [Best practices for the Archiving Option](#)
- [About creating an Archiving Option archive job](#)
- [About vault stores in the Archiving Option](#)
- [About vault store partitions in the Archiving Option](#)
- [About archives in the Archiving Option](#)
- [About archive settings in the Archiving Option](#)
- [About Exchange mailbox groups in archive jobs](#)
- [About searching for data in the archives](#)
- [About restoring items from the archives](#)
- [About deleting items from the archives](#)
- [About backing up Archiving Option components](#)
- [About restoring an Archiving Option component](#)

- [About backing up and restoring the Archiving Option components from a remote media server](#)
- [Preventing the deletion of expired archived items from an archive](#)
- [About synchronizing archive permissions and settings](#)
- [About single instance storage of archived items](#)
- [Editing default settings for archive jobs](#)
- [About moving Archiving Option components to a new location](#)
- [Troubleshooting archive jobs](#)
- [Reports for the Archiving Option](#)

About the Archiving Option

The Archiving Option includes the following features that you can install separately or together:

- The Backup Exec File System Archiving Option, which archives the eligible Windows file system data.
- The Backup Exec Exchange Mailbox Archiving Option, which archives the eligible Exchange mail messages.

To find the data that is eligible for archiving, Backup Exec applies rules to the selected file system shares and folders and to the Exchange mailboxes. Data in the selections is eligible for archiving if it is backed up and meets the criteria that the rules specify. The archive job then sends the data to disk-based vault stores. The data is deleted from its original location on the resource immediately after it is archived, or after you back up the vault store.

You can apply retention categories to the data that is archived that specify how long to keep data in the archives. Backup Exec can automatically delete the archived data that has expired retention dates.

If you install and configure Backup Exec Retrieve, end users can access their own archived files and mail messages. End users click a link that Backup Exec creates when it archives mail messages or files. The link opens Backup Exec Retrieve where end users can search, browse, preview, and retrieve only the mail messages or files that they own.

By archiving data from the backup sets, Backup Exec eliminates additional querying and movement of data on the resources. After Backup Exec deletes the archived data from its original location, you have more disk space, and Backup Exec requires less time for future backup jobs.

The Archiving Option uses Symantec Enterprise Vault technology to archive data. When you install the Archiving Option, some Enterprise Vault services are also installed.

See “[About Enterprise Vault services for the Archiving Option](#)” on page 1422.

See “[Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option](#)” on page 1415.

See “[Installing the Backup Exec Archiving Option](#)” on page 1428.

See “[How the Archiving Option works](#)” on page 1429.

See “[Best practices for the Archiving Option](#)” on page 1433.

Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option

Some requirements apply to both the Backup Exec Exchange Mailbox Archiving Option and the Backup Exec File System Archiving Option. You should also review the requirements that apply to each specific option.

See “[Requirements for the Backup Exec Exchange Mailbox Archiving Option](#)” on page 1416.

See “[Requirements for the Backup Exec File System Archiving Option](#)” on page 1418.

Following are the requirements that apply to both the Backup Exec Exchange Mailbox Archiving Option and the Backup Exec File System Archiving Option (Archiving Option):

- The media server must be part of a domain. You cannot install the Archiving Option on a server in a workgroup.
- The media server must have enough space to store the Archiving Option index files. When you install the Archiving Option, you are prompted to provide a path where the index files are stored. The path must be on a local NTFS volume. Paths on FAT and FAT32 partitions are not supported.

Note: Symantec recommends that you have more RAM available in addition to Backup Exec's base requirements.

See “[System requirements](#)” on page 114.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

The Archiving Option does not support the following:

- The Backup Exec Central Admin Server Option.

Note: You can install the Archiving Option on a central administration server. However, distributed job management for archiving jobs is not supported.

- Archiving from backup sets from the Backup Exec Remote Media Agent for Linux Servers.
- Installation on clustered servers. Additionally, you cannot install Backup Exec on a cluster if you have also selected the Archiving Option for installation.
- Archiving from legacy mailbox backup sets.
- Installation of the Exchange Mailbox Archiving Option on a computer on which Microsoft Exchange Server is installed.

See [“How to calculate disk space requirements for the Exchange Mailbox Archiving Option”](#) on page 1423.

See [“How to calculate disk space requirements for the File System Archiving Option”](#) on page 1425.

See [“About Enterprise Vault services for the Archiving Option”](#) on page 1422.

Requirements for the Backup Exec Exchange Mailbox Archiving Option

Some requirements that apply to both the Backup Exec Exchange Mailbox Archiving Option and the Backup Exec File System Archiving Option are listed separately. Review those requirements in addition to the requirements that apply only to the Exchange Mailbox Archiving Option.

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1415.

Following are the requirements that apply to the Exchange Mailbox Archiving Option:

- You must enter an Exchange Agent license key on the media server for each Exchange Server that you want to archive.
- You must install one of the following versions of Microsoft Outlook on the media server before you install the Archiving Option:
 - Outlook 2007 Service Pack 2 with hotfix kb968858 or later
 - Outlook 2003 Service Pack 3
- Microsoft Outlook must be the default mail client on the media server.

Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option

- You must create a profile when you install Outlook on the media server, and then connect to an Exchange Server mailbox. Outlook may display an error message about a conflicting program. If Outlook offers to fix the problem, choose to do so, and then follow the instructions that are given.
- The Exchange Server backups must have the Granular Recovery Technology (GRT) option enabled.
- The Exchange Server backups must be on one of the following devices:
 - A non-removable backup-to-disk folder.
 - A deduplication storage folder.
 - A storage array in a Storage Provisioning Option environment.
- Archive jobs must have a valid path configured on an NTFS volume that is local to the media server for temporary storage of data. The default path is set to use C:\temp.
See [“Setting defaults for restore jobs”](#) on page 638.
- A mailbox must be configured for exclusive use by Backup Exec on each Exchange Server on which you want to select mailboxes for archiving. Whenever you create an archive job for the Exchange Mailbox Archiving Option you are prompted to enter the name of the system mailbox. The system mailbox is the mailbox that you configure for use by Backup Exec. It does not need to be named 'system' mailbox.
- The following are restrictions for this mailbox:
 - The mailbox must not be used for any other purpose. The Exchange Mailbox Archiving Option requires exclusive access.
 - The mailbox must not be hidden from address lists.
 - The mailbox account must not be disabled.
- The media server domain, and the Exchange Server domains must trust the domain that the Backup Exec service account belongs to.
- You must grant permissions to the Backup Exec service account to the Exchange servers.
See [“About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option ”](#) on page 1419.
- The Backup Exec service account must be a member of the Active Directory domain. Symantec recommends that you use a Backup Exec service account that has domain and local administrator rights on the Exchange Server. You should not make the Backup Exec service account a domain administrator.

Requirements for the Backup Exec File System Archiving Option

Some requirements that apply to both the Backup Exec Exchange Mailbox Archiving Option and the Backup Exec File System Archiving Option are listed separately. Review those requirements in addition to the requirements that apply only to the File System Archiving Option.

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1415.

The following are requirements for the File System Archiving Option:

- The media server must be in the same time zone as the file servers from which eligible data is archived.
- The media server domain, the file server domains, and the Exchange Server domains must trust the domain that the Backup Exec service account belongs to.
- The media server domain must trust the domains that contain the accounts of users that access the file server shares.
- If Backup Exec Retrieve is installed, a trust relationship must also exist for that domain.
- The Backup Exec service account must have local administrative rights on the file server.
- The Backup Exec service account must have Full Control share permissions on the share that is selected for archiving.
- The Backup Exec service account must be granted the following NTFS rights on the folders in the share that is selected for archiving:
 - **Modify**
 - **List Folder Contents**
 - **Read**
 - **Write**
- You must install one of the following Outlook versions on the media server to provide full indexing of .msg files:
 - Outlook 2003 Service Pack 3
 - Outlook 2007 Service Pack 2 with hotfix kb968858 or later

About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option

For the Archiving Option, the Backup Exec service account must access mailboxes on the Exchange Servers that you want to archive. To gain this access, the Backup Exec service account must have permission to access the Exchange Servers.

You can use either of the following methods to grant the permissions that the Backup Exec service account needs to access the mailboxes on Exchange Servers:

- Grant permissions at the organization level or at the Administrative Group level.

Permissions are then propagated automatically to any new Exchange Servers that you add under the level at which the permissions are assigned.

Note: You must have Exchange administrative permissions to grant permissions to other accounts.

- Grant permissions explicitly on each Exchange Server.

If you grant permissions explicitly and then add another Exchange Server, you must grant permissions explicitly on the added server as well.

The Backup Exec service account must also have the Send As permission on the mailbox that you create for Backup Exec's exclusive use. This mailbox, called the system mailbox, must be created on each Exchange Server on which you want to select mailboxes for archiving.

See [“Granting permissions at the Organization level for Exchange Server 2007 for the Archiving Option”](#) on page 1419.

See [“Granting permissions explicitly on each Exchange Server 2007 for the Archiving Option”](#) on page 1420.

See [“Granting permissions at the Organization level for Exchange Server 2003 for the Archiving Option”](#) on page 1421.

See [“Granting permissions at the server level for Exchange Server 2003 for the Archiving Option”](#) on page 1422.

Granting permissions at the Organization level for Exchange Server 2007 for the Archiving Option

You can grant the **Full Control** permission for the Backup Exec service account at the Organization level.

See [“About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option ”](#) on page 1419.

Note: You must have Exchange administrative permissions to grant permissions to other accounts.

To grant permissions at the Organization level for the Backup Exec service account in the Archiving Option

1 On the Exchange Server, click **Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Shell**.

2 Type the following command:

```
Get-OrganizationConfig | Add-ADPermission -User '<Domain Name\Backup Exec service account>' -AccessRights GenericAll -InheritanceType All
```

3 Type the following command:

```
Add-ADPermission -Identity '<system mailbox name>' -User '<Domain Name\Backup Exec service account>' -ExtendedRights 'Send-as'
```

4 To grant **Send As** permission on the mailboxes that you created for Backup Exec's exclusive use, repeat the previous step on the appropriate Exchange Servers.

Granting permissions explicitly on each Exchange Server 2007 for the Archiving Option

You can grant the **Full Control** permission for the Backup Exec service account on each Exchange Server. Perform this procedure on each Exchange Server that you want to archive.

See [“About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option ”](#) on page 1419.

Note: You must have Exchange administrative permissions to grant permissions to other accounts.

To grant permissions explicitly on each Exchange Server 2007 for the Archiving Option

- 1 On the Exchange Server, click **Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Shell**.

- 2 Type the following command:

```
Get-MailboxServer -Identity "<mailbox server name>" > |  
Add-ADPermission -User "<Domain name\Backup Exe service account>"  
-AccessRights GenericAll -InheritanceType All
```

- 3 Type the following command:

```
Add-ADPermission -Identity '<system mailbox name>' -User '<Domain  
Name\Backup Exec service account>' -ExtendedRights 'Send-as'
```

Granting permissions at the Organization level for Exchange Server 2003 for the Archiving Option

You can grant the **Full Control** permission for the Backup Exec service account at the Organization level.

See [“About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option”](#) on page 1419.

Refer to the Microsoft knowledge base for more information on this procedure.

To grant permissions at the Organization level for Exchange Server 2003 for the Archiving Option

- 1 Configure the **ShowSecurityPage** registry setting to enable the display of the **Security** page.
- 2 In the left pane of Microsoft Exchange System Manager, right-click the **Exchange Organization**, and then click **Properties**.
- 3 On the **Security** tab, click **Add**.
- 4 Select the Backup Exec service account to add it to the list.
- 5 Click **OK**.
- 6 In the **Name** list, select the Backup Exec service account.
- 7 In the **Permissions** list, ensure that all of the check boxes in the **Allow** column are selected.
- 8 Select any check boxes that are not selected.
- 9 Click **OK**.

Granting permissions at the server level for Exchange Server 2003 for the Archiving Option

You can grant permissions for the Backup Exec service account at the server level for Exchange Server 2003. Perform this procedure on each Exchange Server that you want to archive.

See [“About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option ”](#) on page 1419.

To grant permissions at the server level for Exchange Server 2003 for the Archiving Option

- 1 In the left pane of the **Microsoft Exchange System Manager**, expand the **Servers** container.
- 2 Right-click the Exchange Server, and then click **Properties**.
- 3 On the **Security** tab, click **Add**.
- 4 Select the Backup Exec service account to add it to the list.
- 5 Click **OK**.
- 6 In the **Name** list, click the Backup Exec service account.
- 7 In the **Permissions** list, ensure that all of the check boxes in the **Allow** column are selected.
- 8 Select any check boxes that are not selected.
- 9 Click **OK**.

About Enterprise Vault services for the Archiving Option

Symantec Enterprise Vault technology is the foundation for the Archiving Option. When you install the Archiving Option, some Enterprise Vault services are also installed. The Enterprise Vault services that run on the media server use the same credentials as the Backup Exec service account.

The following Enterprise Vault services are installed on the media server:

- Enterprise Vault Admin Service
- Enterprise Vault Directory Service
- Enterprise Vault Indexing Service
- Enterprise Vault Storage Service
- Enterprise Vault Task Controller Service

You must always use the Backup Exec Services Manager on the media server to update your Backup Exec credentials. The Backup Exec Service Manager

automatically updates the Enterprise Vault service credentials with the same credentials.

Note: Use of the Windows Services applet to edit the credentials of an Enterprise Vault service or a Backup Exec service is not supported. The use of this applet can leave the Archiving Option unsynchronized with the Backup Exec service account credentials. Errors can occur during the archiving operations.

See [“Changing service account information”](#) on page 107.

See [“Starting and stopping Backup Exec services”](#) on page 168.

How to calculate disk space requirements for the Exchange Mailbox Archiving Option

Backup Exec requires permanent disk space for the following Exchange Mailbox Archiving Option components:

- The vault store partitions.
- The index locations.
- The SQL Server database, which contains the Directory, vault store, and fingerprint databases.

[Table N-1](#) describes the formulas that you can use to estimate the disk space requirements for these components for the Exchange Mailbox Archiving Option.

The following values and variables are used in the formulas:

- N is the number of emails.
- m is the average number of identical copies of attachments across user mailboxes.
- The compression factor for attachments is estimated as 60%.
If the attachments are mostly Office 2007 files, the compression factor to use is 90%.
- The average number of emails that have attachments is estimated at 20%.
- The average size of an email attachment is estimated at 250 KB.

Table N-1 Calculations for disk space requirements for the Exchange Mailbox Archiving Option

Component	Requirements
Vault store partitions	<p>The size of a vault store partition depends on the following items:</p> <ul style="list-style-type: none"> ■ The size of the emails. ■ The type of attachments. ■ The number and size of the attachments. ■ The number of emails with attachments. <p>Note: If single instance storage is enabled, items are shared within and across vault stores and vault store partitions. Shareable parts of a message that exceed the single-instance threshold of 20 KB are shared. These parts include attachments and message bodies. User information and shareable parts under the single instance threshold are not shared.</p> <p>See “About single instance storage of archived items” on page 1494.</p> <p>You can use the following calculations to approximate the disk space requirements of a vault store partition:</p> <ul style="list-style-type: none"> ■ Approximate vault store partition size for which single instance storage is not enabled: $(N \times 16) + (N \times 0.2 \times 0.6 \times 250)$ kilobytes ■ Approximate vault store partition size for which single instance storage is enabled: $(N \times 16) + ((1/m) \times (N \times 0.2 \times 0.6 \times 250))$ kilobytes <p>For example, you want to know the disk space requirements for a vault store partition for 100,000 emails. You estimate that each email attachment is shared across three people on average.</p> <p>If single instance storage of archived items is not enabled, the calculation for the disk space requirements is as follows: $(100000 \times 16) + (100000 \times 0.2 \times 0.6 \times 250)$ kilobytes = 4.6 GB approximately</p> <p>If single instance storage is enabled, the calculation for the disk space requirements is as follows: $(100000 \times 16) + ((1/3) \times 100000 \times 0.2 \times 0.6 \times 250)$ kilobytes = 2.6 GB approximately</p>

Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option

Table N-1 Calculations for disk space requirements for the Exchange Mailbox Archiving Option (*continued*)

Component	Requirements
Indexes	<p>The size of an index is approximately 8% of the total size of the items that are archived. The percentage may be less if there is less content to index. For example, there is less content to index when there are large attachments such as MP3 or .jpeg files.</p> <p>For example, you have 100,000 emails that each have a body size of 8 KB. About 20% of the emails have attachments, each with an average total size of 250 KB. The index size is approximately 450 MB.</p>
Directory database	<p>The Directory database only grows when a new mailbox or share is archived for the first time.</p> <p>The recommended disk space is 500 MB.</p>
Vault store database	<p>The size of a vault store database is approximately $N \times 500$ bytes. The vault store database grows with every item that is archived. Temporary space is used to hold information on the items that have not been backed up or indexed.</p>
Fingerprint database	<p>The fingerprint database is created only if you enable single instance storage of archived items. Backup Exec initially allocates 212 MB for the fingerprint database. The fingerprint database grows with every item that is archived.</p> <p>If the database grows to more than 212 MB, use the following calculation to estimate the disk space that it requires:</p> $1/m \times N \times 0.2 \times 500 \text{ bytes}$ <p>See "About single instance storage of archived items" on page 1494.</p>

How to calculate disk space requirements for the File System Archiving Option

Backup Exec requires permanent disk space for the following File System Archiving Option components:

- Vault store partitions.
- Indexes
- SQL Server database

[Table N-2](#) describes the formulas that you can use to estimate the disk space requirements for these components for the File System Archiving Option.

The following values and variables are used in the formulas:

- N is the number of files.
- m is the average number of identical copies per file.
If m is unknown, use the estimate of 1.2.
- The compression factor for files is estimated as 50%.
This estimate applies to a mix of files that contain mostly Office 2003 documents. Office 2007 documents do not compress but when mixed with non-Office files, the compression average is 80% of the original size. Pure image files are not compressed.

Table N-2 Calculating disk space requirements for the File System Archiving Option components

Component	Disk space requirements
Vault store partition	<p>You can use the following calculations to approximate the disk space requirements of a vault store partition:</p> <ul style="list-style-type: none"> ■ Approximate vault store partition size for which single instance storage of archived items is not enabled: $(N \times 4) + (N \times \text{average file size in kilobytes} \times 0.5)$ kilobytes ■ Approximate vault store partition size for which single instance storage is enabled: $(N \times 4) + ((1/m) \times N \times \text{average file size} \times 0.5)$ kilobytes <p>For example, you want to know the disk space requirement for a vault store partition for 10,000 files. The average size of each file is 250 KB, and the average number of identical copies per file is 1.2.</p> <p>If single instance storage of archived items is not enabled, then the calculation for the disk space requirement is as follows:</p> $(10000 \times 4) + (100000 \times 250 \times 0.5) \text{ kilobytes} = 1.3 \text{ GB approximately}$ <p>If single instance storage is enabled, then the calculation for the disk space requirement is as follows:</p> $(10000 \times 4) + ((1/1.2) \times 10000 \times 250 \times 0.5) \text{ kilobytes} = 1.08 \text{ GB approximately}$

Table N-2 Calculating disk space requirements for the File System Archiving Option components (*continued*)

Component	Disk space requirements
Indexes	<p>You can estimate that the index files require approximately 2% of the total size of the files that are archived. The percentage may be less if there is less content to index. If the files are all compressed image file, indexing is less than if the files are mostly small text messages. Numerous small text messages require disk space requirements similar to those of the Exchange Mailbox Archiving Option for indexing</p> <p>For example, to archive 10 GB of data, at least 200 MB of available disk space is required to store the index files.</p>
Directory database	<p>The Directory database only grows when a new mailbox or share is archived for the first time.</p> <p>The recommended disk space is 1 GB.</p>
Vault store database	<p>The vault store database grows with every item that is archived. Temporary space is used to hold information on the items that have not been backed up or indexed.</p> <p>The size of a vault store database is approximately $N \times 3000$ bytes.</p>
Fingerprint database	<p>The fingerprint database is created only if you enable single instance storage of archived items. The fingerprint database holds the shareable parts of archived items. Shareable parts of an item that exceed the single-instance threshold of 20 KB are shared. For the File System Archiving Option, it is expected that all files are larger than the 20-KB threshold.</p> <p>Backup Exec initially allocates 212 MB for the fingerprint database. The fingerprint database grows with every item that is archived.</p> <p>If the database grows to more than 212 MB, use the following calculation to estimate the disk space that it requires:</p> <p>$1/m \times N \times 500$ bytes</p> <p>See “About single instance storage of archived items” on page 1494.</p>

Installing the Backup Exec Archiving Option

You can install one or both of the following options locally as a separate, add-on component of Backup Exec:

- The Exchange Mailbox Archiving Option
- The File System Archiving Option

Command line switches are also available for silent mode installation of these options.

Before you attempt to install the Archiving Option, verify that all requirements are met.

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 153.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1415.

See [“About Enterprise Vault services for the Archiving Option”](#) on page 1422.

See [“About installing Enterprise Vault on a media server on which the Archiving Option is installed”](#) on page 1429.

About uninstalling or reinstalling the Archiving Option

If you uninstall both the Exchange Mailbox Archiving Option and the File System Archiving Option, the Enterprise Vault files and the Enterprise Vault services that are included in the Archiving Options are removed. All archive-related jobs display a status of **Disabled**. You cannot run, edit, or save a disabled job. You can delete a disabled job.

If you uninstall only one option, then no changes occur for the existing archive jobs. You can continue to edit and run the archive jobs as usual.

If you reinstall one or both options, all previously archived data is available if you specify the same Backup Exec installation folder path that you used in the initial installation. You can rerun any disabled jobs if no changes were made to the Backup Exec Database. Otherwise, if you try to run the jobs again, the jobs fail.

If you want to reinstall the Archiving Option and use new databases, you must reinstall both options. However, all previously archived data becomes unavailable. Archiving selections, rules, and archiving configurations are also unavailable.

You must remove the following items before you can reinstall the Archiving Option with new databases:

- The Enterprise Vault databases that are included in the Archiving Option. If SQL Express is used, these databases are located by default at \Program Files\Symantec\Backup Exec\Data. If the full version of SQL Server is used, the archive databases are in the path that the administrator defined.
- The BEAODatabase.xml file, which is located at \Program Files\Symantec\Backup Exec.
- The index files, which are located by default at \Program Files\Symantec\Backup Exec\ArchiveIndex.

If you reinstall one or both options and you use a different Backup Exec service account than you used for the initial installation, then you must use the Backup Exec Services Manager to update the service account after the reinstallation is complete.

See [“Uninstalling Backup Exec options from the local media server”](#) on page 169.

About installing Enterprise Vault on a media server on which the Archiving Option is installed

If you install Enterprise Vault on a media server on which the Archiving Option is installed, all archiving functionality is unavailable. Archive jobs that are active when Enterprise Vault is installed run to completion, but scheduled archive jobs do not run.

All archive-related jobs display a status of **Disabled**. You cannot run, edit, or save a disabled job. You can delete a disabled job.

If you subsequently uninstall Enterprise Vault, archiving functionality remains unavailable.

How the Archiving Option works

To process an archive job, Backup Exec performs the following actions:

- Reads the latest backups of the file systems and Exchange servers from which you make archive selections.
- Applies the archive rules that you specify to identify the files and mail messages that are eligible for archiving.
- Checks if eligible files already exist in the archives.
If a file already exists in the archives, it is not archived again.
- Adds the data to the archives.
All of the archived data content is indexed to enable fast searching and retrieval of archived items.

- Deletes the archived files and mail messages from their original location. Depending on an option that you specify, deletion occurs immediately after the archive job completes or after the vault store is backed up.
- (Optional) Creates links to Backup Exec Retrieve in the folders that belong to the end user. You must install and configure Backup Exec Retrieve before links can be created.
 See [“How Archiving Option end users retrieve archived data by using Backup Exec Retrieve”](#) on page 1432.

The Archiving Option operations that you can perform are described in the following table:

Table N-3 Operations that you can run

Operation	More information
Create archive jobs to archive file system data and Exchange mail messages to vault stores.	See “About creating an Archiving Option archive job” on page 1434.
Create disk-based vault stores to use as storage devices for the archived data.	See “About vault stores in the Archiving Option” on page 1445.
Restore individual items from the archives.	See “About restoring items from the archives” on page 1468.
Delete individual items from the archives.	See “About deleting items from the archives” on page 1474.
Delete expired archive data from an archive automatically to free disk space, or ensure that archive data is never deleted from an archive.	See “Preventing the deletion of expired archived items from an archive” on page 1493.
Back up the Archiving Option components. These components include vault stores, vault store partitions, archives, databases, and index locations.	See “About backing up Archiving Option components” on page 1477.
Restore the Archiving Option components. These components include vault stores, vault store partitions, archives, databases, and index locations.	See “Restoring an Archiving Option component” on page 1483.
Synchronize the archive permissions with mailbox permissions, and share and folder permissions.	See “About synchronizing archive permissions and settings” on page 1493.

Table N-3 Operations that you can run (*continued*)

Operation	More information
Install Backup Exec Retrieve so that end users can recover archived files on their own.	See “About Backup Exec Retrieve” on page 861. See “How Archiving Option end users retrieve archived data by using Backup Exec Retrieve” on page 1432.

See [“Best practices for the Archiving Option”](#) on page 1433.

Types of data not included in Archiving Option archive jobs

The Archiving Option does not include some types of data in archive jobs.

Table N-4 The types of data that are not included in archive jobs

Archiving Option	Types of data
File System Archiving Option	<p>The following types of data are not included in a file system archive job:</p> <ul style="list-style-type: none"> ■ Hard links ■ Files with alternate streams ■ Reparse points ■ Sparse files ■ Files in Microsoft Distributed File System Replication (DFSR) shares, system folders, or recycle bins ■ Files that have an encrypted, hidden, or system attribute ■ Files that are in mount point directories You can share the root of the mount point target, and then select it for archiving.
Exchange Mailbox Archiving Option	<p>The following types of data are not included in an Exchange mailbox archive job:</p> <ul style="list-style-type: none"> ■ Mail messages that have pending reminders. ■ Any Exchange items other than mail messages, such as address book entries and calendar items. ■ Mail messages in Exchange managed folders, journal mailboxes, or in public folders.

About Archiving Option operation entries in the audit log

Audit logs provide information about the operations that have been performed in Backup Exec.

You can view information about archiving operations for the following:

- Vault stores
- Vault store partitions
- Archive settings
- Retention categories

See [“About audit logs”](#) on page 205.

How Archiving Option end users retrieve archived data by using Backup Exec Retrieve

End users can retrieve archived items by using Backup Exec Retrieve. An online Help system is provided with Backup Exec Retrieve.

End users can do the following from Backup Exec Retrieve:

- Search for archived items.
- View recently archived items.
- Retrieve an archived item.

After you install and configure Backup Exec Retrieve, Backup Exec creates a link from the archived folders or mailboxes to a Backup Exec Retrieve URL.

The link to the Backup Exec Retrieve URL is displayed to the end user as described in the following table:

Table N-5 Where links to Backup Exec Retrieve are displayed to end users

Option	Backup Exec Retrieve Links
Backup Exec File System Archiving Option	A link to Backup Exec Retrieve displays in every folder from which a file is archived.
Backup Exec Exchange Mailbox Archiving Option	A link to Backup Exec Retrieve displays in each mailbox from which a mail message is archived.

If you disable Backup Exec Retrieve, all of the existing links in the archived folders and mailboxes are removed.

See [“Application settings for Backup Exec Retrieve”](#) on page 874.

Access control for end users is based on the following:

- The share permissions and file system permissions for file system data.
- Mailbox and folder permissions for Exchange mailboxes.

See [“About Backup Exec Retrieve”](#) on page 861.

Best practices for the Archiving Option

Following are best practices when you use the Archiving Option:

- Use the default full recovery model for the SQL Server instance that hosts the Backup Exec Database and the Archiving Option databases. All Archiving Option databases that are created on the SQL Server are then also created with the full recovery model.
- Create only one archive job for each server from which you want to archive the data that has been backed up.
- Do not use different media servers to archive files or mailboxes from the same file server or Exchange server.
- Configure backup jobs so that full backups and their associated incremental and differential backups use the same selection list.
- Run archive jobs outside the backup window. That is, don't run archive jobs at the same time that you run backup jobs.
- Consider archiving a smaller amount of data at first, such as a mailbox or a folder. All backup data may be eligible when you run the first archive job. Over a period of time, the amount of eligible archive data lessens, and it becomes a predictable amount.
- Select the file system shares to archive that end users have access to rather than administrative shares. End users can then retrieve their own data by using Backup Exec Retrieve.
- Ensure that a selection is included in only one archive job. Unlike backup jobs, archive jobs cannot share the same selections.
- Ensure that all subdirectories in a selection are included in only one archive job.
- Do not archive the system drive. The Archiving Option does not archive system files.
- If you restore multiple Archiving Option components that include the Directory database, use a separate job to restore the Directory database first. Then, create

one job for all of the remaining Archiving Option components that you want to restore.

- If you restore multiple backup sets to restore a database, use a single restore job and leave the database ready to use.
- If you redirect the restore of Archiving Option components to a new server because of hardware failure, redirect the restore of the Directory database first. Create a separate job to redirect the restore of the database. After the redirected restore of the Directory database is complete, you must run some additional tasks in a separate program called Backup Exec Utility. The tasks in the Backup Exec Utility update the Directory database with the new locations of the components. You should run the Backup Exec Utility tasks before you redirect the restore of any other Archiving Option components.

About creating an Archiving Option archive job

You can create customized archive jobs by setting archive options on the archive job properties. You can also use the default settings that are set when Backup Exec is installed.

Certain types of data are not included in archive jobs.

See [“Types of data not included in Archiving Option archive jobs”](#) on page 1431.

Note: Data must be backed up before it can be archived.

You must perform the following actions before you can run an archive job:

- Create a vault store.
You can create a vault store when you create a job or any time before an archive job runs.
See [“Creating a vault store in the Archiving Option”](#) on page 1446.
- Ensure that the Backup Exec service account has the appropriate permissions to access the file system servers and the Exchange servers that you want to archive.
See [“About the Backup Exec service account”](#) on page 106.

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

See [“Editing default settings for archive jobs”](#) on page 1494.

See [“Best practices for the Archiving Option”](#) on page 1433.

See [“Types of data not included in Archiving Option archive jobs”](#) on page 1431.

Creating an Archiving Option archive job by setting job properties

Create an archive job by setting the properties that you want to use.

See [“About creating an Archiving Option archive job”](#) on page 1434.

To create an Archiving Option archive job by setting job properties

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Archive Tasks**, click **New archive job**.
- 3 In the task pane, under **Source**, do one or both of the following:

To select file system data to archive

Do the following in the order listed:

- Click **File System Selections**.
- Click **The same archive settings to all shares and folders**.
- Select the folders in which you want Backup Exec to find the data that is eligible for archiving.

To select specific shares and folders to which you want to apply different archive settings

Do the following in the order listed:

- Click **File System Selections**.
- Click **Different archive settings to specific shares and folders**.

See [“Applying different archive settings to file system share and folder selections for archive jobs”](#) on page 1459.

To select Exchange mailboxes to archive

Click **Exchange Selections**, and then select the appropriate Exchange Servers.

See [“Exchange selections options for archive jobs”](#) on page 1438.

- 4 In the task pane, under **Destination**, click **Vault Store**.
- 5 Select the appropriate options.
See [“Vault store options for archive jobs”](#) on page 1439.
- 6 In the task pane, under **Settings**, click **General**.
- 7 Select the appropriate options.
See [“General options for archive jobs”](#) on page 1440.

8 Do one of the following:

To archive file system selections

In the task pane, under **Settings**, click **File System**, and select the appropriate options.

See [“File system options for archive jobs”](#) on page 1441.

To archive Exchange selections

In the task pane, under **Settings**, click **Exchange**, and select the appropriate options.

See [“Exchange options for archive jobs”](#) on page 1443.

9 In the task pane, under **Settings**, click **Notification**.

10 Select the appropriate options.

See [“Notification options for jobs”](#) on page 682.

11 Do one of the following:

To run the job now

Click **Run Now**.

To configure scheduling options

In the task pane, under **Frequency**, click **Schedule**.

See [“Schedule options”](#) on page 356.

File System Selections options for archive jobs

You can select the folders or shares where you want Backup Exec to find data to archive. You can apply the same archive settings to all of the selections, or apply different archive settings to different selections.

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

Table N-6 File System Selections options for archive jobs

Item	Description
The same archive settings to all shares and folders	<p>Lets you apply the same retention category and archive rules to all of the eligible files and folders that you select.</p> <p>This option does the following:</p> <ul style="list-style-type: none"> ■ Keeps all of the eligible data for the same amount of time. ■ Uses the same rules to include or exclude all of the eligible data from the archive job. <p>You create the archive settings after you select the files and folders where you want Backup Exec to find data to archive.</p> <p>See “About archive settings in the Archiving Option” on page 1455.</p>
Different archive settings for specific shares and folders	<p>Lets you apply different retention categories and rules to the eligible files in the shares or folders that you select.</p> <p>You create the archive settings after you select the shares and folders where you want Backup Exec to find data to archive.</p> <p>See “Share and Folder Selections options for archive jobs” on page 1437.</p>
Show administrative shares	<p>Displays the administrative shares from which you can select the files and folders where you want Backup Exec to find data to archive. If you select files and folders from administrative shares, end users cannot use Backup Exec Retrieve to restore their own files.</p> <p>See “How Archiving Option end users retrieve archived data by using Backup Exec Retrieve” on page 1432.</p>

Share and Folder Selections options for archive jobs

You can select the file system shares or folders from which you want to archive data. For each selection, you can apply different archive settings.

See [“Applying different archive settings to file system share and folder selections for archive jobs”](#) on page 1459.

Table N-7 Share and Folder Selections options

Item	Description
Share and Folder Selections	Displays the share selections and the folder selections that you want to include or exclude from the archive job.
Type	Displays one of the following types: <ul style="list-style-type: none"> ■ Include Backup Exec searches the share selection or the folder selection for eligible data to include in the archive job. ■ Exclude Backup Exec searches the share selection or the folder selection for eligible data to exclude from the archive job.
Settings	Displays the name of the archive settings that you want to apply to this share selection or this folder selection. See “About archive settings in the Archiving Option” on page 1455.
Include/Exclude	Lets you select the shares and folders that you want to include in or exclude from the archive job.
Remove Selections	Lets you remove the share or folder from the selections list.
Assign Settings	Lets you select the retention category and archive rules to apply to specific share and folder selections. See “About archive settings in the Archiving Option” on page 1455.

Exchange selections options for archive jobs

You can select the Exchange Servers from which you want to archive data.

If the Exchange Servers that you want to archive do not appear in the list, ensure that the following items are properly configured:

- A Backup Exec Agent for Microsoft Exchange Servers license key for the Exchange Server has been entered on the media server.
See [“Viewing license information”](#) on page 174.
- The Remote Agent for Windows Systems that is installed on the Exchange Server publishes to the media server.
See [“About publishing the Remote Agent for Windows Systems to media servers”](#) on page 1951.

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1415.

Enter System Mailbox options for archive jobs

You can type the name of a mailbox for exclusive use by Backup Exec to log on to the Exchange Server.

Type the name using the following format:

SystemMailbox@domain.com

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1415.

Vault store options for archive jobs

When you create a new archive job, you must assign a vault store where Backup Exec stores the archived data.

See [“Creating a vault store in the Archiving Option”](#) on page 1446.

Table N-8 Vault store options for archive jobs

Item	Description
Server	Displays the name of the servers that are selected in the job.
Vault Store	<p>Displays the vault store where Backup Exec stores the archived data.</p> <p>See “About vault stores in the Archiving Option” on page 1445.</p> <p>You must assign a vault store to the server if one is not assigned.</p>

Table N-8 Vault store options for archive jobs (*continued*)

Item	Description
Assign Vault Store	Displays the available vault stores or lets you create a new vault store. See “Vault store selections” on page 1449. If you change the assigned vault store, the change only affects the mailboxes or shares that you archive after you reassign the vault store. The shares and mailboxes that already have an archive in the previously assigned vault store continue to be archived to that same archive.

General options for archive jobs

You can select general options for archive jobs.

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

Table N-9 General options for archive jobs

Item	Description
Job name	Displays the name for the archive job.
Job priority	Displays the priority of the access to the devices for this job. See “About job priority” on page 195.

Table N-9 General options for archive jobs (*continued*)

Item	Description
Archive from encrypted backup data	<p>Lets Backup Exec archive the data from the backup sets that are encrypted. The archived data is stored as decrypted data in the vault store. The data in the backup set remains encrypted.</p> <p>Only common encryption keys can be used to decrypt a backup set during an archive job. If a restricted key is used, then eligible items in the backup set are not archived.</p> <p>Note: This option applies only for the File System Archiving Option.</p> <p>See “About restricted keys and common keys in encryption” on page 417.</p>
Archive backup data that was created in the last x days	<p>Lets Backup Exec archive the data only from the backup sets that are as old as the specified number of days.</p> <p>The default number of days is 30.</p> <p>Use this option to limit Backup Exec to relevant backup sets in which to find eligible data to archive.</p> <p>Note: Backup Exec searches the backup sets of the specified server to find the data to archive. If backup jobs use the same selection list, then Backup Exec archives the data from the latest full backup and any subsequent incremental or differential backups.</p>

File system options for archive jobs

You can select options for file system archive jobs

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

Table N-10 File system options for archive jobs

Item	Description
Allow archiving from backup data that is on a tape device	<p>Lets Backup Exec archive the data from the backup sets that are on tapes.</p> <p>The tapes that contain the backup data that you want to archive must be available. The media server must have access to them in a tape drive or in a robotic library slot. Otherwise, the archive job completes with exceptions.</p>
Retention category	<p>Displays the retention category that applies to the file system selections in the archive job. A retention category specifies the time period for which you want to keep archived items.</p>
New	<p>Displays the information for you to create a new retention category.</p> <p>See “Retention category properties” on page 1459.</p>
Rule	<p>Displays the name of the rule that you specify.</p>
Type	<p>Indicates if the rule includes or excludes the specified data in the archive job.</p>
New	<p>Lets you create a new archive rule to add to the list of rules in the file system archive settings.</p> <p>See “Archive rules in the File System Archiving Option” on page 1443.</p>
Edit	<p>Lets you edit an existing archive rule.</p>
Delete	<p>Deletes an archive rule from the list of rules in the file system archive settings.</p>
Move Up	<p>Moves a rule up in the list of rules. An item is archived according to the first rule for which it meets the criteria. The top rule in the list is the first rule that applies.</p>
Move Down	<p>Moves a rule down in the list of rules.</p>

Archive rules in the File System Archiving Option

You can configure the rules that specify the characteristics of the data to include or exclude in the archive job.

Table N-11 Options for rules in the File System Archiving Option

Item	Description
Include when archiving	Specifies that the files that meet the requirements that you select are included in the archive job. This option is enabled by default.
Exclude from archiving	Specifies that the files that meet the requirements that you select are excluded from the archive job.
Files of the following types	Specifies the types of file to include or exclude from the archive job. You can type your own rule or use a predefined rule.
Files not accessed in	Includes or excludes the files that have not been accessed in a specified number of days. The default is to include the files that have not been accessed in 30 days in the archive job.
Files not modified in	Includes or excludes the files that have not been modified in a specified number of days.
Files not created in	Includes or excludes the files that have not been created in a specified number of days.
File size	Includes or excludes the files that are greater than or equal to, or less than or equal to a specified size. The default is to include the files that are greater than or equal to 10 MB in the archive job.

Exchange options for archive jobs

You can configure an archive job for Exchange mailboxes.

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

Table N-12 Exchange options for archive jobs

Item	Description
Exchange Server Name	Displays the name of the Exchange Server that contains the mailbox selections that you want to archive.
System Mailbox	<p>Displays the name of the system mailbox on the Exchange Server for Backup Exec to log on to.</p> <p>If a system mailbox is not assigned, you must assign one.</p> <p>You must configure a mailbox for exclusive use by Backup Exec on each Exchange Server on which you want to select mailboxes for archiving.</p> <p>See “About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option ” on page 1419.</p>
Assign System Mailbox	<p>Lets you assign a system mailbox for exclusive use by Backup Exec on the Exchange Server.</p> <p>See “Enter System Mailbox options for archive jobs” on page 1439.</p>
Select a domain	Displays the domains that you can select.
Mailbox Group	<p>Displays the names of the mailbox groups in the selected domain that this job archives.</p> <p>See “About Exchange mailbox groups in archive jobs” on page 1461.</p> <p>If a mailbox group does not appear, you must create one.</p>
New	<p>Lets you create a mailbox group to add to the archive job.</p> <p>See “Mailbox group options” on page 1462.</p>
Edit	<p>Lets you edit the selected mailbox group.</p> <p>See “Mailbox group options” on page 1462.</p>

Table N-12 Exchange options for archive jobs (*continued*)

Item	Description
Delete	Deletes the selected mailbox group from the list of mailbox group to archive.
Move Up	Moves the mailbox group up in the list of mailbox groups. The archive settings apply to the mailbox groups in the order in which the mailbox groups are listed. A mailbox that belongs to multiple groups is archived according to the archive settings of the highest-level group that it is in. See “About Exchange mailbox groups in archive jobs” on page 1461.
Move Down	Moves the selected mailbox group lower in the list.

About vault stores in the Archiving Option

A vault store is a disk-based container for the archived data that Backup Exec archives from one server. When you create an archive job, you select a vault store as the device to which you want to send the archived data. A vault store contains at least one vault store partition that is the physical location where the archived items are stored. You can create additional vault store partitions for a vault store when it requires more disk space.

Each vault store has an associated database. The database holds information about the archives in the vault store and all the items that are stored in each archive. For example, when an item is archived, the vault store's database is updated with this information. Single instance storage-related information is contained in the fingerprint databases for all of the vault stores.

The following vault store properties let you manage the deletion of archived items:

- Delete an item from its original location on the resource immediately after it is archived, or after the vault store is backed up.
See [“About deleting archived data from its original location”](#) on page 1477.
- Delete the archived items that have expired retention periods from specific archives in vault stores.
See [“Preventing the deletion of expired archived items from an archive”](#) on page 1493.

You can back up vault stores and their associated databases, along with other Archiving Option components.

See [“About single instance storage of archived items”](#) on page 1494.

See [“About backing up Archiving Option components”](#) on page 1477.

See [“Creating a vault store in the Archiving Option”](#) on page 1446.

See [“Editing or viewing vault store properties”](#) on page 1447.

See [“About vault store partitions in the Archiving Option”](#) on page 1451.

Creating a vault store in the Archiving Option

You can create a vault store for Backup Exec to store data from archive jobs.

See [“About vault stores in the Archiving Option”](#) on page 1445.

To create a vault store in the Archiving Option

- 1 On the navigation bar, click **Devices**.
- 2 In the task pane, under **Device Tasks**, click **Configure Devices Assistant**.
- 3 Under **Archiving Option**, click **Vault Store**.
- 4 Select the appropriate options.

See [“New vault store options”](#) on page 1446.

New vault store options

You can create a new vault store.

See [“Creating a vault store in the Archiving Option”](#) on page 1446.

Table N-13 New vault store options

Item	Description
Name	Lets you specify the name of the vault store where Backup Exec stores archived data. Note: Do not name a vault store with the same name that is already in use by an Archiving Option component, such as Fingerprint Databases or All Partitions . A vault store that has the same name as another Archiving Option component can cause errors when you make backup selections. Backup job or restore job failures can also occur.
Description	Lets you specify a description of the vault store. You can edit this field to change the description.
Path	Lets you specify the path name where Backup Exec automatically creates the first vault store partition.
Immediately after archiving	Deletes the archived item from its original location after the item is archived in the vault store. See “About deleting archived data from its original location” on page 1477.
After vault store backup	Deletes the archived item from its original location after the vault store is backed up and the next archive job runs. See “About deleting archived data from its original location” on page 1477.

Editing or viewing vault store properties

You can edit or view vault store properties.

See [“About vault stores in the Archiving Option”](#) on page 1445.

To edit or view vault store properties

- 1 On the navigation bar, click **Devices**.
- 2 Select the vault store for which you want to edit the properties.

- 3 In the task pane, under **General Tasks**, click **Properties**.
- 4 Edit the information as appropriate.
See “[Vault store properties](#)” on page 1448.

Vault store properties

You can edit the vault store properties.

See “[Editing or viewing vault store properties](#)” on page 1447.

Table N-14 Vault store properties

Item	Description
Name	Displays the name of the vault store where Backup Exec stores archived data. You can edit this field to change the name. Note: Do not name a vault store with the same name that is already in use by an Archiving Option component, such as Fingerprint Databases or All Partitions . A vault store that has the same name as another Archiving Option component can cause errors when you make backup selections. Backup job or restore job failures can also occur.
Description	Displays a description of the vault store. You can edit this field to change the description.
Database name	Displays the name of the database that is associated with this vault store.
State	Displays the following states: <ul style="list-style-type: none">■ Available Archive jobs can send data to this vault store.■ Being deleted The vault store is in the process of being deleted. Archive jobs cannot send data to this vault store.■ In backup mode A backup or restore job is running for the vault store.

Table N-14 Vault store properties (*continued*)

Item	Description
Item deletion mode	<p>Designates when to delete archived items from their original locations.</p> <p>You can choose to delete the item immediately after it is archived or delete it after the vault store is backed up. If you delete an item immediately after it is archived, the item is deleted from its original location after the archive job has successfully completed.</p> <p>See “About deleting archived data from its original location” on page 1477.</p>
Archive Count	Displays the number of archives that the vault store contains.
Total size	Displays the total size of all of the items that are archived in the vault.

Vault store selections

You can select a vault store for an archive job destination.

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

Table N-15 Vault store selections

Item	Description
Vault Store	<p>Displays the available vault stores that you can assign to a server. Backup Exec stores the data that is archived from the server in the vault store that you select.</p> <p>If you change the assigned vault store, the change only affects the mailboxes or shares that you archive after you reassign the vault store.</p> <p>The shares and mailboxes that already have an archive in the previously assigned vault store continue to be archived to that same archive.</p>

Table N-15 Vault store selections (*continued*)

Item	Description
New	Lets you create a new vault store that you can then assign to a server. See “Vault store properties” on page 1448.

About deleting an Archiving Option vault store

You can delete a vault store if one of the following conditions apply:

- It is not assigned to any servers on which you have made archive selections.
- It is the only vault store and you have deleted all of the archive jobs.

When you delete a vault store, you cannot cancel or undo the operation.

When you delete a vault store, all partitions, archives, and archived items in that vault store are also deleted. You must reassign another vault store to all of the servers that were assigned to the vault store that you want to delete.

See [“Deleting a vault store”](#) on page 1450.

Deleting a vault store

You can delete a vault store from Backup Exec.

See [“About deleting an Archiving Option vault store”](#) on page 1450.

To delete a vault store

- 1 On the navigation bar, click **Devices**.
- 2 Select the vault store that you want to delete.
- 3 In the task pane, under **General Tasks**, click **Delete**.
- 4 If no other vault stores exist, you must do one of the following:
 - Delete all of the existing archive jobs before you can delete this vault store.
See [“Deleting scheduled jobs”](#) on page 571.
 - Create a new vault store, assign it to all of the affected archived servers, and then delete the selected vault store.
See [“Creating a vault store in the Archiving Option”](#) on page 1446.

About vault store partitions in the Archiving Option

A vault store partition represents the physical location where the archived items are stored. A vault store can contain one or more vault store partitions. Backup Exec creates one vault store partition in each vault store by default.

As the data in a vault store grows, you can create more vault store partitions to provide additional capacity. You can specify a local drive or a network share as a location for a vault store partition. You cannot specify a path that is a subdirectory in the path for another vault store partition.

A vault store can contain many vault store partitions, but only one partition is open at a time. As data is archived, it is stored in the open partition. You can specify a vault store partition as open or closed by editing the partition properties.

You can restore archived items from closed partitions, as well as delete the archived items that are in closed partitions.

Backup Exec searches the vault store partitions daily to delete the archived items that have expired retention periods. You can specify the time at which this daily operation runs.

See [“Preventing the deletion of expired archived items from an archive”](#) on page 1493.

See [“Creating a vault store partition”](#) on page 1451.

See [“Editing vault store partition properties”](#) on page 1452.

See [“About vault stores in the Archiving Option”](#) on page 1445.

See [“About archives in the Archiving Option ”](#) on page 1453.

Creating a vault store partition

You can create a new vault store partition.

See [“About vault store partitions in the Archiving Option”](#) on page 1451.

To create a vault store partition

- 1 On the navigation bar, click **Devices**.
- 2 Expand the vault store.
- 3 Right-click **Partitions**, and then on the shortcut menu, click **New vault store partition**.
- 4 Enter the appropriate information.

See [“Vault store partition properties”](#) on page 1452.

Editing vault store partition properties

You can change the state of a vault store partition to open or closed. You can also edit the name and description of a vault store partition.

See [“About vault store partitions in the Archiving Option”](#) on page 1451.

To edit vault store partition properties

- 1 On the navigation bar, click **Devices**.
- 2 Expand the vault store that contains the vault store partition that you want to edit.
- 3 In the right pane, select the vault store partition that you want to edit.
- 4 In the task pane, under **General Tasks**, click **Properties**.
- 5 Edit the appropriate information.

See [“Vault store partition properties”](#) on page 1452.

Vault store partition properties

A vault store partition represents the physical location where the archived items are stored. You can create a new vault store partition or change the state of an existing vault store partition.

See [“Creating a vault store partition”](#) on page 1451.

See [“Editing or viewing vault store properties”](#) on page 1447.

Table N-16 Vault store partition properties

Item	Description
Name	Displays the name of the vault store partition.
Description	Displays a description of the vault store partition.

Table N-16 Vault store partition properties (*continued*)

Item	Description
Location	<p>Displays the path name where the vault store partition is located.</p> <p>The path can be on a local drive or on a network share. You cannot specify a path that is a subdirectory in the path for another vault store partition.</p> <p>For example, you can create a vault store partition on C:\vault store 1. However, you cannot create another vault store partition on C:\vault store 1\vault store 2.</p> <p>Ensure that the Backup Exec service account has full permissions for the path.</p> <p>See “About the Backup Exec service account” on page 106.</p>
State	<p>Displays one of the following states:</p> <ul style="list-style-type: none"> ■ Open New archived data is stored in this vault store partition. ■ Closed New archived data cannot be stored in this vault store partition. <p>See “About vault store partitions in the Archiving Option” on page 1451.</p>

About archives in the Archiving Option

An archive is a logical group of archived items. Items in an archive are stored in different vault store partitions depending on which partition is open at the time that the item is archived. Each archived file system share has its own archive, and each archived Exchange mailbox has its own archive. Backup Exec creates the archives when it creates an archive job.

You cannot back up archives. You can only back up the vault store partitions.

See [“Editing archive properties”](#) on page 1454.

See [“Deleting an archive”](#) on page 1455.

Editing archive properties

You can edit archive properties.

See “[About archives in the Archiving Option](#)” on page 1453.

To edit archive properties

- 1 On the navigation bar, click **Devices**.
- 2 Expand the vault store, and then select **Archives**.
- 3 In the right pane, select the archive for which you want to edit the properties.
- 4 In the task pane, under **General Tasks**, click **Properties**.

See “[Archive properties](#)” on page 1454.

Archive properties

You can view archive properties. You can also edit the setting to let Backup Exec automatically delete the archived items that have expired retention periods.

See “[Editing archive properties](#)” on page 1454.

Table N-17 Archive properties

Item	Description
Name	Displays the name of the file share or Exchange mailbox that is archived.
Type	Displays one of the following types of archive: <ul style="list-style-type: none">■ File share■ Exchange mailbox
Status	Displays one of the following statuses as appropriate: <ul style="list-style-type: none">■ Available■ Being created■ Being deleted
Server	Displays the name of the server on which the archive is stored.

Table N-17 Archive properties (*continued*)

Item	Description
Automatically delete archived items that have expired retention periods	Lets Backup Exec delete the archived items that have expired retention periods from the archives. You can set a time at which Backup Exec deletes these items daily. See “Editing default settings for archive jobs” on page 1494. Uncheck this check box if you do not want archived items to be automatically deleted from this archive. This option is enabled by default.

Deleting an archive

You can delete an archive. However, if you delete an archive from Backup Exec, all of the archived data in the archive is also deleted.

See [“About archives in the Archiving Option ”](#) on page 1453.

To delete an archive

- 1 On the navigation bar, click **Devices**.
- 2 Expand the vault store, and then select **Archives**.
- 3 In the right pane, select the archive that you want to delete.
- 4 In the task pane, under **General Tasks**, click **Delete**.
- 5 When you are prompted if you want to delete the archive, click **Yes**.

About archive settings in the Archiving Option

Archive settings let you apply the following criteria to the file system shares or folders or to Exchange mailboxes:

- The retention category that specifies how long to keep the data in the archives.
- The rules that determine if data is eligible for archiving.

For example, you can specify that only the mail messages that are older than six months are archived for a mailbox selection.

You can create archive settings for the following selections:

- Exchange mailboxes
- File system shares
- File system folders within the shares

Note: You can name each group of archive settings that you create. You can only name archive settings when you select the option **Apply different archive settings for specific folders** when you create an archive job.

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

See [“Applying different archive settings to file system share and folder selections for archive jobs”](#) on page 1459.

Archive settings options

You can specify settings to apply to the file system shares or folders in which you want Backup Exec to search for eligible data to archive.

See [“Applying different archive settings to file system share and folder selections for archive jobs”](#) on page 1459.

Table N-18 Archive settings options

Item	Description
Name	Specifies the name of the archive settings to apply to Exchange mailbox selections or file system selections. You can apply these same archive settings to other selections.
Retention Category	Specifies the name of the retention category to apply to the selections.
New	Lets you create a new retention category. See “Retention category properties” on page 1459.
Include when archiving	Specifies that the files that meet the requirements that you select are included in the archive job. This option is enabled by default.

Table N-18 Archive settings options (*continued*)

Item	Description
Exclude from archiving	Specifies that the files that meet the requirements that you select are excluded from the archive job.
Files of the following types	Specifies the types of file to include or exclude from the archive job. You can type your own rule, or use a predefined rule.
Files not accessed in	Includes or excludes the files that have not been accessed in a specified number of days. The default is to include the files that have not been accessed in 30 days in the archive job.
Files not modified in	Includes or excludes the files that have not been modified in a specified number of days.
Files not created in	Includes or excludes the files that have not been created in a specified number of days.
File size	Includes or excludes the files that are greater than or equal to or less than or equal to a specified size. The default is to include the files that are greater than or equal to 10 MB in the archive job.
Add Rule	Adds a rule to the list of rules in the archive settings. This rule applies when you run the archive job for the file system selections.
Delete Rule	Deletes a rule from the list of rules in the archive settings.
Move Up	Moves a rule up in the list of rules. An item is archived according to the first rule for which it meets the criteria. The top rule in the list is the first rule that applies.
Move Down	Moves a rule down in the list of rules.

About retention categories for archived items

Use retention categories to specify the period of time for which you want to keep items in the archives. You can give the retention categories meaningful names, such as Business or Personal. Retention categories make it easier for you to retrieve

items because you can search for them by their category name. Each retention category has a retention period, which indicates how long you want to retain the items that are archived with this retention category.

For example, you can create a retention category named Finance Data Retention and set it to retain archived data for seven years.

The retention period starts on the date that the item is archived. Backup Exec runs a daily operation that deletes all items that have expired retention periods. You can prevent this operation from running on specific archives.

See [“Preventing the deletion of expired archived items from an archive”](#) on page 1493.

You cannot delete retention categories. You can edit retention categories, including the retention periods.

Changes that you make to a retention category apply to the following:

- All items to which the retention category is already applied.
- Any new items to which you apply the retention category.
See [“Editing a retention category ”](#) on page 1458.

You can create retention categories as needed when you create an archive job. You can also specify a retention category to use as the default setting for all archive jobs. If you do not specify a retention category, then a default retention category with a retention period of infinite is applied to an archive job.

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

See [“Editing a retention category ”](#) on page 1458.

See [“Editing default settings for archive jobs”](#) on page 1494.

Editing a retention category

You can edit an existing retention category. Changes apply to existing archived items as well as to new items to which you apply the retention category.

See [“About retention categories for archived items”](#) on page 1457.

To edit a retention category

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Archive**.
- 3 In the **Default retention category** field, select the retention category that you want to edit.
- 4 Click **Edit**.

- 5 Edit the appropriate information.
See [“Retention category properties”](#) on page 1459.
- 6 Click **OK**.

Retention category properties

Create a retention category to specify the period of time for which you want to keep items in the archives.

See [“About retention categories for archived items”](#) on page 1457.

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

See [“Editing default settings for archive jobs”](#) on page 1494.

Table N-19 Retention category properties

Item	Description
Name	Displays the name of the retention category.
Description	Displays a description of the retention category.
Infinite	Retains the item in the archives for an infinite amount of time. The retention period starts from the date that the item is archived.
For a period of	Retains the item in the archives for a specified period of time. The retention period starts from the date that the item is archived.

Applying different archive settings to file system share and folder selections for archive jobs

You can select specific file system shares and folders and apply different archive settings to them in the same archive job.

See [“About archive settings in the Archiving Option”](#) on page 1455.

To apply different archive settings to file system share and folder selections for archive jobs

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Archive Tasks**, click **New archive job**.
- 3 In the task pane, under **Source**, click **File System Selections**.

- 4 Click **Different archive settings to specific shares and folders**.
- 5 Click **Include/Exclude**.
- 6 Complete the options as appropriate.
See “[Include/Exclude Selections options for archive jobs](#)” on page 1460.
- 7 Click **OK**.
- 8 Click the share selection or folder selection for which you want to assign archive settings, and then click **Assign Settings**.
- 9 Complete the options as appropriate.
See “[Archive settings options](#) ” on page 1456.
- 10 Click **OK**.
- 11 Click **Include/Exclude** again and make selections as needed.
- 12 Click each share selection or folder selection, and then click **Assign Settings** and create the archive settings that you want to apply.
- 13 Continue setting job properties to complete the archive job
See “[Creating an Archiving Option archive job by setting job properties](#)” on page 1435.

Include/Exclude Selections options for archive jobs

You can include the file system shares and folders in which Backup Exec can search for data to archive. You can also exclude the file system shares and folders that you do not want to include in the archive job.

See “[Applying different archive settings to file system share and folder selections for archive jobs](#)” on page 1459.

Table N-20 Include/Exclude Selections options

Item	Description
Include when archiving	Lets you include the selections that are eligible for archiving.
Exclude from archiving	Lets you exclude the selections from the archive job.
All Resources	Lists the resources that you can select to include in or exclude from the archive job.

Table N-20 Include/Exclude Selections options (*continued*)

Item	Description
Show administrative shares	Displays the administrative shares from which you can select the files and folders where you want Backup Exec to find data to archive. If you select files and folders from administrative shares, end users cannot use Backup Exec Retrieve to restore their own files. See “How Archiving Option end users retrieve archived data by using Backup Exec Retrieve” on page 1432.

About Exchange mailbox groups in archive jobs

A mailbox group contains the selections on the Exchange Server that you want to archive.

A mailbox group consists of user mailboxes to which you want to assign the same archive settings. For example, you can add a single user to a mailbox group, or you can add the entire Exchange organizational unit to a mailbox group.

In the **Archive job properties**, in the **Exchange** settings, Backup Exec applies the archive settings sequentially to each mailbox group in the list. The archive settings of the first mailbox group that a mailbox is found in are applied to that mailbox.

The order of the mailbox groups is important. You should arrange the mailbox groups that have specific selections of users, groups, and distribution lists at the top of the list. Arrange the mailbox groups that contain the least specific selections at the bottom of the list. For example, a mailbox group that contains specific users should be listed before a mailbox group that contains a user group. In turn, a mailbox group that contains a user group should be listed before a mailbox group that contains the whole Exchange organizational unit. For example, you want to ensure that the correct archive settings are applied to the users that are in multiple groups.

You would arrange the following example mailbox groups in the order listed:

- The Managers group contains individual user accounts and requires all messages to be archived.
- The Some Users group contains some users in an organizational unit and requires messages to be archived from the last two months.

- The All Users group contains the entire Exchange organizational unit and requires messages to be archived from the last six months.

You can select the following items to archive in a mailbox group:

- Distribution lists
- User groups
- Users

You can create mailbox groups when you create an archive job for Exchange Server mailboxes, or at any time from the **Archive Job Defaults** dialog box.

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

See [“Managing Exchange mailbox groups”](#) on page 1463.

Mailbox group options

For an archive job, create a mailbox group that contains the selections on the Exchange Servers that you want to archive. You can also specify the retention category and archive rules for each group.

See [“About Exchange mailbox groups in archive jobs”](#) on page 1461.

See [“Creating an Archiving Option archive job by setting job properties”](#) on page 1435.

Table N-21 Mailbox group options

Item	Description
Mailbox group name	Designates the name of the mailbox group.
Retention category	Lets you specify the retention category for the mailbox group. The default setting is the default retention category, which has a retention period of infinite. See “About retention categories for archived items” on page 1457.
New	Lets you create a new retention category. See “Retention category properties” on page 1459.

Table N-21 Mailbox group options (*continued*)

Item	Description
Are older than	Indicates that the items that are older than the specified time are archived. The default setting is one year.
Are larger than and older than	Indicates that the items that are larger than the specified size and are older than the specified time are archived. You should archive larger mail messages more often than other messages. Specify a lesser amount of time in this option than in the previous Are older than option. The default setting is 1 MB and one year.
Archive only messages with attachments	Indicates that the messages with attachments are archived. The option is enabled by default.
Archive unread messages	Indicates that the messages that have not been read are archived.

Managing Exchange mailbox groups

You can configure and manage mailbox groups for archive jobs for the Exchange Mailbox Archiving Option.

See [“About Exchange mailbox groups in archive jobs”](#) on page 1461.

To manage Exchange mailbox groups

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Archive**.
- 3 Click **Manage Mailbox Groups**.
- 4 Enter the necessary information as appropriate.

See [“Manage mailbox groups options”](#) on page 1463.

Manage mailbox groups options

You can configure or edit the mailbox groups that contain the selections for an archive job.

See [“Managing Exchange mailbox groups”](#) on page 1463.

Table N-22 Manage mailbox group options

Item	Description
Select a domain	Displays the domains that you can select.
Mailbox Group	<p>Displays the names of the mailbox groups in the selected domain that this job archives.</p> <p>See “About Exchange mailbox groups in archive jobs” on page 1461.</p> <p>If there are no mailbox groups, you must create one before you can run an archive job.</p>
New	<p>Lets you create a mailbox group to add to an archive job.</p> <p>See “Mailbox group options” on page 1462.</p>
Edit	<p>Lets you edit the selected mailbox group.</p> <p>See “Mailbox group options” on page 1462.</p>
Delete	Deletes the selected mailbox group from the list of mailbox groups to archive.
Move Up	<p>Moves the selected mailbox group up in the list.</p> <p>The retention category and the archive rules apply to the mailbox groups in the order in which the mailbox groups are listed. A mailbox that belongs to multiple groups is archived according to the archive settings of the highest-level group that it is in.</p> <p>See “About Exchange mailbox groups in archive jobs” on page 1461.</p>
Move Down	Moves the selected mailbox group lower in the list.

About searching for data in the archives

You can search the archives to find and select the data that you want to delete or restore from the archives. The archives contain access control restrictions, but these restrictions are not applied when you search from the Backup Exec Administration Console. The search displays all archived versions of the data.

You can specify criteria such as content, retention categories, or retention periods. You can also restrict the search to one archive or to all archives that are associated with a server.

See [“Searching for data in the archives”](#) on page 1465.

Searching for data in the archives

You can search the archives to find data.

See [“About searching for data in the archives”](#) on page 1464.

To search for data in the archives

- 1 On the **Edit** menu, click **Search Archives**.
- 2 Select the archive that you want to search.
- 3 Do one of the following:

To search for file system data in the archives

On the **File System** tab, enter the appropriate information.

See [“Search Archives options for file system selections”](#) on page 1466.

To search for Exchange messages in the archives

On the **Exchange** tab, enter the appropriate information.

See [“Search archives options for Exchange selections”](#) on page 1467.

- 4 Click **Find Now**.
- 5 Select other actions as appropriate.
See [“Search Archives options”](#) on page 1465.

Search Archives options

You can search for the specific items that are in the archives.

See [“Searching for data in the archives”](#) on page 1465.

Table N-23 Search Archives options

Item	Description
Search archive	Designates the archive in which Backup Exec searches for items.

Table N-23 Search Archives options (*continued*)

Item	Description
Apply	Applies the version of the file that you select in the search results window to the list of selections in the View by Resource tab.
Find Now	Starts the search for the specified items.
Stop	Cancel the search for the specified items.
New Search	Clears any existing criteria so that you can specify new criteria to search for items.
Close	Exits the Search Archives dialog box.

Search Archives options for file system selections

You can search for the file system items that are in the archives.

See [“Searching for data in the archives”](#) on page 1465.

Table N-24 Search archives options for file system selections

Item	Description
File name	Searches for the files that match this text. Leave this field blank to search all files. You can use wildcard characters. Use a question mark (?) to represent any single characters. Use an asterisk (*) to represent any number of characters. For example, to include all files with the exe extension, type *.exe.
Path	Searches for the files in the specified path. Leave this field blank to search all directories.
File content	Searches for the files that have content that match this text.
File size	Searches for the files that match the specified size.
File modified	Searches for the files that are created or modified in the specified time period.

Table N-24 Search archives options for file system selections (*continued*)

Item	Description
Archived	Searches for the archived files that match the range of dates.
Will exceed retention	Searches for the files with the retention periods that match the range of dates.
Retention category	Searches for files in the specified retention category.

Search archives options for Exchange selections

You can search for the Exchange items that are in the archives.

You can use wildcard characters. Use a question mark (?) to represent any single character. Use an asterisk (*) to represent any number of characters.

See [“Searching for data in the archives”](#) on page 1465.

Table N-25 Search archives options for Exchange selections

Item	Description
Subject	Searches for the mail messages that have matching text in the subject line.
Content	Searches for the mail messages that have matching text in the content line.
From	Searches for the mail messages that have matching text in the From field.
To	Searches for the mail messages that have matching text in the To field.
Email size	Searches for the mail messages that match the email size that you specify.
Has attachments	Searches for the mail messages that have attachments.
Received	Searches for the received mail messages that match the range of dates.
Archived	Searches for the archived mail messages that match the range of dates.

Table N-25 Search archives options for Exchange selections *(continued)*

Item	Description
Will exceed retention	Searches for the mail messages with the retention periods that match the range of dates.
Retention category	Searches for mail messages in the specified retention category.

About restoring items from the archives

You can perform the following restore operations for archived items:

- Restore files to their original locations or to another file server.
- Restore mail messages to the original mailbox or to another mailbox on the Exchange server.

Note: The mailbox must already exist on the server to which you want to restore the mail messages.

If you enable Backup Exec Retrieve, end users can retrieve their own data by using Backup Exec Retrieve.

The archives can contain multiple versions of the same item. To restore a specific version of the item, you must select it individually. Otherwise, Backup Exec restores the latest version of an item. You can distinguish between versions of the same file by checking the modified time of an item.

Note: Access permissions for archived data are not restored.

See [“Restoring items from archives”](#) on page 1468.

Restoring items from archives

You can restore data from archives by selecting the job properties that you want to use.

See [“About restoring items from the archives”](#) on page 1468.

To restore data from archives

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Archive Tasks**, click **Restore from Archives**.
- 3 Select the data that you want to restore from the archives.
See [“Selections options for restoring items from archives”](#) on page 1470.
- 4 Complete the following options as necessary:

To redirect archived files to a computer other than the one from which the data was archived

In the **Properties** pane, under **Destination**, do any of the following:

- Click **File Redirection** to redirect archived file sets.
See [“File redirection options to restore items from archives”](#) on page 1472.
- Click **Microsoft Exchange Redirection**.
See [“Microsoft Exchange redirection options to restore items from archives”](#) on page 1473.

To set general options for the restore from archives job

In the **Properties** pane, under **Settings**, click **General**.

See [“General options to restore items from archives”](#) on page 1470.

To set options for Exchange data for the restore from archives job

In the **Properties** pane, under **Settings**, click **Microsoft Exchange**.

See [“Microsoft Exchange options to restore items from archives”](#) on page 1471.

To set a network interface and protocol to use for the restore from archives job

In the **Properties** pane, under **Settings**, click **Network and Security**.

See [“Network and security restore options”](#) on page 617.

To set commands to run before or after the restore from archives job

In the **Properties** pane, under **Settings**, click **Pre/Post Commands**.

See [“Default Pre/Post Commands options”](#) on page 400.

To configure notification when the restore from archives job completes

In the **Properties** pane, under **Settings**, click **Notification**.

See [“Notification options for jobs”](#) on page 682.

5 Do one of the following:

To run the restore from archives job now

Click **Run Now**.

To schedule the restore from archives job for later

In the **Properties** pane, under **Frequency**, click **Schedule**.

See [“Schedule options”](#) on page 356.

Selections options for restoring items from archives

You can select the data that you want to restore from the archives.

See [“Restoring items from archives ”](#) on page 1468.

Table N-26 Selections options for restore from archives jobs

Item	Description
Include subdirectories	Selects the contents of all the subfolders when a directory is selected.
Show file details	Displays the details about the files that can be restored from the archives.
Preview pane	Displays the preview pane at the bottom of the dialog box.
Search archives	Lets you search for the data that you want to restore from the archives. See “Search Archives options” on page 1465.
View by Resource	Displays archived data by the resource from which it is archived. This feature is useful for finding the files that are located on a specific computer.

General options to restore items from archives

You can specify a name and a job priority for the job.

See [“Restoring items from archives ”](#) on page 1468.

Table N-27 General options to restore items from archives

Item	Description
Job name	Displays a name that identifies this job in the job schedule.
Job priority	Displays the priority of the access to the devices for this job. See “About job priority” on page 195.
Restore over existing files	Overwrites the files that are on the destination drive that have the same name as the files that you want to restore from the archives. Use this option only when you are sure that you want to restore an older version of a file from the archives.
Skip if file exists	Prevents Backup Exec from overwriting the files that are on the destination drive with files from the archives that have the same names.
Overwrite the file on disk only if it is older	Prevents Backup Exec from overwriting the files that are on the destination drive if they are more recent than the files from the archives. This option is enabled by default.
Preserve tree	Restores the files from the archives with their original directory structure intact. This option is enabled by default. If you clear this option, all data including subdirectories is restored from the archives to the path that you specify in the Redirection dialog box. You may want to clear this option when you restore several subdirectories or individual files from the archives. You should not clear this option if you restore an entire share from the archives.

Microsoft Exchange options to restore items from archives

You can restore over existing messages and folders when you restore mail messages from the archives.

Note: The mailbox that you restore to must already exist. It is not created as part of the restore job.

See [“Restoring items from archives”](#) on page 1468.

The option **When restoring individual mail messages, restore over existing messages** replaces an existing message with the message that you restore from the archives. A new object ID is not created for the restored message. Only the contents and properties of the message are replaced.

This option is not enabled by default.

If this option is not enabled, or if the original message does not exist, then the message is recreated as a new message. Backup Exec creates a new object ID for the recreated message.

If this option is not enabled and if the original message does exist, then the message is not restored from the archives.

File redirection options to restore items from archives

You can restore file system data from the archives to another drive or path other than where the data was originally backed up.

See [“Restoring items from archives”](#) on page 1468.

Table N-28 File redirection options to restore items from archives

Item	Description
Redirect file sets	Lets you restore the data from the archives to a drive or path other than where the data was originally backed up.
Restore to drive	Designates the destination drive to which you want to restore the data from the archives.
Browse (...)	Lets you view local and network drives.
Server logon account	Displays the current logon account that the media server uses.
Change	Lets you use a different logon account or create a new one. See “About configuring logon accounts” on page 184.

Table N-28 File redirection options to restore items from archives (*continued*)

Item	Description
Clear	Lets you clear this field.
Restore to path	Designates the destination path of the device that is listed in the Restore to Drive field to which you want to restore data from the archives. To retain the original directory structure, ensure that the Preserve tree option is enabled. See “Restoring items from archives ” on page 1468.
Path logon account	Displays the logon account that is required for the destination path.
Change	Lets you use a different logon account or create a new one. See “About configuring logon accounts” on page 184.
Clear	Lets you clear this field.

Microsoft Exchange redirection options to restore items from archives

You can redirect the restore of Exchange Mailbox Archiving items from the archives.

See [“Restoring items from archives ”](#) on page 1468.

Table N-29 Microsoft Exchange redirection options to restore items from archives

Item	Description
Redirect Exchange sets	Lets you restore mail messages and folders from the archives to a drive or path other than where the data was originally backed up.
Restore to server	Specifies the name of the computer to which you want to restore the data. The name of the computer uses the format \\server name.

Table N-29 Microsoft Exchange redirection options to restore items from archives (*continued*)

Item	Description
Server logon account	Displays the current logon account that the media server uses.
Change	Lets you use a different logon account or create a new logon account. See “About configuring logon accounts” on page 184.
Clear	Lets you clear this field.
Redirect mailboxes	Lets you restore mailboxes from the archives to a mailbox other than where the data was originally backed up.
Restore to mailbox	Specifies the name of the mailbox to which you want to redirect the restore. The mailbox must already exist on the server to which you want to restore the data from the archives.
Mailbox logon account	Displays the logon account that is required for the destination mailbox. Click Clear to clear this field.
Change	Lets you use a different logon account or create a new logon account.
Clear	Lets you clear the field.

About deleting items from the archives

You can delete archived files and mail messages from the archives. If you need to free some disk space, you can delete items from the archives before their retention periods expire.

The archives can contain multiple versions of the same item. To delete a specific version of the item, you must select it individually. Otherwise, Backup Exec deletes the latest version of an item. You can distinguish between versions of the same file by checking the modified time of an item.

You can delete only files and mail messages from the archives. To delete an entire archive, you must delete it from the **Devices** view.

Additionally, Backup Exec searches the vault store partitions daily to delete the archived items that have expired retention periods. You can specify the time at which this daily operation runs.

See [“Editing default settings for archive jobs”](#) on page 1494.

See [“Deleting an archive”](#) on page 1455.

See [“Deleting items from the archives”](#) on page 1475.

See [“About searching for data in the archives”](#) on page 1464.

Deleting items from the archives

You can delete specific items from the archives.

See [“About deleting items from the archives”](#) on page 1474.

To delete items from the archives

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Archive Tasks**, click **Delete from archives**.
- 3 Select the data that you want to delete.
See [“Selections options to delete items from the archives”](#) on page 1476.
- 4 In the **Properties** pane, under **Settings**, click **General**.
- 5 Select the appropriate options.
See [“General options to delete items from the archives”](#) on page 1476.
- 6 In the **Properties** pane, under **Settings**, click **Network and Security**.
- 7 Select the appropriate options.
See [“Network and security restore options”](#) on page 617.
- 8 In the **Properties** pane, under **Settings**, click **Pre/Post Commands**.
- 9 Select the appropriate options.
See [“Default Pre/Post Commands options”](#) on page 400.

- 10** To send notification when the job completes, in the **Properties** pane, under **Settings**, click **Notification**.

Select the appropriate options.

See [“Notification options for jobs”](#) on page 682.

- 11** If you want to run the job now, click **Run Now**. Otherwise, in the **Properties** pane, under **Frequency**, click **Schedule** to set the scheduling options that you want to use.

See [“Schedule options”](#) on page 356.

Selections options to delete items from the archives

You can specify the criteria to search for the items that you want to delete from the archives.

See [“Deleting items from the archives”](#) on page 1475.

Table N-30 Selections options to delete items from the archives

Item	Description
Include subdirectories	Selects the contents of all the subfolders when a directory is selected.
Show file details	Displays the details about the files that can be deleted from the archives.
Preview pane	Displays the preview pane at the bottom of the dialog box.
Search archives	Enables you to find the archived items that you want to delete from the archives.
View by Resources	Displays archived data by the resource from which it was archived. This feature is useful for finding the files that were located on a certain server or workstation.

General options to delete items from the archives

You can enter a name for the job to delete items from the archives, and select a job priority.

See [“Deleting items from the archives”](#) on page 1475.

Table N-31 General options to delete items from the archives

Item	Description
Job name	Displays the name of the job.
Job priority	Displays the priority of the access to the devices for this job. See “About job priority” on page 195.

About deleting archived data from its original location

When you create a vault store, you can specify when to delete the archived data from its original location.

You can let Backup Exec do one of the following:

- Delete the item from its original location immediately after it is archived.
If the data is lost before the vault store is backed up, the only version of the data is on the backup set.
- Delete it after the vault store is backed up and the next archive job runs.

If Backup Exec deletes an item immediately after it is archived, the item is deleted from its original location after the archive job has successfully completed. If the item is modified after it is archived but before it is backed up, then it is not deleted from its original location.

See [“Editing or viewing vault store properties”](#) on page 1447.

See [“Creating a vault store in the Archiving Option”](#) on page 1446.

See [“Vault store properties”](#) on page 1448.

About backing up Archiving Option components

You can select any or all of the Archiving Option components for backup. If you select all of the components for backup in the same job, recovery time is faster. However, if you create multiple backup jobs for the components, the backup jobs run faster.

The Archiving Option components that you can select in the backup selections view are described in the following table, along with recommendations for backup:

Table N-32 Backing up Archiving Option components

Component	Description
Archiving Option Components	<p>Archiving Option Components contain all of the components that are associated with the Archiving Option. Symantec recommends that you select Archiving Option Components to back up all of the Archiving Option environment.</p>
Backup Exec Archiving Site	<p>The Backup Exec Archiving Site is a logical representation of an installation of the Archiving Option. A media server can have only one Archiving Site. If you select this component for backup, the Directory database is also automatically backed up.</p>
Directory database	<p>The Directory database is a Microsoft SQL Server database that contains configuration data and information about the archives.</p> <p>After the database is populated, the amount of data in the Directory database changes very little over time.</p> <p>You should back up the Directory database after you add or remove any Archiving Option component. You should also back up the Directory database if you change the location of any component. Configuration changes can include creating vault stores, creating vault store partitions, and changing vault store partition statuses.</p>
Index location	<p>The index location stores all of the archived data content that is indexed to enable fast searching and retrieval of archived items. The indexing data is stored in index files in the location that is specified when the Archiving Option is installed.</p> <p>You should back up the index location on a regular basis.</p>

Table N-32 Backing up Archiving Option components (*continued*)

Component	Description
Vault store group	<p>The vault store group is a logical entity. If you select it for backup, all of the vault databases, vault store partitions, and the fingerprint databases are backed up. Because these components are closely related, you should consider selecting the vault store group to back up all of these components together.</p>
Fingerprint databases	<p>The fingerprint databases contains the single instance storage-related information for all of the vault stores in the vault store group.</p> <p>If you enable single instance storage of archived items, you should back up the fingerprint databases on a regular basis.</p> <p>See “About single instance storage of archived items” on page 1494.</p>
Vault store	<p>The vault store is a logical entity. If you select it for backup, all of the vault databases and the vault store partitions are backed up.</p>
All partitions	<p>A vault store partition represents the physical location where the archived items are stored. A vault store can contain one or more vault store partitions. If you select All Partitions for backup, then all of the vault store partitions in the vault store are selected for backup.</p> <p>Note: When you back up an open partition, the vault store database is automatically backed up.</p> <p>You should back up the vault store partitions on a regular basis.</p> <p>See “About vault store partitions in the Archiving Option” on page 1451.</p>

Table N-32 Backing up Archiving Option components (*continued*)

Component	Description
Vault store databases	<p>The vault store databases are the Microsoft SQL Server databases that contain configuration data and information about the archives. Each vault store has an associated database. Each vault store database contains an entry for each item that is archived in the associated vault store. If an item is deleted from the archive, then the references to it are deleted from the vault store database.</p> <p>You should back up the vault store databases on a regular basis.</p>

You can also back up and restore the Archiving Option components from a remote media server on which license keys are not installed.

See [“About backing up and restoring the Archiving Option components from a remote media server”](#) on page 1491.

See [“Backing up Archiving Option components”](#) on page 1481.

See [“Editing default settings for archive jobs”](#) on page 1494.

See [“About consistency checks for Archiving Option databases”](#) on page 1480.

About consistency checks for Archiving Option databases

Backup Exec automatically checks the physical consistency of an Archiving Option database before a backup job and after a restore job. Any consistency check failures are reported in the Backup Exec job log. Backup Exec uses Microsoft SQL Server's utility Physical Check Only for consistency checks of Archiving Option databases.

For more information about the Physical Check Only utility, see the Microsoft SQL Server documentation.

About disabling backup mode for Archiving Option components

When you back up the Directory database, ensure that the Archiving Option components are not in backup mode.

See [“Editing or viewing vault store properties”](#) on page 1447.

If a component is in backup mode, you must take it out of backup mode by running the task **Disable Backup Mode on Archiving Option entities** in the Backup Exec Utility.

See [“Running the Backup Exec Utility for an Archiving Option component”](#) on page 1491.

Backing up Archiving Option components

You can back up all Archiving Option components, or select the components individually for backup.

See [“About backing up Archiving Option components”](#) on page 1477.

Note: If a backup job for an Archiving Option component runs at the same time that an archive job runs, the archive job fails.

To back up Archiving Option components

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Backup Selections** pane, under the media server on which the Archiving Option is installed, expand **Archiving Option Components**.
- 4 Do one of the following:

To back up all of the components that are associated with the Archiving Option

Do the following in the order listed:

- Ensure that the Archiving Option components are not in backup mode. See [“About disabling backup mode for Archiving Option components”](#) on page 1480.
- Select **Archiving Option Components**.

To back up the individual components that are associated with the Archiving Option

Do the following in the order listed:

- Expand **Archiving Option Components**.
- Expand the components.
- Select the components that you want to back up.
See [“About backing up Archiving Option components”](#) on page 1477.
- If you select the Directory database for backup, ensure that the Archiving Option components are not in backup mode.
See [“About disabling backup mode for Archiving Option components”](#) on page 1480.

- 5 In the **Properties** pane, under **Settings**, click **Archive**.
- 6 Select a backup method.
See [“Backup job properties for archive jobs”](#) on page 1482.
- 7 In the **Properties** pane, select other backup options as appropriate.
See [“Creating a backup job by setting job properties”](#) on page 332.

Backup job properties for archive jobs

You can select a backup method for an Archiving Option component.

See [“Backing up Archiving Option components”](#) on page 1481.

For the Directory database, Backup Exec performs a full backup instead of a differential backup, even if you select a differential backup method. When you select the incremental backup method for any of the databases, the transaction logs are backed up and then truncated.

Note: When you select the Backup Exec Archiving Site for backup, the Directory database is also backed up.

See [“About backup methods”](#) on page 274.

See [“About consistency checks for Archiving Option databases”](#) on page 1480.

About restoring an Archiving Option component

You can restore any of the following Archiving Option components:

- Directory database
- Vault store databases
- Fingerprint databases
- Vault store partition
- Index location

Review the scenarios in the following table to find the best procedure to restore an Archiving Option component.

Table N-33 Methods for restoring an Archiving Option component

Method	More information
If a data loss occurs, and you want to restore an Archiving Option component to the same location	See “Restoring an Archiving Option component” on page 1483.
If hardware fails and a data loss occurs, and you want to restore an Archiving Option component to a different location	See “About redirecting the restore of Archiving Option components” on page 1485.
If you want to move components to new hardware, such as to a new SQL Server or a new disk	See “About moving Archiving Option components to a new location” on page 1498.

See [“About consistency checks for Archiving Option databases”](#) on page 1480.

Restoring an Archiving Option component

If a data loss occurs, you can restore one or more Archiving Option components to the same location where they originally existed.

If you restore multiple components that include the Directory database, Symantec recommends that you use a separate job to restore the Directory database first. Then, create one job for all of the remaining Archiving Option components that you want to restore.

To restore an Archiving Option component

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.

- 3 In the **View by Resource** pane, select the backup sets that contain the data that you want to restore.
- 4 (Optional) If you restore a database, do the following in the order listed:
 - On the **Properties** pane, under **Settings**, click **Archive**.
 - Select the options as appropriate.
See “[Restore job properties for Archiving Option databases](#)” on page 1484.
- 5 Complete other properties as appropriate.
See “[Restoring data by setting job properties](#)” on page 605.
- 6 Click **Run Now**, or schedule a time to run the restore job.
See “[Schedule options](#)” on page 356.
- 7 If you restore the Archiving Option database, restart the Enterprise Vault services to re-establish the connection with the Directory database. Do the following in the order listed:
 - Click **Tools > Backup Exec Services**.
 - On the Backup Exec Services Manager, click **Restart all services**.

Restore job properties for Archiving Option databases

You can set options to restore Archiving Option databases to the same location where they originally existed.

Note the following items when you restore Archiving Option databases:

- You should not perform any other archive-related operations when you restore Archiving Option databases.
- You should restart the Enterprise Vault services after you restore the Archiving Option Directory database.

See “[Restoring an Archiving Option component](#)” on page 1483.

Note: You should not perform any other archive-related operations when you restore Archiving Option databases.

Table N-34 Archive restore job properties

Item	Description
Leave the database ready to use; additional transaction logs or differential backups cannot be restored	<p>Lets the restore operation roll back all incomplete transactions when you restore the last full, differential, or log backup. After the restore operation, the database is ready for use. If you do not select this option, the database is left in an intermediate state and is not usable.</p> <p>If you select this option, you cannot continue to restore backups. You must restart the restore operation from the beginning.</p> <p>See “About restoring SQL databases and file groups” on page 1293.</p>
Leave the database nonoperational; additional transaction logs or differential backups can be restored	<p>Indicates that you have additional differential or transaction log backups to be restored in another restore job.</p> <p>If you restore the vault store database, you may be prompted to stop the Enterprise Vault Storage Service before you can continue the restore job.</p> <p>See “About Enterprise Vault services for the Archiving Option” on page 1422.</p>

See [“About consistency checks for Archiving Option databases”](#) on page 1480.

About redirecting the restore of Archiving Option components

You can run a redirected restore job if you want to restore one or more Archiving Option components to a different location. The procedures for redirecting restore jobs include running tasks in Backup Exec Utility to update the new locations of the restored components. Review the procedures before you create a redirected restore job.

The following table lists the possible scenarios and the associated redirected restore solutions for Archiving Option components.

Table N-35 Redirected restore solutions for Archiving Option components

Scenario	Solution
The SQL Server that hosts the databases fails and a data loss occurs.	Redirect the restore of the Archiving Option databases to a new SQL Server. See “Redirecting a restore of the Archiving Option databases” on page 1486.
The local drive or network share that hosts the vault store partition fails and a data loss occurs.	Redirect the restore of a vault store partition to a different path in a local drive or network share. See “Redirecting the restore of an Archiving Option vault store partition” on page 1488.
The disk that contains the index files fails and a data loss occurs.	Redirect the restore of the index files to a new location. See “Redirecting the restore of Archiving Option index files” on page 1489.

See [“About restoring an Archiving Option component”](#) on page 1483.

Redirecting a restore of the Archiving Option databases

You can redirect the restore of the Archiving Option databases to a new SQL Server.

Symantec recommends that you use one job to restore all of the backup sets for the Directory database. If necessary, you can use multiple jobs to restore all of the backup sets. If you use multiple jobs, ensure that you leave the Directory database ready to use before you run the **Change Database Location** task in Backup Exec Utility.

Note: All Archiving Option databases must be on the same SQL Server. If you redirect the restore of one of the databases, you must restore all of the databases to the same location.

See [“About redirecting the restore of Archiving Option components”](#) on page 1485.

To redirect a restore of the Archiving Option databases

- 1 Create a restore job.
See [“Restoring data by setting job properties”](#) on page 605.
- 2 Select the appropriate full backup set and any related differential and incremental backup sets to restore the Directory database.
See [“About selecting data to restore”](#) on page 626.
- 3 Enter other information on the **Restore Job Properties** dialog box as appropriate.
See [“Restoring data by setting job properties”](#) on page 605.
- 4 In the **Properties** pane, under **Settings**, click **Archive**.
- 5 Select **Leave the database ready to use; additional transaction logs or differential backups cannot be restored**.
- 6 In the **Properties** pane, under **Destination**, click **Archive Redirection**.
- 7 Click **Redirect to a new Microsoft SQL Server**.
- 8 Type the path to the SQL Server to which you want to redirect the restore.
- 9 Click **Run Now**, or schedule a time to run the redirected restore job.
See [“Scheduling jobs”](#) on page 356.
- 10 When the redirected restore job is complete, start the **Backup Exec Utility**.
See [“Running the Backup Exec Utility for an Archiving Option component”](#) on page 1491.
- 11 In the **Backup Exec Utility** task pane, under **Archiving Option Tasks**, click **Change Database Location**.
- 12 In **Destination SQL server instance**, type the name of the new SQL Server.
- 13 Click **OK**.
- 14 After the operation completes, exit the Backup Exec Utility.
- 15 Create a restore job.
See [“Restoring data by setting job properties”](#) on page 605.
- 16 Select the appropriate full backup set and any related differential and incremental backup sets to restore the vault store databases and the fingerprint databases.

- 17 Enter other information on **Restore Job Properties** as appropriate. Do not configure information on the **Archive Redirection** dialog box.
- 18 Click **Run Now**, or schedule a time to run the redirected restore job.
See [“Schedule options”](#) on page 356.

Redirecting the restore of an Archiving Option vault store partition

You can redirect the restore of a vault store partition to a different path in a local drive or network share.

If you restore a vault store partition that has an **Open** status, its vault store database is automatically restored.

See [“About redirecting the restore of Archiving Option components”](#) on page 1485.

If a vault store partition needs more disk space, you can create a new partition.

See [“Creating a vault store partition”](#) on page 1451.

To redirect the restore of an Archiving Option vault store partition

- 1 Start the **Backup Exec Utility**.
See [“Running the Backup Exec Utility for an Archiving Option component”](#) on page 1491.
- 2 In the **Backup Exec Utility** task pane, under **Archiving Option Tasks**, click **Change Vault Partition Path**.
- 3 Select the name of the vault store partition.
- 4 In **New Vault Store Partition Path**, type the new path to which you want to restore the vault store partition.
- 5 Ensure that **Move Vault Store Partition Files** is not selected.
- 6 Click **OK**.
- 7 On the Backup Exec Administration Console, create a restore job.
See [“Restoring data by setting job properties”](#) on page 605.
- 8 Select the appropriate full backup set and any related differential and incremental backup sets to restore the vault store partition.

- 9 In **Restore Job Properties**, select other restore options as appropriate. Do not configure information on the **Archive Redirection** dialog box.
See [“Restoring data by setting job properties”](#) on page 605.
- 10 Click **Run Now** or schedule a time to run the redirected restore.
See [“Schedule options”](#) on page 356.

Redirecting the restore of Archiving Option index files

You can redirect the restore of the index files to a new location.

Note: You must locate the index files on a local NTFS drive.

See [“About redirecting the restore of Archiving Option components”](#) on page 1485.

To redirect the restore of Archiving Option index files

- 1 Start the **Backup Exec Utility**.
See [“Running the Backup Exec Utility for an Archiving Option component”](#) on page 1491.
- 2 In the **Backup Exec Utility** task pane, under **Archiving Option Tasks**, click **Change Index Location**.
- 3 In **New Index Location**, type the new path to which you want to restore the index files.
- 4 Ensure that **Move Index Files** is not selected.
- 5 Click **OK**.
- 6 On the Backup Exec Administration Console, create a restore job.
See [“Restoring data by setting job properties”](#) on page 605.
- 7 Select the appropriate backup sets to restore the index files.
- 8 In the **Restore Job Properties** pane, select other restore options as appropriate.
See [“Restoring data by setting job properties”](#) on page 605.
- 9 In the **Properties** pane, under **Destination**, click **Archive Redirection**.
- 10 Click **Restore index files to a new location**.
- 11 Type the path of the new location to which you want to restore the index files.
- 12 Click **Run Now** or schedule a time to run the redirected restore job.
See [“Schedule options”](#) on page 356.

Archive redirection options for Archiving Option components

You can redirect the restore of the Archiving Option Directory database and redirect the restore of index files.

The procedures for redirecting restore jobs include running tasks in Backup Exec Utility to update the new locations of the restored components. Review the procedures before you create a redirected restore job.

See [“About redirecting the restore of Archiving Option components”](#) on page 1485.

Note: You can redirect the restore of a vault store partition by running a task in Backup Exec Utility.

See [“Redirecting the restore of an Archiving Option vault store partition”](#) on page 1488.

Table N-36 Archive redirection options for Archiving Option components

Item	Description
Redirect to a new Microsoft SQL Server	<p>Redirects the restore of the Archiving Option Directory database to a different SQL Server.</p> <p>After you redirect the restore of the Directory database, you must run a procedure that is in a separate program called Backup Exec Utility. The procedure in Backup Exec Utility updates the Directory database with information about the new location of the database. Review the procedure for redirecting the restore of a database before you create the job.</p> <p>See “Redirecting a restore of the Archiving Option databases” on page 1486.</p>
Server	<p>Lets you specify the name of the server to which you want to redirect the restore job.</p>
Instance	<p>Lets you specify the name of the instance of the SQL Server to which you want to redirect the restore job.</p>

Table N-36 Archive redirection options for Archiving Option components
(continued)

Item	Description
Restore index files to a new location	<p>Redirects the restore of the index files to a new location.</p> <p>To redirect the restore of index files to another location, you must first run a procedure that is in a separate program called Backup Exec Utility. The procedures in Backup Exec Utility update the Directory database with information about the new location of the index files.</p> <p>See “Redirecting the restore of Archiving Option index files” on page 1489.</p>
Path	<p>Lets you specify the path name to which you want to redirect the restore job for the index files.</p>

Running the Backup Exec Utility for an Archiving Option component

You must run the Backup Exec Utility to complete some operations for an Archiving Option component.

See [“About redirecting the restore of Archiving Option components”](#) on page 1485.

To run the Backup Exec Utility

- 1 From the Backup Exec installation directory, double-click **BEUtility.exe**.
- 2 In the **Properties** pane, under **Archiving Option Tasks**, click the appropriate task.
- 3 Click **Help** for information about a task.

About backing up and restoring the Archiving Option components from a remote media server

You can back up and restore the Archiving Option components from a remote media server on which license keys are not installed. You can also edit the backup job default settings for the Archiving Option components.

The remote media server that you use to back up the Archiving Option components does not require license keys for the following options:

- File System Archiving Option
- Exchange Mailbox Archiving Option

You must provide the credentials of the Backup Exec service account on the media server on which the Archiving Option is installed.

See [“About backing up Archiving Option components”](#) on page 1477.

See [“Creating a backup job by setting job properties”](#) on page 332.

See [“Restoring data by setting job properties”](#) on page 605.

See [“Editing backup job default settings for Archiving Option components from a remote media server”](#) on page 1492.

Editing backup job default settings for Archiving Option components from a remote media server

You can edit the default backup method for the Archiving Option components from a remote media server.

See [“About backing up and restoring the Archiving Option components from a remote media server”](#) on page 1491.

To edit backup job default settings for the Archiving Option

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Archive**.
- 3 Enter the appropriate information.

See [“Backup job default settings for the Archiving Option”](#) on page 1492.

Backup job default settings for the Archiving Option

You can select a default backup method for the Archiving Option components from a remote media server.

See [“Editing backup job default settings for Archiving Option components from a remote media server”](#) on page 1492.

For the Directory database, Backup Exec performs a full backup instead of a differential backup, even if you select the differential backup method.

When you select the incremental backup method for the SQL databases, the transaction logs are backed up and then truncated.

Preventing the deletion of expired archived items from an archive

Backup Exec deletes the archived items that have expired retention periods from a specific archive. You can clear this option to prevent Backup Exec from deleting expired archived items.

See [“About archives in the Archiving Option”](#) on page 1453.

To prevent the deletion of expired archived items from an archive

- 1 On the navigation bar, click **Devices**.
- 2 Expand the vault store that contains the archive.
- 3 Click **Archives**, and then in the right pane, right-click the archive.
- 4 Click **Properties**.
- 5 Uncheck the option **Automatically delete archived items that have expired retention periods**.
- 6 Click **OK**.

About synchronizing archive permissions and settings

Backup Exec runs a daily synchronization task for the Exchange Mailbox Archiving Option. Synchronization associates the correct archive settings to each mailbox in all of the mailbox groups. This task also ensures that archive permissions are synchronized with the mailbox permissions for each mailbox that is archived.

For the File System Archiving Option, archive permissions are synchronized with share and folder permissions for each file that is archived.

You can specify the time of day to run this operation. Symantec recommends that you schedule the archive jobs to run at a different time than the synchronization operation.

An alert is sent to the administration console when the synchronization operation completes. The alert displays the summary statistics of the operation, and contains a link to the operation's job log.

Note: The Exchange Servers and file system servers must be online and accessible by the media server for synchronization to occur.

See [“Editing default settings for archive jobs”](#) on page 1494.

About single instance storage of archived items

Single instance storage of the archived items lets Backup Exec identify the shareable parts of an item. An example of a shareable part is a message attachment or the contents of a document. Backup Exec then stores the parts separately, and only once. When Backup Exec identifies a shareable part that is already stored in a vault store, it references the stored shareable part instead of archiving it again.

If single instance storage is enabled, items are shared within and across vault stores and vault store partitions. The vault store partitions may be on different device types. Shareable parts of a message that exceed the single-instance threshold of 20 KB are shared. These parts include attachments and message bodies. The user information and the shareable parts that are under the single instance storage threshold are not shared.

Enabling this option can provide a significant reduction in the storage space that is required for archived items. If you enable single instance storage, you should back up the fingerprint databases. Single instance storage-related information is contained in the fingerprint databases for all of the vault stores.

See [“Enabling single instance storage of archived items”](#) on page 1494.

Enabling single instance storage of archived items

You can enable single instance storage of archived items.

See [“About single instance storage of archived items”](#) on page 1494.

To enable single instance storage of archived items

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Archive**.
- 3 Ensure that **Enable single instance storage of archived items** is selected.

See [“Archive job default settings”](#) on page 1495.

Editing default settings for archive jobs

You can change the default settings for all archive jobs.

To edit default settings for archive jobs

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Archive**.

- 3 Select the appropriate options.
See [“Archive job default settings”](#) on page 1495.
- 4 Click **OK**.

Archive job default settings

You can use the default settings that Backup Exec sets during installation for the Archiving Option. Or you can choose your own default settings.

See [“Editing default settings for archive jobs”](#) on page 1494.

Table N-37 Archive job default settings

Item	Description
Backup method	<p>Displays the backup method to use to back up the Archiving Option components.</p> <p>See “About backup methods” on page 274.</p> <p>For the Directory database, Backup Exec performs a full backup instead of a differential backup, even if you select a differential backup method.</p> <p>When you select the incremental backup method for the SQL databases, the transaction logs are backed up and then truncated.</p> <p>See “About backing up Archiving Option components” on page 1477.</p>
Allow archiving from backup data that is on a tape device	<p>Lets Backup Exec archive the data from the backup sets that are on tapes.</p> <p>The tapes that contain the backup data that you want to archive must be available to the media server. The media server must have access to the tape drive or a robotic library slot. Otherwise, the archive job completes with exceptions.</p>

Table N-37 Archive job default settings (*continued*)

Item	Description
<p>Archive from encrypted backup data</p>	<p>Lets Backup Exec archive the data from the backup sets that are encrypted. The archived data is stored as decrypted data in the vault store. The data in the backup set remains encrypted.</p> <p>This option is not checked by default.</p> <p>Only common encryption keys can be used to decrypt a backup set during an archive job. If a restricted key is used, then eligible items in the backup set are not archived.</p>
<p>Archive backup data that was created in the last x days</p>	<p>Lets Backup Exec archive the data only from the backup sets that are as old as the specified number of days.</p> <p>The default number of days is 30.</p> <p>Use this option to limit Backup Exec to relevant backup sets in which to find eligible data to archive.</p> <p>Note: Backup Exec searches the backup sets of the specified server to find the data to archive. If backup jobs use the same selection list, then Backup Exec archives the data from the latest full backup and any subsequent incremental or differential backups.</p>
<p>Default retention category</p>	<p>Displays the retention category that applies to the Backup Exec archive jobs by default. A retention category specifies the period of time for which you want to keep items in the archives.</p> <p>You can edit a retention category to change the retention period.</p> <p>See “Editing a retention category ” on page 1458.</p> <p>The default retention category specifies a retention period of infinite.</p> <p>See “About retention categories for archived items” on page 1457.</p>

Table N-37 Archive job default settings (continued)

Item	Description
New	Lets you create a new retention category that you can apply to the Backup Exec archive jobs. See “Retention category properties” on page 1459.
Delete archived items that have expired retention periods every day at	Indicates the time at which Backup Exec searches the vault store partitions to delete the archived items that have expired retention periods. The default time is 4:00 A.M. For individual archives, you can prevent Backup Exec from automatically deleting the expired archived items. See “Preventing the deletion of expired archived items from an archive” on page 1493.
Synchronize archive permissions and mailbox group members every day at	Indicates the time at which Backup Exec synchronizes the correct archive settings and archive permissions to each mailbox in all of the mailbox groups. The default time is 3:00 A.M. See “About synchronizing archive permissions and settings” on page 1493.

Table N-37 Archive job default settings (*continued*)

Item	Description
Enable single instance storage of archived items	<p>Lets Backup Exec identify the shareable parts of an item, such as a message attachment or the contents of a document. Backup Exec then stores the parts separately, and only once. When Backup Exec identifies a shareable part that is already stored in a vault store, it references the stored shareable part instead of archiving it again.</p> <p>Enabling this option can provide a significant reduction in the storage space that is required for archived items.</p> <p>See “About single instance storage of archived items” on page 1494.</p> <p>If you enable this option, you should back up the fingerprint databases. Single instance storage-related information is contained in the fingerprint databases for all of the vault stores.</p> <p>See “About backing up Archiving Option components” on page 1477.</p>
Manage Mailbox Groups	<p>Lets you create or edit mailbox groups.</p> <p>See “About Exchange mailbox groups in archive jobs” on page 1461.</p>

About moving Archiving Option components to a new location

You can use the Backup Exec Utility to move Archiving Option components to a new location. Ensure that no other archive-related operations are running when you move a component.

If you must move a component because the hardware that hosts the component has failed, you should use a redirected restore job.

See [“About redirecting the restore of Archiving Option components”](#) on page 1485.

Table N-38 Moving Archiving Option components to a new location

Component	More information
Index location	<p>You can move an index location if the disk where the index files are stored runs out of space.</p> <p>Use the Change Index Location task in Backup Exec Utility.</p>
Databases	<p>You can move databases to a different SQL Server. For example, you can move the databases if the current SQL Server becomes overloaded.</p> <p>Use the Change Database Location task in Backup Exec Utility.</p>
Vault store partitions	<p>You can move vault store partitions if you must remove the current drive or network share that contains the partition.</p> <p>Note: If a vault store partition only requires more disk space, you can create a new partition and designate it as open.</p> <p>See “Creating a vault store partition” on page 1451.</p> <p>Use the Change Vault Store Partition Path task in Backup Exec Utility.</p>

See [“Running the Backup Exec Utility for an Archiving Option component”](#) on page 1491.

Troubleshooting archive jobs

If there are issues with archive jobs, you can find information in the following sources:

- Backup Exec job logs.
See [“Viewing the properties for completed jobs”](#) on page 572.
- Enterprise Vault event log that is located in the Windows Event Viewer.
See [“Viewing the Enterprise Vault event log for Archiving Option events”](#) on page 1500.
- Backup Exec diagnostic utilities.
See [“About the Backup Exec diagnostic application”](#) on page 804.

An Exchange Mailbox Archiving Option job may not find data to archive for the following reasons:

- Only back up sets for which the Granular Recovery Technology option was enabled and that reside on backup-to-disk folders can be archived.
- The associated Exchange mail stores may not be backed up, or the mailbox or user may have been deleted in the last 14 days.

A File System Archiving Option job can only find data to archive if backup sets are on disk, or if the option **Allow archiving from backup data that is on a tape device** is selected.

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1415.

Viewing the Enterprise Vault event log for Archiving Option events

You can view the Windows Event Viewer to review the Enterprise Vault event log for information about Archiving Option events. Enterprise Vault generates many log entries. You must take some action to make sure that the log files do not grow too large. For information on how to control the log file size, see the Windows Event Viewer help.

Reports for the Archiving Option

The reports in the following table are available to help you monitor your Archiving Option environment.

See [“About reports in Backup Exec”](#) on page 690.

Table N-39 Reports for the Archiving Option

Report	Description
Vault Store Usage Summary	Displays the archived items that are in each vault store and the total size of the vault store.
Vault Store Usage Details	Displays the archives that are in each store and the size of each archive.
File System Archive Settings	Displays the archive settings that are applied to archive selections for each server.
Exchange Mailbox Group Archive Settings	Displays the archive settings that are applied to mailbox groups in each domain.

Table N-39 Reports for the Archiving Option (*continued*)

Report	Description
Archive Selections by Archive Rules and Retention Categories	Displays the archive rules and retention categories that are applied to each archive selection.
Archive Job Success Rate	Displays the number of archive jobs that ran successfully.
Failed Archive Jobs	Displays a list of archive jobs that failed recently.
Overnight Archive Summary	Displays a summary of archive jobs for the last 24 hours.

Symantec Backup Exec Central Admin Server Option

This appendix includes the following topics:

- [How CASO works](#)
- [How CASO and the Shared Storage Option work together](#)
- [Requirements for installing CASO](#)
- [How to choose the location for CASO device and media data](#)
- [Installing the CASO central administration server](#)
- [Installing a managed media server from the central administration server in CASO](#)
- [About upgrading an existing CASO installation](#)
- [Changing a Backup Exec media server to a central administration server](#)
- [Changing a media server to a managed media server](#)
- [Changing a managed media server to a stand-alone media server](#)
- [Running the Backup Exec Utility for CASO operations](#)
- [Uninstalling Backup Exec from the central administration server in CASO](#)
- [Uninstalling Backup Exec from a managed media server](#)
- [About configuring CASO](#)

- [How to use media server pools in CASO](#)
- [About copying jobs instead of delegating jobs in CASO](#)
- [Requirements for duplicate backup data and synthetic backup jobs in CASO](#)
- [How centralized restore works in CASO](#)
- [Media Servers view in CASO](#)
- [Pausing a managed media server in CASO](#)


How CASO works

The Symantec Backup Exec 2010 - Central Admin Server Option (CASO) is installed as a separate, add-on component of Backup Exec 2010. If your organization includes more than one Backup Exec media server, you can benefit from using CASO.

When CASO is installed in a Backup Exec environment, one media server, known as the central administration server, delegates jobs to managed media servers across the network. Job delegation is the automatic load balancing of jobs across available managed media servers in the CASO environment.

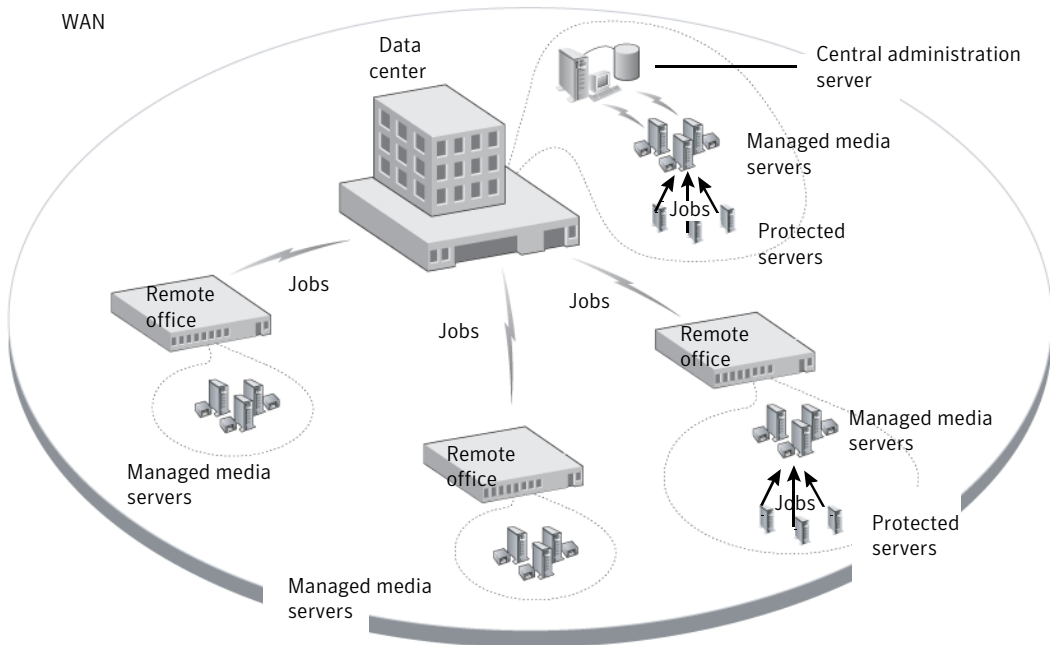
The following graphic shows a local area network (LAN) environment with a central administration server and several managed media servers.

Figure 0-1 CASO-configured Backup Exec environment - LAN



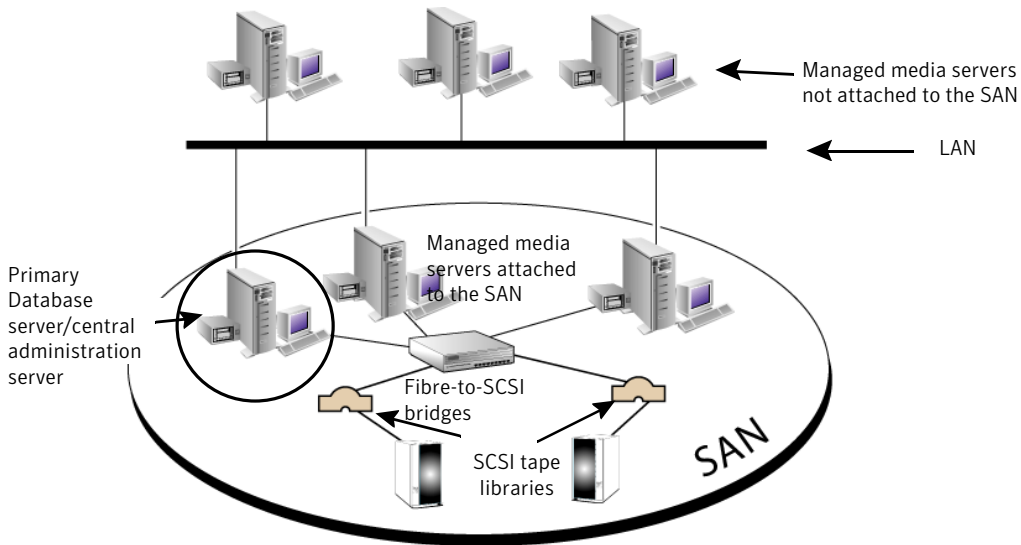
The same communications that occur over a LAN between the central administration server and the managed media servers take place over a WAN.

Figure 0-2 CASO-configured Backup Exec environment - WAN



The following graphic shows CASO installed in a Backup Exec SAN Shared Storage Network environment.

Figure 0-3 CASO-configured Backup Exec environment - SAN Shared Storage Network



All backup information in the CASO environment can be centralized on the central administration server. You can filter this information to display it for each managed media server, or to display it by media server pools (groups of managed media servers).

The managed media servers are managed by the central administration server. They perform the actual processing of backup and restore jobs. You create jobs on the central administration server by associating policies and selection lists. Then, you target the jobs to run on a managed media server or in a media server pool. The jobs are delegated, or load-balanced, across the available storage devices on the managed media server or media server pool. Multiple media servers can share a device when sharing is enabled. Centralized restore jobs can also be delegated to managed media servers.

See [“About sharing storage”](#) on page 444.

Additionally, the central administration server can function as a managed media server and process delegated jobs. A managed media server can also run jobs that are created locally at its local Administration Console.

CASO includes the following additional functionality:

- Centralized operations, such as backup and restore jobs, job monitoring, and reporting.

- Centralized information, such as device and media data, job logs, job histories, and alerts.
- Centralized creation of policies and selection lists, and the associations between them. Also, the ability to copy settings to a managed media server for local job operations. A persistent network connection between the central administration server and the managed media server is not necessary.
- Media server pools, so that operations can be performed on specific groups of managed media servers and their attached storage devices.
- A configurable catalog location to enable centralized, distributed, or replicated catalogs.

See “[How to choose the location for CASO device and media data](#)” on page 1509.

See “[Installing the CASO central administration server](#)” on page 1512.

See “[About upgrading an existing CASO installation](#)” on page 1521.

How CASO and the Shared Storage Option work together

The Backup Exec Central Admin Server Option (CAS) and the Shared Storage Option (SSO) can be used together to provide the following:

- The ability to centrally monitor and manage multiple Backup Exec media servers with CASO.
- The ability to share a centralized tape library for LAN-free backup between multiple Backup Exec media servers on a SAN with SSO.

Separate license keys can be purchased and entered for both CASO and SSO for installation on a Backup Exec media server. The CASO server and the SSO Primary server should be installed on the same Backup Exec media server to centrally manage all shared devices. Additional SSO licenses are required for each Backup Exec media server that shares the centralized device on the SAN.

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

Requirements for installing CASO

The system requirements (except RAM) for running CASO follow the minimum requirements for running Backup Exec 2010. However, processor speed, memory, and disk space requirements may increase based on the number of managed media

servers being managed, the number of protected servers being backed up, and the amount of catalog storage space required.

On the computer on which you install the central administration server, 512 MB RAM is required. 1 GB RAM is recommended. Other applications on the media server also require a certain amount of physical RAM to function properly. The requirements for RAM may also increase when the central administration server manages more media servers or tape hardware.

Ensure the following:

- You must have administrative rights on computers on which you want to install the Central Admin Server Option (CASO).
- When installing CASO on media servers in multiple domains, the Backup Exec service account should be in the trusted domain and have administrative rights on all media servers that are to be managed by the central administration server.

If the Backup Exec database for the central administration server is installed on a SQL Server instance on a different computer, the account must be a domain account with local administrative privileges on that computer as well.

- The central administration server and the managed media servers must be part of a domain or domains. The Central Admin Server Option is not supported in a workgroup.
- Use only NetBIOS computer names for managed media servers and central administration servers. You cannot enter fully qualified domain names or IP addresses as server names.

See [“System requirements”](#) on page 114.

See [“Installing the CASO central administration server”](#) on page 1512.

How to choose the location for CASO device and media data

During the installation of the Managed Media Server feature, you are prompted to keep the managed media server’s device and media data either on the central administration server or in a database on the managed media server.

The following table compares how CASO tasks are performed depending on the location of the managed media server’s device and media data:

Table O-1 Comparison of CASO tasks

Task	Device and media data on the central administration server	Device and media data on the managed media server
<p>Delegate jobs from the central administration server to the managed media server</p> <p>See “About job delegation in CASO” on page 1544.</p>	<p>Yes</p>	<p>No</p> <p>Instead, you can create jobs on the central administration server, and then copy them to the managed media server.</p> <p>See “About copying jobs instead of delegating jobs in CASO” on page 1551.</p>
<p>Manage storage devices and media on the managed media server from the central administration server</p>	<p>Yes</p>	<p>No</p>
<p>Hold, delete, run, cancel, and change the priority of copied jobs from the central administration server if the option to monitor jobs is enabled on the managed media server</p> <p>See “About configuring CASO” on page 1529.</p>	<p>Yes</p>	<p>Yes</p>

Table O-1 Comparison of CASO tasks (*continued*)

Task	Device and media data on the central administration server	Device and media data on the managed media server
<p>Monitor jobs that are created on the local managed media server if the option to monitor jobs is enabled on the managed media server</p> <p>See “About configuring CASO” on page 1529.</p>	Yes	Yes
<p>Send job status updates, job logs, and job histories to the central administration server if the option to monitor jobs is enabled on the managed media server</p> <p>See “About configuring CASO” on page 1529.</p>	Yes	Yes
<p>Centralize, distribute, or replicate the catalog</p> <p>See “Changing the CASO catalog location” on page 1542.</p>	Yes	<p>No</p> <p>Only a distributed catalog location can be selected.</p>
<p>Run centralized restore</p> <p>See “How centralized restore works in CASO” on page 1552.</p>	Yes	<p>No</p> <p>You can browse the backup sets and run restore operations for the managed media server from the central administration server.</p>

Note: In a CASO environment, you can add an NDMP Server only to a central administration server or a managed media server on which the device and media database is located.

See [“About upgrading an existing CASO installation”](#) on page 1521.

See [“How CASO works”](#) on page 1504.

See [“Running the Backup Exec Utility for CASO operations”](#) on page 1527.

Installing the CASO central administration server

Before you start the installation, review the information about the location of device and media data.

See [“How to choose the location for CASO device and media data”](#) on page 1509.

During the installation of the Managed Media Server feature, you are prompted to keep the managed media server’s device and media data either on the central administration server or in a database on the managed media server. Your choice affects how you can manage jobs in the CASO environment.

To install the Central Admin Server Option (CASO), install the central administration server first, and then install the managed media servers.

To install the central administration server

- 1 Install Backup Exec and CASO on the server that you want to be the central administration server.
See [“Installing a custom installation of Backup Exec”](#) on page 119.
- 2 Enter the CASO license key when the prompt appears.
- 3 After Backup Exec and CASO are installed on the central administration server, start Backup Exec and verify that the **Media Servers** view displays on the navigation bar.
- 4 From the **Media Servers** view, verify that Central Administration Server displays in the **Media Server Type** column for the media server on which you installed CASO.
- 5 On the task pane, under **Media Server Installation Tasks**, click **Configure managed media server defaults**.

- 6 On the **Managed Media Servers Defaults** dialog box, choose the settings that you want applied to each new managed media server that you install.
See [“Setting defaults for managed media servers”](#) on page 1531.
- 7 Install a managed media server.
See [“Installing a managed media server from the central administration server in CASO ”](#) on page 1513.

Installing a managed media server from the central administration server in CASO

After installing the central administration server, you can push-install a managed media server feature to a stand-alone server.

If the managed media server is not displayed on the **Media Server** view after you follow these instructions, and if your network contains firewalls, you may need to open some ports between the central administration server and the managed media server.

Before you install a managed media server, decide where to locate the device and media database for the managed media server. During the installation of the managed media server, you are prompted to keep the managed media server’s device and media data on the central administration server or in a database on the managed media server. Your choice affects how you can manage jobs in the CASO environment.

See [“How to choose the location for CASO device and media data ”](#) on page 1509.

To push-install a managed media server from the central administration server

- 1 On the central administration server’s navigation bar, click **Media Servers**.
- 2 In the task pane, under **Media Server Installation Tasks**, click **Install additional managed media servers**.
- 3 In the **Remote Server** field, type the name of the managed media server that you want to add, or click **Browse** to locate the server.
- 4 Click **Add to List**.

5 Under **Remote Computer Logon Credentials** , complete the fields as follows:

User Name	Type the user name for an account that has administrative rights on the remote computer.
Password	Type the password for an account that has administrative rights on the remote computer.
Domain	Select the domain in which the remote computer is located.

6 Click **Next**.

7 Do one of the following:

- If you do not have license keys for Backup Exec and its options
- Go to <https://licensing.symantec.com> to activate the product.
After you activate the product, Symantec sends license keys to you. License keys are required to install Backup Exec and its options. You can access the Web site from any computer that has Internet access.
 - When you receive your license keys, go to step 8.

If you have license keys for Backup Exec and its options

Go to step 8.

8 Select one of the following methods to enter license keys:.

- | | |
|------------------------------------|---|
| To enter license keys manually | Do the following in the order listed: <ul style="list-style-type: none">■ Type a license key into the license key field.■ Click Add.■ Repeat for each license key for each option or agent that you want to add. |
| To import license keys from a file | Do the following in the order listed: <ul style="list-style-type: none">■ Click Import from file.■ Select the besernum.xml file. |
| To install an evaluation version | Do the following in the order listed: <ul style="list-style-type: none">■ Leave the license key field blank.■ Proceed to step 9. |

9 Click **Next**.

The license keys you entered are saved to the besernum.xml file.

10 On the **Backup Exec Features** list, select **Managed Media Server**.

See [“About Backup Exec’s standard features”](#) on page 112.

11 Do one of the following:

To change the directory where the Backup Exec files are installed In the **Destination Folder** field, type the name of the directory.

To accept the default directory (recommended) Proceed to step [12](#).

Symantec recommends that you do not select a mount point as the destination directory because if you delete the mount point, Backup Exec is uninstalled.

12 Click **Next**.

13 Provide a user name, password, and domain for an Administrator account that the Backup Exec system services can use, and then click **Next**.

See [“About the Backup Exec service account”](#) on page 106.

14 Select to install a local SQL Express instance or to use an existing instance of SQL Server 2005 (SP3) or SQL Server 2008, and then click **Next**.

See [“About Microsoft SQL Server 2005 Express Edition components installed with Backup Exec”](#) on page 111.

15 Type the name of the central administration server with which this managed media server will communicate.

If you configure a managed media server as a secondary server in the SAN SSO, make the primary server the central administration server. The primary server must be the same server as the central administration server.

16 Select where to keep the device and media data for this managed media server:

On the central administration server

Use this option if you want to do the following:

- Delegate jobs to this managed media server.
- Manage all of the storage devices and media from the central administration server.

On the managed media server

Use this option in the following situations:

- A persistent network connection is not available between the central administration server and the managed media server.
- You want to reduce network traffic slightly because of a low-bandwidth network connection.

You cannot delegate jobs from the central administration server to this managed media server, but you can copy jobs to this managed media server. The copied jobs can then run without a network connection to the central administration server.

See [“How to choose the location for CASO device and media data ”](#) on page 1509.

If you select the SAN SSO option and the managed media server option during installation, all Backup Exec catalog and database functions are centralized by default.

17 Click **Next**.

18 Review the information about device drivers, and then click **Next**.

19 After Backup Exec validates the remote computers, you can change the list in any of the following ways:

To manually add one remote computer

Click **Add**, and then click **Add a Single Server**.

To manually add multiple remote computers

Click **Add**, and then click **Add Multiple Servers with the Same Settings**.

To add multiple remote computers by importing an existing list of computers

Click **Import and Export**, and then select one of the following options:

- Select **Import from File** to enable Backup Exec to add the names of the remote computers from a selected list.
- Select **Import Servers Published to this Media Server** to enable Backup Exec to add the names of all the remote computers that are set up to publish to this media server.

You must enter remote computer logon credentials for the list of remote computers.

To change the product that you selected to install or to change other properties you selected for this installation	Select the remote computer that you want to change, and then click Edit .
To delete a remote computer from the list	Select the remote computer that you want to delete, and then click Delete .
To save this list of remote computers and the associated remote computer logon credentials	Verify that Save the server list for future remote install sessions is checked. This option enables the names of all of the remote computers and their credentials to be added automatically the next time you want to install Backup Exec or options to these remote computers.
To save this list of remote computers to an XML file	Click Import and Export , and then click Export to File . You can select the location to save the XML file. This option is useful if you want to use the same list for multiple media servers. When you import the list, you must re-enter the remote computer logon credentials.
To fix the errors that were located during the validation	Right-click the name of the computer, and then click Fix Errors .
To enable Backup Exec to attempt to re-validate an invalid remote computer	Right-click the name of the computer, and then click Retry Validation .

20 After all of the computers are validated, click **Next**.

21 Read the Backup Exec installation review, and then click **Install**.

22 Click **Next**, and then click **Finish**.

If you did not restart the remote computer, you may need to do it now in order for the configuration to take effect.

23 (Optional) Install device drivers for the tape storage devices that are connected to the server.

See [“About configuring tape devices by using the Tape Device Configuration Wizard”](#) on page 453.

After restarting the managed media server, the Backup Exec central administration server and the managed media server begin communicating with one another. The managed media server defaults that you set from the central administration server are applied.

- 24 On the central administration server's navigation bar, click **Media Servers**.
- 25 Make sure that the managed media server name is displayed on the right pane.

If the managed media server is not displayed on the **Media Server** view, and if your network contains firewalls, you may need to open some ports between the central administration server and the managed media server.

About installing a CASO managed media server across a firewall

A managed media server may be installed outside the firewall that the central administration server is installed in or in a different firewall.

The following rules apply to the managed media servers that are installed across a firewall:

- Port 3527 must be open in both directions to enable communication for the Backup Exec Server service.
- Port 10000 must be open for the Remote Agent for Windows Systems, which allows browsing for remote selections.
- A SQL port must be open in both directions to the central administration server's database to enable database connections.
- A static port must be used.

The Backup Exec SQL instance is configured by default to use a dynamic port. Each time SQL Server is started, the port number can change. You must change the dynamic port to a static port. After you change the configuration of the port from dynamic to static, you must add the static port to the Windows Firewall Exceptions list.

See your Windows operating system documentation.

See ["Changing the dynamic port on the SQL Express instance in CASO to a static port"](#) on page 1518.

See ["Opening a SQL port in CASO for a SQL 2005 or 2008 instance"](#) on page 1520.

Changing the dynamic port on the SQL Express instance in CASO to a static port

You must change the port on which the Backup Exec SQL Express instance for the central administration server is running from a dynamic port to a static port. Then, create an alias for the managed media server to allow it to connect to the SQL port on the central administration server. After changing the port, you must restart the Backup Exec and Microsoft SQL services on the central administration server.

To change the dynamic port for a SQL Express instance to a static port

- 1 On the central administration server, click **Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**.
- 2 Expand **SQL Server 2005 Network Configuration**.
- 3 Click **Protocols for BKUPEXEC**, and then in the right pane, double-click **TPC/IP**.
- 4 On the TCP/IP Properties dialog box, click the IP Addresses tab.
- 5 Under IPAll, in TCP Dynamic Ports, remove the value and leave the field blank.
- 6 Under IPAll, type in a port number for TCP Port.
The port number can be between 1025 and 65535 and must not be in use by another application.
See “[Troubleshooting restore issues](#)” on page 797.
- 7 Under the heading for the specific network interface card that is being used, such as IP1 or IP2, change Enabled from No to Yes.
- 8 Under that same heading, in TCP Dynamic Ports, remove the value of 0, and type the same port number you entered for TCP Port.
- 9 Click **Apply**.
- 10 You must restart the Backup Exec and SQL services.
- 11 Create an alias for the managed media server to allow it to connect to the SQL port on the central administration server.
See “[Creating an alias for a managed media server when a SQL Express instance is used](#)” on page 1519.

Creating an alias for a managed media server when a SQL Express instance is used

You must change the port on which the Backup Exec SQL Express instance for the central administration server is running from a dynamic port to a static port. Then, create an alias for the managed media server to allow it to connect to the SQL port on the central administration server. After changing the port, you must restart the Backup Exec and Microsoft SQL services on the central administration server.

To create an alias when a SQL Express instance is used

- 1 On the managed media server, click **Start > SQL Server Configuration Manager**.
- 2 Expand **SQL Native Client Configuration**.
- 3 Click **Aliases**, and then double-click the alias name that contains the central administration server name and the Backup Exec SQL instance name.
- 4 On the alias properties dialog box, enter the appropriate information as described in the following table:

Alias Name	Type the name of the central administration server and the Backup Exec SQL instance name using the format server name\instance name.
Port No	Type the port number of the remote Backup Exec SQL Server instance that you noted in the previous procedure.
Protocol	Select TCP/IP .
Server	Type the name of the central administration server and the Backup Exec SQL instance name using the format server name\instance name.

- 5 Click **Apply**, and then click **OK**.
- 6 Close the SQL Server Configuration Manager utility.

Opening a SQL port in CASO for a SQL 2005 or 2008 instance

You must find the port number on which the Backup Exec SQL 2005 or 2008 instance for the central administration server is running, and then create an alias for the managed media server.

To open a SQL port for a SQL 2005 or 2008 instance

- 1 On the central administration server, go to `\Program Files\Microsoft SQL Server\80\Tools\Binn` and double-click **svrnetcn.exe**.
- 2 On the General tab, select the Backup Exec SQL instance.
- 3 Under Enabled Protocols, select **TCP/IP**, and then click **Properties**.

- 4 Note the port number that is displayed.
- 5 Create an alias for the managed media server to allow it to connect to the SQL port on the central administration server.

See “[Creating an alias for a managed media server when a SQL 2005 or SQL 2008 instance is used](#)” on page 1521.

Creating an alias for a managed media server when a SQL 2005 or SQL 2008 instance is used

You must find the port number on which the Backup Exec SQL 2005 or 2008 instance for the central administration server is running, and then create an alias for the managed media server.

To create an alias when a SQL 2005 or 2008 instance is used

- 1 On the managed media server, to create an alias for the managed media server, go to `\Windows\System32` and double-click **cliconfg.exe**.
- 2 On the Alias tab, click **Add**.
- 3 In the Server alias field, type:

`server name\instance name`
- 4 Under Network libraries, select **TCP/IP**.
- 5 In the Server name field, type:

`server name\instance name`
- 6 Uncheck **Dynamically determine port**.
- 7 In the Port number field, type the port number of the remote Backup Exec SQL Server instance.

About upgrading an existing CASO installation

In an existing CASO environment, upgrade the central administration server, and then upgrade the managed media servers.

If necessary, you can perform rolling upgrades in the CASO environment. That is, you can upgrade the central administration server from Backup Exec 12.5 to Backup Exec 2010 first, and then upgrade the managed media servers from Backup Exec 12.5 to Backup Exec 2010 over a period of time. You must have the most recent Backup Exec service pack to perform rolling upgrades.

Note: Forward compatibility is not supported in rolling upgrades. Therefore, any system that runs Backup Exec 12.5 cannot protect a system that runs Backup Exec 2010.

Symantec recommends that you do not keep a mix of versions in the CASO installation for an extended time. Key functionality for administering managed media servers is missing in a mixed-version environment, and decreases your ability to properly administer the CASO environment.

Note: If the managed media server has the SAN SSO Option installed, you cannot perform rolling upgrades.

After you upgrade the central administration server to Backup Exec 2010, the following operations are supported on managed media servers that run Backup Exec 12.5:

- Backup
- Restore
- Inventory
- Catalog

See [“About CASO catalog locations”](#) on page 1541.

See [“Changing the CASO catalog location”](#) on page 1542.

See [“Upgrading an existing CASO central administration server”](#) on page 1522.

See [“Upgrading an existing CASO managed media server”](#) on page 1523.

Upgrading an existing CASO central administration server

The central administration server must be upgraded before any managed media servers are upgraded.

See [“About upgrading an existing CASO installation”](#) on page 1521.

Before upgrading Backup Exec, run a database maintenance job to delete job histories and catalogs that you no longer need in order to shorten the upgrade window.

See [“Configuring database maintenance”](#) on page 209.

Note: Symantec recommends that you stop all Backup Exec services on each managed media server before you upgrade the central administration server.

To upgrade an existing central administration server

- 1 Verify that the latest service pack for Backup Exec is installed.
- 2 Place all scheduled jobs on hold on the central administration server and the managed media servers.
See [“Placing all scheduled occurrences of an active job on hold”](#) on page 563.
- 3 Allow all active jobs to complete.
- 4 From the installation media browser, select the option to install Symantec Backup Exec.
- 5 On the Welcome panel, click **Next**.
- 6 Select **I accept the terms of the license agreement**, and then click **Next**.
- 7 Check **Local Install**, and then click **Install Backup Exec software and options**.
- 8 Click **Next**.
- 9 Follow the prompts in the wizard.
- 10 On the Back Up Existing Catalog and Data page, enter or browse to a directory to which all existing catalogs and data will be backed up. The default location is:

```
C:\Program Files\Symantec\Backup Exec\Data
```


If you do not want to keep previous catalogs and data, click **Do not back up previous data and catalogs**.
- 11 Click **Next** to continue.
An upgrade summary is displayed. When the upgrade is complete, communication with the managed media servers is automatically enabled.
- 12 Release all jobs from hold.
See [“Placing all scheduled occurrences of an active job on hold”](#) on page 563.
- 13 Upgrade some or all of the managed media servers.

Upgrading an existing CASO managed media server

The central administration server must be upgraded before any managed media servers are upgraded.

See [“About upgrading an existing CASO installation”](#) on page 1521.

Before upgrading Backup Exec, run a database maintenance job to delete the job histories and the catalogs that you no longer need. This practice shortens the upgrade window.

See [“Configuring database maintenance”](#) on page 209.

Table O-2 Upgrading an existing CASO managed media server

Step	Description
Step 1	Verify that the latest service pack for Backup Exec 12.5 is installed.
Step 2	Pause the managed media server to prevent the central administration server from delegating jobs to it. If jobs are running, let them finish or cancel them before beginning the upgrade.
Step 3	<p>On the managed media server that you want to upgrade, do one of the following:</p> <ul style="list-style-type: none"> ■ Choose where to keep device and media data. See “Installing a managed media server from the central administration server in CASO” on page 1513. ■ Keep the managed media server’s device and media data centralized on the central administration server. To move device and media data to a database on the managed media server at a later time, you must run the Backup Exec Utility. Backup Exec Utility prompts you to retarget jobs and media server pools to use the new device and media data location on the local managed media server. See “Changing a media server to a managed media server” on page 1526.
Step 4	<p>Resume the managed media server.</p> <p>See “Pausing a managed media server in CASO” on page 1561.</p> <p>See “Running the Backup Exec Utility for CASO operations” on page 1527.</p>

Changing a Backup Exec media server to a central administration server

You can change a stand-alone Backup Exec media server to a central administration server.

To change a Backup Exec media server to a central administration server

- 1 On the media server that you want to be the central administration server, start Backup Exec.
- 2 On the **Tools** menu, click **Install Options and License Keys on this Media Server**.
- 3 On the **Welcome** screen, click **Next**.
- 4 Check **Local Install**, and then click **Next**.
- 5 Do one of the following:

If you do not have license keys for Backup Exec and its options

Do the following in the order listed:

- Go to <https://licensing.symantec.com> to activate the product.

After you activate the product, Symantec sends license keys to you. License keys are required to install Backup Exec and its options. You can access the Web site from any computer that has Internet access.

- When you receive your license keys, go to step 6.

If you have license keys for Backup Exec and its options

Go to step 6.

- 6 Select one of the following methods to enter license keys:

To enter license keys manually

Do the following in the order listed:

- Type a license key into the License Key field.
- Click **Add**.
- Repeat for each license key for each option or agent that you want to add.

To import license keys from a file

Do the following in the order listed:

- Click **Import from file**.
- Select the besernum.xml file.

To install an evaluation version Leave the license key field blank.

A license key is not required for a fully functional Trial version.

7 Click **Next**.

The license keys you entered are saved to the besernum.xml file, located in the Windows or WINNT directory.

8 On the **Backup Exec Features** list, under **Backup Exec Options**, select **Central Admin Server Option**.

9 Click **Next**.

10 Enter a user name, password, and domain of an account that has local administrative privileges for the Backup Exec services to use.

11 Click **Next**.

12 Read the Backup Exec installation review, and then click **Install**.

13 Click **Finish**.

Changing a media server to a managed media server

You can change a stand-alone Backup Exec media server to a managed media server.

Note the following exceptions:

- If a central administration server is already set for a managed media server, you must use the Backup Exec Utility to change to another central administration server.
- If the Backup Exec SAN Shared Storage Option is installed, the Set Central Administration Server option is not available on the secondary server.

If the managed media server is not displayed on the Media Server view after you follow these instructions, and if your network contains firewalls, you may need to open some ports between the central administration server and the managed media server.

To change a media server to a managed media server

- 1 Verify that the central administration server is running.
- 2 Start Backup Exec at the stand-alone media server.
- 3 On the **Tools** menu, click **Set Central Administration Server**.

- 4 Enter the name of the central administration server.
An informational alert appears stating that the media server will be managed by the specified central administration server.
- 5 Click **OK**.
- 6 Restart the Backup Exec media server.
- 7 On the central administration server's navigation bar, click **Media Servers**.

Changing a managed media server to a stand-alone media server

You can change a managed media server to a stand-alone media server by deleting it from the Media Servers view.

To change a managed media server to a stand-alone media server

- 1 On the central administration server, from the navigation bar, click **Media Servers**.
- 2 Select the managed media server that you want to delete.
- 3 On the task pane, under **General Tasks**, click **Delete**.

Running the Backup Exec Utility for CASO operations

To move the location of the device and media data, or to set a different central administration server for managed media servers, you must run a separate application called Backup Exec Utility.

Use the Backup Exec Utility only with the guidance of Symantec Technical Support. Improper use of this utility can result in configuration changes that may prevent Backup Exec from running.

To run the Backup Exec Utility

- 1 From the Backup Exec installation directory, in `\Program Files\Symantec\Backup Exec`, double-click **BEUtility**.
- 2 Refer to the help for information about performing tasks. On the Backup Exec Utility menu, click **Help**.

Uninstalling Backup Exec from the central administration server in CASO

Before you uninstall Backup Exec from the central administration server, you must delete all managed media servers from the Media Servers view on the central administration server.

Caution: Failure to uninstall in the following sequence may result in long delays when shutting down Backup Exec services during the uninstall of Backup Exec on the managed media servers.

To uninstall Backup Exec from the central administration server

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Select a managed media server.
- 3 On the task pane, under **General Tasks**, click **Delete**.
- 4 Repeat these steps for each managed media server that is displayed in the **Media Servers** view.
- 5 After deleting all managed media servers, uninstall Backup Exec on the central administration server.

See [“Uninstalling Backup Exec”](#) on page 168.

Uninstalling Backup Exec from a managed media server

You must delete the managed media server from the Media Servers view on the central administration server before you uninstall Backup Exec.

To uninstall Backup Exec from a managed media server

- 1 On the central administration server, from the navigation bar, click **Media Servers**.
- 2 Select the managed media server that you want to delete.
- 3 On the task pane, under **General Tasks**, click **Delete**.
- 4 After deleting the managed media server from the **Media Servers** view on the central administration server, uninstall Backup Exec from the managed media server.

See [“Uninstalling Backup Exec”](#) on page 168.

About configuring CASO

After installing CASO, you can do the following to configure your CASO environment.

Table O-3 Configuring the CASO environment

If you want to	Do this
Set defaults for a managed media server	See “Setting defaults for managed media servers” on page 1531.
Accommodate a low bandwidth network connection or reduce network traffic	See “About reducing network traffic in CASO” on page 1530.
Customize thresholds for unresponsive managed media servers to enable job recovery, and customize how often managed media servers send active job status updates	See “Setting communication thresholds and active job status updates for CASO ” on page 1533.
Customize job log settings and job history information to remain on managed media servers. The information also can be copied and moved automatically to the central administration server.	See “Copying logs and histories to the central administration server” on page 1536.
Change the location of the device and media data	See “Running the Backup Exec Utility for CASO operations” on page 1527.
Delegate jobs from the central administration server to any available devices	See “Creating device pools” on page 516.
Perform operations on a group of managed media servers	See “How to use media server pools in CASO ” on page 1545.
View alerts that are generated on a managed media server	See “How alerts work in CASO” on page 1538.

Table O-3 Configuring the CASO environment (*continued*)

If you want to	Do this
Configure notification when alerts occur	See “About alerts and notification in CASO” on page 1540.
View default error-handling rules for recovering failed jobs	See “About recovering failed jobs in CASO” on page 1560.
Let a delegated backup job use any network interface to access remote agents if the selected network interface is unavailable	See “Enabling managed media servers to use any available network interface card ” on page 1540.

About reducing network traffic in CASO

To accommodate a low bandwidth network connection or to reduce network traffic, you can do the following:

- Reduce the frequency of job status updates that are sent from the managed media servers to the central administration server.
- Prevent jobs that are created on the local managed media servers from being monitored by the central administration server.
- Reduce the frequency that job logs and job histories are sent from the managed media servers to the central administration server.
- Increase the amount of time that Backup Exec waits before changing the media server’s status if the media server becomes unresponsive.
- Keep the catalogs on the managed media server (distributed). If there is a persistent network connection between the central administration server and the managed media server, then you can browse the catalog and perform restore operations from both servers, regardless of the catalog location.

See [“Setting defaults for managed media servers”](#) on page 1531.

See [“Copying logs and histories to the central administration server”](#) on page 1536.

See [“Setting communication thresholds and active job status updates for CASO ”](#) on page 1533.

See [“Changing the CASO catalog location”](#) on page 1542.

Setting defaults for managed media servers

Backup Exec sets communication defaults automatically. However, you can change the default settings.

See [“About reducing network traffic in CASO”](#) on page 1530.

To set defaults for managed media servers

- 1 On the central administration server’s navigation bar, click **Media Servers**.
- 2 Do one of the following:

To set defaults to apply to a managed media server when it is installed On the task pane, under Media Server Installation Tasks, click **Configure managed media server defaults**.

To set defaults for an existing managed media server Right-click the managed media server for which you want to configure settings, and then click **Properties**

To set defaults for a group of managed media servers in a pool Do the following in the order listed:

- Select or create the media server pool that contains the managed media servers to which you want to apply new settings.
- Right-click the media server pool, and then click **Properties**.

- 3 On the **Settings** tab, select the appropriate options.
See [“Default settings for managed media servers”](#) on page 1531.
- 4 Click **OK**.

Default settings for managed media servers

Backup Exec sets communication defaults automatically. However, you can change the default settings.

See [“Setting defaults for managed media servers”](#) on page 1531.

Table O-4 Default settings for managed media servers

Item	Description
Apply these settings to all of the managed media servers in the pool	Sets the defaults for a group of managed media servers in a pool. This option appears only for media server pools.

Table O-4 Default settings for managed media servers (*continued*)

Item	Description
Fast connection with the central administration server	Configures frequent communications between the central administration server and the managed media server. By default, when you choose this setting, job status updates are sent every 10 seconds to the central administration server. Job logs and job histories are sent whenever a job on the managed media server completes.
Slow connection with the central administration server	Configures less frequent communications between the central administration server and the managed media server. By default, when you choose this setting, job status updates are sent every 120 seconds to the central administration server. Job logs and job histories are sent only when a job on the managed media server fails.
Custom settings	<p>Enables the Edit Custom Settings option, which sets specific defaults for job recovery thresholds, for sending job status updates, and for sending job logs and histories.</p> <p>See “Setting communication thresholds and active job status updates for CASO” on page 1533.</p> <p>See “Copying logs and histories to the central administration server” on page 1536.</p>
Edit Custom Settings	<p>Sets specific defaults for job recovery thresholds, for sending job status updates, and for sending job logs and histories.</p> <p>See “Setting communication thresholds and active job status updates for CASO” on page 1533.</p> <p>See “Copying logs and histories to the central administration server” on page 1536.</p>
Monitor jobs that are created on the local managed media server in addition to the jobs that are delegated from the central administration server	<p>Enables you to view the jobs that are created on the local managed media server as well as delegated jobs.</p> <p>You can also hold, delete, run, cancel, and change the priority order of the jobs that are created on or copied to the local managed media server.</p> <p>Jobs that are created from policies on the managed media server cannot be deleted from the central administration server.</p>

Table O-4 Default settings for managed media servers (*continued*)

Item	Description
Display alert when the difference in seconds between the clocks on the managed media server and the central administration server is more than	<p>Enables Backup Exec to create an alert if the clock on the managed media server differs from the clock on the central administration server. An alert is generated when the number of seconds indicated is exceeded.</p> <p>CASO monitors the internal computer clocks on both the managed media servers and the central administration server. If time differences develop between the central administration server and the managed media servers, jobs could run at unexpected times. To prevent problems, the time that is reported on managed media server clocks should match the time that is reported on the central administration server clock. If you receive time difference alerts, reset the managed media server system clock to match the system clock on the central administration server.</p> <p>If you change the system time on either the managed media server or the central administration server, you must restart the Backup Exec services on the system.</p>

Setting communication thresholds and active job status updates for CASO

The communication statuses determine how the central administration server processes current and future jobs that are delegated to an unresponsive managed media server.

See [“What happens when CASO communication thresholds are reached”](#) on page 1536.

You can change the thresholds that trigger the communication statuses when managed media servers become unresponsive. You can also set how often the managed media server sends active job status updates to the central administration server. The frequency affects the network traffic.

To set communication thresholds and the frequency of active job status updates

- 1 On the central administration server’s navigation bar, click **Media Servers**.
- 2 Do one of the following:

To set defaults to apply to a managed media server when it is installed

Do the following in the order listed:

- On the task pane, under Media Server Installation Tasks, click **Configure managed media server defaults**.

To set defaults for an existing managed media server

Do the following in the order listed:

- Right-click the managed media server for which you want to configure settings, and then click **Properties**.
- Click the Settings tab.

To set defaults for a group of managed media servers in a pool

Do the following in the order listed:

- Select or create the media server pool that contains the managed media servers to which you want to apply new settings.
- Right-click the media server pool, and then click **Properties**.
- Click the Settings tab.
- Check **Apply these settings to all of the managed media servers in the pool**.

3 On the **Settings** tab, click **Custom settings**, and then click **Edit Custom Settings**.

4 On the **Configuration** tab, change the defaults as appropriate.

See “[Default configuration settings for managed media servers](#)” on page 1534.

5 Click **OK**.

Default configuration settings for managed media servers

You can change the thresholds that trigger the communication statuses when managed media servers become unresponsive. You can also set how often the managed media server sends active job status updates to the central administration server. The frequency affects the network traffic.

See “[Setting communication thresholds and active job status updates for CASO](#)” on page 1533.

Table 0-5 Default configuration settings for managed media servers

Item	Description
Communication Stalled (no more jobs are queued to the managed media server)	<p>Indicates the amount of time before the managed media server's status changes to Communication Stalled if the managed media server is unresponsive.</p> <p>The central administration server does not delegate jobs to the managed media server when it has a status of Communication Stalled. Job delegation resumes if the managed media server returns to an Enabled status before the threshold is exceeded.</p> <p>The default threshold is five minutes.</p>
No Communication (jobs are recovered from the managed media server)	<p>Indicates the amount of time before the managed media server's status changes from Communication Stalled to No Communication.</p> <p>When the status of the managed media server changes from Communication Stalled to No Communication, the central administration server marks the active jobs on the managed media server as Failed. The custom error-handling rule Recovered Jobs is applied to any job that is active when the No Communication status appears.</p> <p>The default threshold is 15 minutes.</p>
Send active job status updates in seconds to the central administration server	<p>Sends a job status update to the central administration server. You can adjust the number of seconds that a managed media server waits between sending job status updates to the central administration server. To preserve network bandwidth when many jobs are running, increase the amount of time between job update statuses. Decrease the amount of time if you want to send more updates.</p> <p>The default is 10 seconds, which provides near real-time monitoring. This setting is recommended only for fast network connections.</p> <p>For low-bandwidth network connections, consider a setting of 120 seconds. This frequency allows updates to be displayed for a medium-sized job while still significantly decreasing the network traffic caused by job status updates.</p> <p>If you uncheck the check box, job status updates are not sent. Job progress is not displayed on the central administration server. When the job is complete, the Job History view on the central administration server is updated.</p>

What happens when CASO communication thresholds are reached

In a CASO environment, communications that occur between managed media servers and the central administration server can sometimes be disrupted even if network communications are normal. If job-related communication disruptions occur between a managed media server and the central administration server, the managed media server's communication status changes from Enabled to Stalled. The jobs waiting to be processed by the managed media server are held in the managed media server's job queue until the communications are restored.

You can set the amount of time that Backup Exec waits before changing the managed media server's status if it becomes unresponsive. The configuration settings use time thresholds that when exceeded, change the managed media server statuses that are reported to the central administration server. These statuses include Stalled and No Communication. When a managed media server's status changes to Stalled or No Communication, the central administration server changes how it handles current and future jobs delegated to the stalled managed media server.

For example, if communications from a managed media server are not received at the central administration server after the set amount of time, the central administration server changes the media server's communication status to Stalled. Job delegation to the managed media server is suspended as it continues to wait for the managed media server to return to an Enabled status. Jobs are delegated to other managed media servers that are represented in the destination device or media server pool.

CASO continues to monitor the amount of time during which no communications are received from the managed media server. After a set amount of time passes after a Stalled status appears, CASO changes the status of the managed media server to No Communication. CASO marks the jobs as Failed, and then begins job recovery by invoking the custom error handling rule Recovered Jobs for any job that is active at the time the No Communication status appears.

See [“Setting communication thresholds and active job status updates for CASO”](#) on page 1533.

Copying logs and histories to the central administration server

During and after job processing, job log and job history information is generated for each job that is processed at each managed media server. By default, this information is stored locally at the managed media server where the jobs are processed. However, both job log and job history information can be copied and sent automatically to the central administration server, depending on the options you select in the Job Logs and Histories tab.

To copy logs and histories to the central administration server

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Do one of the following:

To set defaults to apply to a managed media server when it is installed On the task pane, under Media Server Installation Tasks, click **Configure managed media server defaults**.

To set defaults for an existing managed media server Right-click the managed media server for which you want to configure settings, and then click **Properties**.

To set defaults for a group of managed media servers in a pool Do the following in the order listed:

- Select or create the media server pool that contains the managed media servers to which you want to apply new settings.
- Right-click the media server pool, and then click **Properties**.
- On the **Settings** tab, check **Apply these settings to all of the managed media servers in the pool**.

- 3 On the **Settings** tab, click **Custom settings**, and then click **Edit Custom Settings**.
- 4 On the **Job Logs and Histories** tab, select the appropriate options.
See "[Job Logs and Histories for managed media servers](#)" on page 1537.
- 5 Click **OK**.

Job Logs and Histories for managed media servers

By default, job logs and job histories are stored locally at the managed media server where the jobs are processed. However, you can set up Backup Exec to send job logs and job histories to the central administration server.

See "[Copying logs and histories to the central administration server](#)" on page 1536.

Table O-6 Job Logs and Histories for managed media servers

Item	Description
Never	Keeps the job log on the managed media server where the job ran.
Every day at	Sends a copy of the job log to the central administration server every day at the time you specify.

Table O-6 **Job Logs and Histories** for managed media servers (*continued*)

Item	Description
On job completion	Sends a copy of the job log to the central administration server when the job completes.
Only if the job fails	Sends a copy of the job log to the central administration server only if the job fails.
Only when required by job type	Sends a copy of the job history to the central administration server whenever the managed media server processes a job type that requires the central administration server to store the job history. The job types Set Copy and Synthetic Backup require job histories to be stored on the central administration server.
Every day at	Sends a copy of the job history to the central administration server every day at a specified time.
On job completion	Sends a copy of the job history to the central administration server when the jobs complete.
Only if the job fails	Sends a copy of the job history to the central administration server only if the job fails.

How alerts work in CASO

In a Central Admin Server Option environment (CASO), alerts generated on a managed media server are automatically rolled up to the central administration server. To see those alerts on the central administration server, you must configure alert categories to enable or disable alerts on each managed media server in the CASO environment, and on the central administration server itself.

See [“Configuring alerts on the central administration server”](#) on page 1539.

After you respond to and clear the active alert on the central administration server, the alert is cleared on the managed media server as well.

If you enable Backup Exec alerts on a managed media server without enabling alerts on the central administration server, alerts appear only on the managed media server where they are generated; they will not appear on the central administration server.

Enable and configure alerts at the central administration server, and then copy the alert configurations to a managed media server. When the alert is generated on a managed media server, it appears on both the managed media server and the central administration server.

See [“Copying alerts to managed media servers”](#) on page 1539.

On the central administration server, you can view alerts for all managed media servers, or you can filter the alerts to view only those for a specific managed media server or media server pool.

Configuring alerts on the central administration server

In a Central Admin Server Option environment (CASO), alerts generated on a managed media server are automatically rolled up to the central administration server. To see those alerts on the central administration server, you must configure alert categories to enable or disable alerts on each managed media server in the CASO environment, and on the central administration server itself. If you enable Backup Exec alerts on a managed media server without enabling alerts on the central administration server, alerts appear only on the managed media server where they are generated; they will not appear on the central administration server.

See [“How alerts work in CASO”](#) on page 1538.

To configure specific alerts on the central administration server

- 1 Start Backup Exec on the central administration server.
- 2 On the **Tools** menu, click **Alert Categories**.
- 3 Scroll through the list of alert categories until you find the category you want to configure.
- 4 Under **Category properties**, check **Enable alerts for this category**.
- 5 Click **OK**.

Copying alerts to managed media servers

Enable and configure alerts at the central administration server, and then copy the alert configurations to a managed media server. When the alert is generated on a managed media server, it appears on both the managed media server and the central administration server.

To copy specific alerts to managed media servers.

- 1 From the **Tools** menu, select **Copy Settings to Media Servers**.
- 2 Under **Select Settings to Copy**, select the check box for **Alert Configuration**.
- 3 Click **Add**.
- 4 Enter the name of a managed media server to which the configuration will be copied.
- 5 Click **OK**.

- 6 On the **Copy Settings** dialog box, click **OK**.

An alert on the central administration server will confirm that the copy succeeded.

- 7 Click **OK** to clear the active alert.

About alerts and notification in CASO

In a Central Admin Server Option (CASO) environment, you can configure a notification on either the central administration server or the managed media server. Regardless of where you configure the notification, if it is for a delegated job, it is sent by the central administration server.

You can choose to notify the local administrator of the managed media server, or the administrator of the central administration server, or both.

See [“About alerts and notifications”](#) on page 644.

Enabling managed media servers to use any available network interface card

By default, jobs that are delegated or copied to a managed media server from the central administration server use the network and security settings that are set on the managed media server.

However, you can select an option on the central administration server to let a job use any network interface to access remote agents if the selected network interface is unavailable. Enabling this option for a backup job lets the managed media server use an alternative network interface to run important backup jobs that would otherwise fail to run.

To enable managed media servers to use any available network interface card

- 1 On the central administration server’s navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 In the task pane, under Settings, click **Network and Security**.
- 4 Check **Allow managed media server to use any network interface to access remote agents**.
- 5 Click **OK**.

About CASO catalog locations

In the CASO environment, you can choose the catalog location. Regardless of the catalog location, if a persistent network connection is available between the central administration server and the managed media server, then you can browse the backup sets in the catalog and perform restore operations from both servers.

The following catalog locations are available:

Table O-7

Item	Description
Distributed	<p>Image files, which are small files that contain information about the backup set, are distributed to the central administration server from every managed media server. History files, which contain detailed information about the backup set, remain on the managed media server.</p> <p>Note: It is important that you back up the catalog files on the managed media server since most catalog information is kept here when the distributed catalog location is used.</p> <p>When the catalog is distributed, the view of the restore selections on the central administration server displays only the backup set at the volume level. Backup set details are not displayed if the managed media server that created this backup set is not available, but the whole volume can be restored from the central administration server.</p> <p>A distributed catalog provides increased performance, default centralized restore capability, and decreased network traffic. If a managed media server does not have a persistent connection to the central administration server, then whenever the managed media server does connect, the image files in the catalog are automatically distributed to the central administration server. The temporary increase in network traffic caused by the catalog distribution is not significant.</p>

Table O-7 (continued)

Item	Description
Centralized	All catalog files and information for the managed media server are kept on the central administration server.
Replicated	<p>All catalog files are replicated from the managed media server to the central administration server. Both the managed media server and the central administration server store the catalog produced by the managed media server.</p> <p>Deletions of catalog files are replicated between the managed media server and the central administration server only when the catalog files are deleted by Backup Exec according to the catalog settings. If catalog files on the managed media server are deleted as a result of a backup job or a manual deletion, the deletions are replicated the next time that the catalogs are synchronized.</p>

When choosing the catalog location, consider the following:

- If there is enough available disk space on the managed media server to keep a distributed or replicated catalog.
- If there is enough network bandwidth to handle the traffic generated by a centralized or replicated catalog. Centralized and replicated catalogs require a high bandwidth network connection.
- If it is important for your data recovery needs to keep catalog information in one location. For example, when the catalog location is centralized or replicated, all catalog information is kept in one location, which makes it easier to back up. When the catalog location is distributed, most catalog information is kept on the managed media server.

See [“Changing the CASO catalog location”](#) on page 1542.

Changing the CASO catalog location

You can change the catalog location to distributed, centralized, or replicated in the CASO environment.

Changing the catalog location can cause catalog files to be copied or moved between the managed media server and the central administration server, which can increase network traffic. Additionally, you must restart the Backup Exec services on the managed media server on which the catalog location is changed.

See [“About CASO catalog locations”](#) on page 1541.

See [“Changing application settings for the catalogs”](#) on page 601.

See [“How centralized restore works in CASO”](#) on page 1552.

See [“How CASO restores data that resides on multiple devices”](#) on page 1553.

To change the catalog location

- 1 On the central administration server, click **Media Servers**.
- 2 Right-click the managed media server for which you want to change the catalogs, and then click **Properties**.
- 3 On the **media server properties** dialog box, in the **Advanced** tab, select the catalog location you want to use.

See [“Advanced properties for managed media servers”](#) on page 1543.

Advanced properties for managed media servers

You can change the catalog location to distributed, centralized, or replicated in the CASO environment.

See [“Changing the CASO catalog location”](#) on page 1542.

Table O-8 Advanced properties for managed media servers

Item	Description
Managed media server (distributed)	<p>Distributes the catalog files between the central administration server and the managed media server.</p> <p>From the central administration server, you cannot browse backup sets on a managed media server in a catalog that was created from Backup Exec version 9x or earlier.</p> <p>If device and media data are kept in a local database on the managed media server, then the distributed location is the only catalog location that is available.</p> <p>Select this option if you have a low-bandwidth network connection.</p>

Table 0-8 Advanced properties for managed media servers (*continued*)

Item	Description
Central administration server (centralized)	Keeps all catalog files on the central administration server. A high-bandwidth network connection should be available if this location is selected.
Both servers (replicated)	<p>Replicates all catalog files from the managed media server to the central administration server.</p> <p>If a managed media server is unavailable and the device is not shared (that is, if SAN Shared Storage Option is not installed on it), you can still browse the catalog from the central administration server. However, you cannot run a restore job since the managed media server is unavailable.</p> <p>A high-bandwidth network connection should be available if this location is selected.</p>

About job delegation in CASO

Job delegation is the automatic load balancing of jobs among the various storage devices that are attached to the Backup Exec managed media servers. The job is created on the central administration server, but can be run on any managed media server.

When the storage devices are logically grouped in device pools, and as the storage devices become available, they process jobs that are delegated from the central administration server. For example, if a device pool contains two storage devices and one is busy processing a job, the central administration server automatically delegates another job to the idle storage device.

Jobs are automatically created and submitted to the central administration server's job queue after a policy is applied to a selection list. Queued jobs are processed in priority order. Depending on job parameters and system configuration, the central administration server then delegates jobs to available storage devices in a selected device pool.

Templates within a policy can be sent to either a device pool, a specific device, or to devices in a media server pool. Device pools can consist of devices that are attached to a single managed media server, or they can consist of devices from different managed media servers. The advantage of job delegation is realized when a template is sent to a drive pool that spans managed media servers. When multiple managed media servers and their devices are available to the central administration server for job delegation, the efficiency of Backup Exec is greatly

improved because job processing does not have to wait for a specific device or managed media server to become available.

See [“About configuring CASO”](#) on page 1529.

See [“How to use media server pools in CASO ”](#) on page 1545.

How to use media server pools in CASO

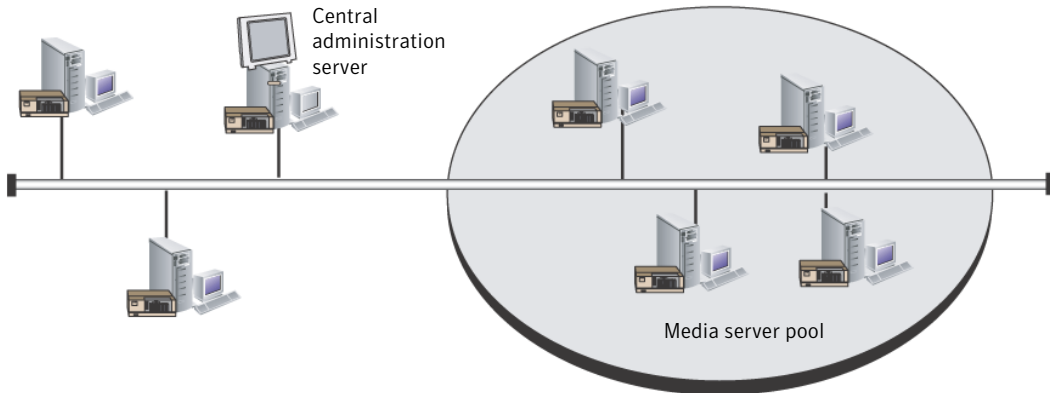
In a CASO environment, you can group multiple managed media servers together into media server pools. If you create a pool of managed media servers, all of the device pools on those managed media servers are available for job delegation. If there are multiple devices attached to each of the managed media servers in the media server pool, you can create multiple, smaller device pools that are made up of fewer storage devices. Use this method to send some jobs to a specific device pool in the media server pool, and send other jobs to a different device pool in the same media server pool.

Media server pools can contain multiple managed media servers or just one managed media server. A managed media server can belong to more than one media server pool. The central administration server can be used as a managed media server and can be included in the media server pool.

Any managed media server or media server in a pool must be able to access the destination device for the backup. If there is no intersection between the device and the managed media server or media server pools, the job does not run. The Job Monitor displays the following status: Ready; No media server available in media server pool.

This graphic shows a media server pool.

Figure O-4 An example of a CASO-configured media server pool inside a corporate network



Use a media server pool to do any of the following:

- Apply settings to all of the managed media servers in the pool.
See [“Applying settings to all managed media servers in a pool in CASO”](#) on page 1550.
- Restrict the backup jobs to a specific set of managed media servers and their attached storage devices.
See [“Restricting the backup of a selection list to specific devices in CASO”](#) on page 1546.
- View or filter data to be displayed for media servers in a pool, such as alerts, statistics, or active jobs.
See [“Viewing active job and alert statistics for a media server pool in CASO”](#) on page 1549.
- Copy configuration settings and logon information from a media server to all media servers in a pool.
See [“Copying configuration settings to another media server”](#) on page 199.

Restricting the backup of a selection list to specific devices in CASO

You can run a job on devices that are on a specific managed media server or on devices that are in a group of managed media servers. This filter lets you control where certain jobs are delegated. For example, to always run backups of Exchange databases only on the devices that are attached to managed media servers in a pool named Exchange Backups, you could select this option, and then select the Exchange Backups media server pool.

To restrict the backup of a selection list to specific devices

- 1 Do either of the following:
 - Create a backup job by setting job properties.
See [“Creating a backup job by setting job properties”](#) on page 332.
 - Create a backup selection list.
See [“Creating selection lists”](#) on page 296.
- 2 On the **Properties** pane, under **Destination**, click **Device and Media**.
- 3 Check **Restrict backup of the selection list to devices on the following media server or media servers in a pool**.
- 4 Select a media server or media server pool.
- 5 Continue setting job properties for the job.

Creating a media server pool in CASO

You can group, or pool, media servers.

See [“How to use media server pools in CASO”](#) on page 1545.

See [“Adding managed media servers to a media server pool in CASO”](#) on page 1547.

To create a media server pool

- 1 On the central administration server’s navigation bar, click **Media Servers**.
- 2 On the task pane, under **Media Server Pool Tasks**, click **New media server pool**.
- 3 Enter a media server pool name and description.
- 4 Select the media servers you want to add to the pool.
- 5 Click **OK**.

Adding managed media servers to a media server pool in CASO

You can add managed media servers to existing media server pools.

See [“Creating a media server pool in CASO”](#) on page 1547.

To add managed media servers to a media server pool

- 1 On the central administration server’s navigation bar, click **Media Servers**.
- 2 Select a media server pool to which you want to add managed media servers.
- 3 On the task pane, under **Media Server Pool Tasks**, click **Add media server**.

- 4 Select the media servers you want to add to the pool.
- 5 Click **OK**.

Renaming a media server pool in CASO

You can rename a media server pool at any time.

To rename a media server pool

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Right-click the media server pool that you want to rename, and then click **Rename**.
- 3 When prompted, enter a new name for the media server pool.
- 4 Click **OK**.

Deleting a media server pool in CASO

You can delete a media server pool at any time.

To delete a media server pool

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Right-click the media server pool that you want to delete, and then click **Delete**.
- 3 When prompted to confirm the delete operation, click **Yes**.

Removing a managed media server from a media server pool in CASO

Removing a managed media server deletes it from a media server pool, but does not remove it from the All Managed Media Servers node.

To remove a managed media server from a media server pool

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Select a media server pool from which you want to remove a managed media server.
- 3 On the right pane, select the managed media server that you want to remove.
- 4 On the task pane, under **Media Server Pool Tasks**, click **Remove media server**.
- 5 Click **OK**.

Viewing general properties for a media server pool in CASO

On the General tab of the media server pool properties, you can view the name, description, and creation date of this media server pool.

To view general properties for a media server pool

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Right-click the media server pool, and then click **Properties**.
- 3 On the **General** tab, view the properties.
- 4 Click **OK**.

Viewing active job and alert statistics for a media server pool in CASO

On the Statistics tab of the media server pool properties, you can view the number of devices, jobs, and alerts for this media server pool. You can also click links to immediately view active jobs and alerts.

To view active job and alert statistics for a media server pool

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Right-click the media server pool, and then click **Properties**.
- 3 On the **Statistics** tab, view the properties.
See "[Statistics properties for a media server pool in CASO](#)" on page 1549.
- 4 Click **OK**.

Statistics properties for a media server pool in CASO

You can view the number of devices, jobs, and alerts for a media server pool.

See "[Viewing active job and alert statistics for a media server pool in CASO](#)" on page 1549.

Table O-9 **Statistic** properties for a media server pool in CASO

Item	Description
Number of devices	Indicates the number of storage devices in this media server pool. When multi-drive robotic libraries are attached, each drive in the robotic library represents a separate device.
Number of active jobs	Indicates the number of jobs currently in progress in this media server pool.

Table O-9 **Statistic** properties for a media server pool in CASO (*continued*)

Item	Description
View jobs...	Shows all current jobs and job histories for this media server pool. The Job Monitor view on the navigation bar is displayed, and the filters for current jobs and job histories are reset to display jobs for this media server.
Number of error alerts	Indicates the number of active error alerts for this media server.
Number of warning alerts	Indicates the number of active warning alerts for this media server.
Number of informational alerts	Indicates the number of active information alerts for this media server.
View alerts...	Shows all active alerts for this media server. The Alerts view on the navigation bar is displayed, and the filter is reset to display active alerts for this media server.

Applying settings to all managed media servers in a pool in CASO

You can apply the same settings to all of the managed media servers that are in a media server pool.

To apply settings to all managed media servers in a pool

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Select or create a media server pool that contains the managed media servers to which you want to apply the settings.
- 3 Right-click the media server pool, and then click **Properties**.
- 4 On the **Settings** tab, check **Apply these settings to all of the managed media servers in the pool**.

- 5 Make any changes to the defaults.
See [“How CASO works”](#) on page 1504.
See [“Setting communication thresholds and active job status updates for CASO ”](#) on page 1533.
See [“Copying logs and histories to the central administration server”](#) on page 1536.
- 6 Choose any setting that you want to apply to all of the managed media servers in the selected media server pool, and then click **OK**.

About copying jobs instead of delegating jobs in CASO

If the managed media server’s device and media data are kept on a local database on the managed media server, the central administration server cannot delegate jobs to it. Instead, you can copy policies, selection lists, and configuration settings from the central administration server to the managed media server. A persistent network connection to the central administration server is not needed when the jobs are run locally on the managed media server.

If you associate the policies and selection lists on the central administration server, jobs are created there. You can then copy the jobs to the managed media server. You can also copy the policies and selections lists to the managed media server, and then associate them. The jobs are created on the managed media server.

Use the same names for objects on the central administration server and all of the managed media servers that you want to copy jobs to. For example, use the same name for a device pool on the central administration server and on the managed media server. Then, it is not necessary to customize settings or names for each managed media server that you copy jobs to.

See [“Copying configuration settings to another media server”](#) on page 199.

See [“Copying jobs, selection lists, or policies”](#) on page 554.

See [“Setting defaults for managed media servers”](#) on page 1531.

Requirements for duplicate backup data and synthetic backup jobs in CASO

A recurring job from a policy that contains a Duplicate Backup Sets template or a Synthetic Backup template must be run on the same managed media server where the job was initially run. The jobs produced from these templates require

access to the media that contains the backup sets that were produced from the preceding jobs.

If you change the target device of the templates within a policy, the jobs resulting from that template can then be delegated to a different managed media server. However, if the destination managed media server is still a valid candidate for delegation, it will be used.

If you do not change the destination devices in the templates, and if the targeted managed media server is unavailable, the jobs remain queued, waiting for the targeted managed media server to become available. If the targeted managed media server is no longer configured as a managed media server, then you can re-delegate the jobs.

See [“About creating jobs using policies and selection lists”](#) on page 544.

See [“Adding a duplicate backup template to a policy”](#) on page 550.

See [“About creating a synthetic backup by copying the example policy”](#) on page 906.

How centralized restore works in CASO

Depending on whether the required storage media resides in storage devices or is stored offsite, initiating restore operations from the central administration server can be an automated process with little user intervention necessary.

When you use centralized restore with online media, you make restore selections and set job properties at the central administration server. During the data selection process, CASO determines the media required to complete the restore operation, and then queries the Backup Exec device and media database to determine the identity of the storage device where the primary media required for the job currently resides. After you make your selections, set restore job properties, and launch the restore job, CASO begins the restore operation by delegating the jobs to the central administration server or managed media servers that control the selected storage devices. If the data being restored spans multiple storage media, you are prompted to load additional media as needed to successfully complete the restore operation.

When you use centralized restore with offline media, you make restore selections and set job properties at the central administration server. During the data selection process, CASO determines the media required to complete the restore operation, and then queries the Backup Exec device and media database to determine the identity of the storage device where the primary media required for the job currently resides. If the media is not found in a storage device, the media is considered offline. CASO then presents you with a selection of drive pools and storage devices that are compatible with the type of media being used during

the restore operation, thus giving you the flexibility of choosing a storage device in which to load your media.

After noting the identity and location of the storage device you have selected to run the job, you do the following:

- Submit the restore job on hold as a scheduled job
- Retrieve the media, place it in the storage device
- Remove the job from hold at the central administration server, at which time the restore job begins.

CASO then delegates the job to the managed media server that controls the selected storage device. If the data being restored spans multiple storage media, you are prompted to load additional media as needed to successfully complete the restore operation.

Before restore operations from the central administration server can be initiated, the following requirements must be met:

- Managed media server communication status must be Enabled.
- Managed media servers must be online with all media server statuses showing Online.

See [“About media in Backup Exec”](#) on page 219.

How CASO restores data that resides on multiple devices

If the data selected for restore is located on a single device attached to a managed media server, then a single restore job is created at, and then delegated from, the central administration server. However, if the data being selected for restore is located on multiple devices in the CASO environment, then the single restore job is split into separate restore jobs, depending on the number of devices involved.

All split restore jobs have the same name as the original job, but are differentiated and linked with a subscript numeral that is appended to the job name.

For example, if you create a restore job and the data you select for restore resides in one device on a managed media server, CASO creates one restore job. However, if you create one restore job and the data you select resides on two or more devices that are attached to a managed media server, CASO creates two or more restore jobs.

The following graphic shows a single restore job that is split into multiple jobs.

Figure 0-5 Example of subscript numerals linking split restore jobs

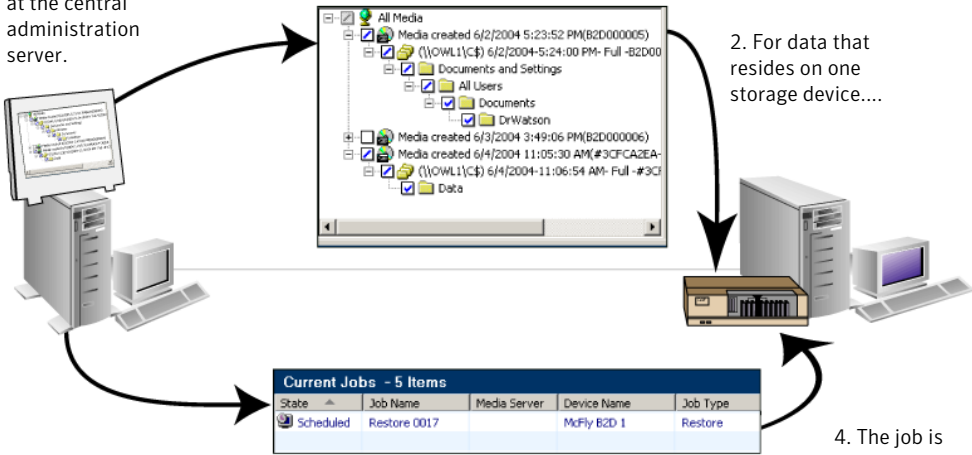
Single restore job split into multiple jobs, are visually linked using a subscript numeral appended to the original job name.

Current Jobs - 5 Items							
State	Job Name	Media Server	Device Name	Job Type	Job Status	Priority	Percent ...
Scheduled	Restore 0015		McFly B2D 1	Restore	Ready; N...	Medium	None
Scheduled	Restore 0016		McFly B2D 1	Restore	Ready; N...	Medium	None
Scheduled	Restore 0016(2)		Seagate 1	Restore	Ready; N...	Medium	None
Scheduled	Restore 0017		McFly B2D 1	Restore	Ready; N...	Medium	None
Scheduled	Restore 0017(2)		Seagate 1	Restore	Ready; N...	Medium	None

The following graphic shows how CASO restores data that is stored on a single device.

Figure 0-6 For data stored on a single storage device

1. Data selections for restore are made at the central administration server.

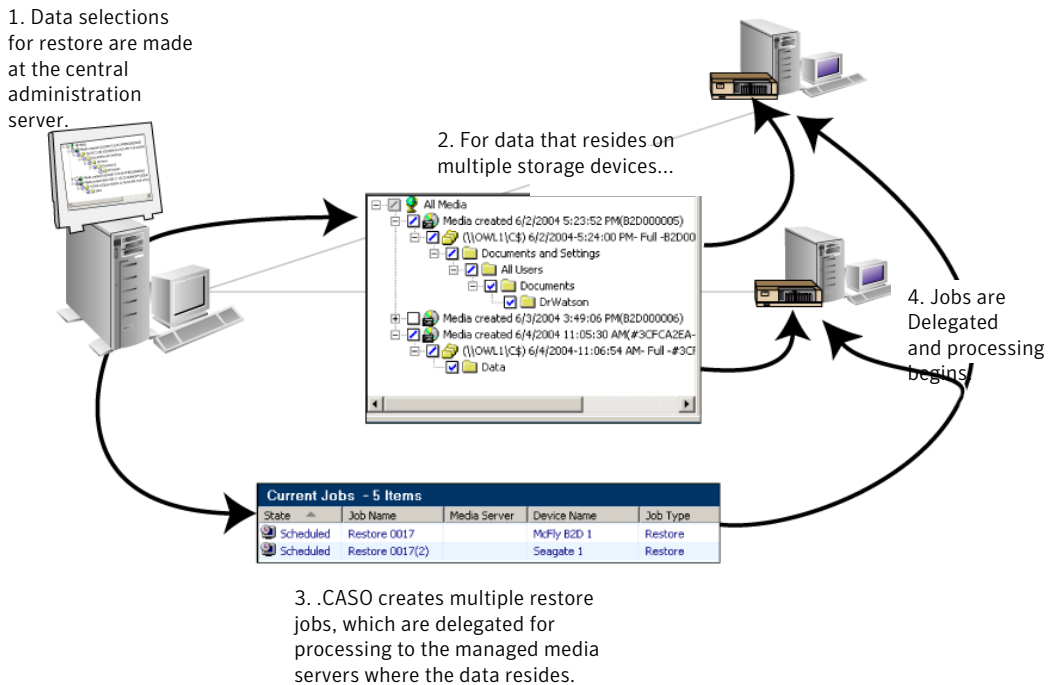


2. For data that resides on one storage device...

3. CASO creates one restore job, which is delegated for processing at the managed media server where the data resides.

4. The job is delegated and processing begins.

The following graphic shows how CASO restores data that is stored on multiple devices.

Figure 0-7 For data stored on multiple storage devices

Best practices for centralized restore in CASO

Symantec recommends the following for best practices when using centralized restore:

- Select only one resource to restore for each job.
- Select the same restore device or media server for all of the selections that are in the same restore job.
- Select a media server that has compatible devices for all media that is required for the restore job.

See [“Restoring from the CASO central administration server”](#) on page 1555.

Restoring from the CASO central administration server

Before you create a restore job, review the best practices for centralized restore.

See [“Best practices for centralized restore in CASO”](#) on page 1555.

See [“How centralized restore works in CASO”](#) on page 1552.

To restore from the central administration server

- 1 Create a restore selection list.
 See [“Creating a restore selection list”](#) on page 628.
- 2 On the **Restore Job Properties** dialog box, in the task pane, under **Source**, click **Device and Media**.
- 3 Enter or change the information as appropriate:
 See [“Device and Media properties for a CASO restore job”](#) on page 1556.
- 4 Set other restore job properties from the **Properties** pane.
 See [“Restoring data by setting job properties”](#) on page 605.
- 5 Click **Run Now** to start the restore operation.

Device and Media properties for a CASO restore job

You can view the following information about a CASO restore job:

- A list of media that is required.
- The location of the media.
- The names of the devices or the media servers that are possible candidates to process the restore job.

See [“Restoring from the CASO central administration server”](#) on page 1555.

Table O-10 Device and Media properties for a CASO restore job

Item	Description
Media or Resource	Displays a list of media that is required for the restore, or the name of the resource that you selected for restore.
Media Location	<p>Displays the location of the media. If the media is listed as offline or unknown, you must retrieve the media, select a device in the Restore Device or Media Server column, and then place the media in a device that the managed media server can access.</p> <p>If the data that is selected for restore resides in a media vault, then Offline is displayed.</p> <p>If the data that is selected for restore resides in an unknown media location, then Unknown is displayed because the media cannot be found in any of the compatible storage devices that are candidates to run the job.</p>

Table O-10 Device and Media properties for a CASO restore job (*continued*)

Item	Description
Restore Device or Media Server	<p>Displays the names of the devices or the media servers that are compatible with the media being restored, and that are possible candidates to process the restore job.</p> <p>To support restoring from the central administration server when the device and media database is on the managed media server, all media servers, including the central administration server, are listed in this column.</p> <p>Note: When you restore an Oracle database in a CASO environment, you must use the managed media server that was used for the original database backup.</p>

Media Servers view in CASO

After CASO is installed, you can perform tasks on the managed media servers from the **Media Servers** view on the central administration server. Managed media servers are also displayed in the **Job Monitor** view, in **Current Jobs** and in **Job History**.

The following information is displayed in the **Media Servers** view.

Table O-11 Media Servers view

Item	Description
Name	<p>Displays the name of the managed media server or central administration server. Along with the name, an icon representing a managed media server or central administration server is used to help you quickly differentiate between the two.</p> <p>See “Media Servers view in CASO” on page 1557.</p>
Description	<p>Displays a user-defined description of the managed media server or the central administration server.</p>

Table O-11 **Media Servers view (continued)**

Item	Description
Communication Status	<p>Displays the status of communications between the managed media server and the central administration server.</p> <p>Statuses include the following:</p> <ul style="list-style-type: none"> ■ Enabled - Communications about jobs between the managed media server and the central administration server are working properly. ■ Disabled - Communications between the managed media server and the central administration server have been disabled by the user. ■ Stalled - Communications between the managed media server and the central administration server have not occurred within the configured time threshold. See “What happens when CASO communication thresholds are reached” on page 1536. ■ No communication - No communication about jobs is being received at the central administration server from the managed media server. The configured time threshold has been reached. Jobs that are sent to the managed media server are recovered. Possible causes can range from a network failure to hardware failure in either the managed media server or the central administration server. See “About recovering failed jobs in CASO” on page 1560. ■ N/A - This status appears when the computer that is displayed in the Name column is a central administration server.
Media Server Status	<p>Displays the present status of a managed media server.</p> <p>Valid statuses include the following:</p> <ul style="list-style-type: none"> ■ Online - All managed media server services are running and communication with the central administration server is functioning as expected. ■ Paused - The managed media server has been placed in a Paused state by the user. Jobs are not delegated to the managed media server when it is in a Paused state. ■ Unavailable - A state where no communications have been received by the central administration server from the managed media server. ■ Offline - All managed media server services have been stopped; the media server cannot run jobs in an offline state.
Media Server Type	<p>Displays Backup Exec’s description of the media server. It can be either a managed media server or a central administration server.</p>

Table O-11 Media Servers view (continued)

Item	Description
Monitor Jobs Created Locally	Displays Yes if the option Monitor jobs that are created on the local managed media server in addition to the jobs that are delegated from the central administration server is enabled; otherwise, No is displayed. If this option is enabled, you can hold, delete, run, cancel, and change the priority of jobs that were copied from the central administration server. See “ How CASO works ” on page 1504.
Catalog Location	Displays the location of the catalog. See “ Changing the CASO catalog location ” on page 1542.
Version	Displays the version of Backup Exec that is installed and running on the media server.
Operating System	Displays the type of operating system that is installed and running on the media server.
Operating System Build	Displays the build number of the operating system that is installed on the media server.

Icons are used in the Media Server view to help you quickly identify the operational status of Backup Exec managed media servers that appear in the Results pane.

You can find a list of icons that appear in the CASO Media Server view at the following URL:

<http://entsupport.symantec.com/umi/V-269-12>

About managing jobs in CASO

You can locate the server on which jobs were created on the **Job Setup** view or the **Job Monitor** view. The column labeled Created On indicates whether the job was created on the central administration server or on the managed media server.

If the option to monitor jobs that are created on the local managed media server is enabled, you can hold, delete, run, cancel, and change the priority of jobs that were copied from the central administration server. However, if you create jobs on the managed media server by associating selection lists and policies, then you can only delete those jobs from the managed media server on which the policy originated.

See “[Deleting a job created from a policy](#)” on page 546.

See “[About copying jobs instead of delegating jobs in CASO](#)” on page 1551.

See [“Setting defaults for managed media servers”](#) on page 1531.

About recovering failed jobs in CASO

The Backup Exec error-handling rule named Recovered Jobs is a custom error-handling rule that is used by CASO to recover jobs that failed because of issues with internal job communications. This rule is created when Backup Exec is installed and is enabled by default.

The retry options for this rule are to retry the job twice, with an interval of five minutes between the retry attempts. During the first retry attempt, CASO attempts to re-delegate the jobs to another available managed media server.

If this attempt fails, CASO makes a second attempt at finding another available managed media server to process the jobs. If no managed media servers are available, the final job disposition is to place the job on hold until you have fixed the error condition.

Note: If you target a job to a media server pool that contains multiple managed media servers and a job failure occurs, the recovery process uses only the managed media servers in the media server pool. Managed media servers that are not in the media server pool are not used for job recovery.

A CASO job with a status of No Communication that is failed and then recovered by Backup Exec is displayed in the Backup Exec job history view in gray, with a job status of Recovered. A CASO job that failed due to errors in normal daily activities is displayed in red text, the same as other failed jobs.

When you open the job history entry for a Recovered job, the reason for the failure is listed as Job Errors, with an explanation of the type of internal communication error that occurred. The job history entry also indicates that the job was recovered.

Note: Job logs are not created for jobs that are recovered.

The following table describes the CASO error codes that are selected by default for the Recovered Jobs custom error-handling rule:

Table 0-12 Error codes for Recovered Jobs custom error-handling rule

Error code	Description
0xE000881B JOBDISPATCH	The displayed message is: Job failed while being dispatched. The job will be recovered.
0xE000881D JOB_CASO_QUEUE FAILURE	The displayed message is: The job could not be delegated to the destination managed media server. The managed media server may not be online, or there may be a communications failure. The job will be recovered.
0xE000881E JOB_CASO_REMOTEMMS_STARTFAILURE	The displayed message is: The job failed to start on the destination managed media server, possibly because a database error occurred. The job will be recovered.

See [“About error-handling rules”](#) on page 590.

See [“Custom error-handling rule for recovered jobs”](#) on page 594.

See [“Setting communication thresholds and active job status updates for CASO ”](#) on page 1533.

Pausing a managed media server in CASO

You can pause and resume a managed media server from the central administration server.

Pausing a managed media server prevents the central administration server from delegating jobs to it. When paused, the managed media server’s status changes from Online to Paused, and is reflected as such in the Media Server Status column of the Results pane.

Caution: When installing Backup Exec options at a managed media server, the managed media server must be paused, so that no further jobs are delegated to it from the central administration server while the installation process occurs. If jobs are running, let them finish or cancel them before beginning the installation.

To pause a managed media server

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Select the managed media server you want to pause or resume.
- 3 In the task pane, under **Media Server Tasks**, click **Paused**.

Resuming a paused managed media server in CASO

When you resume a paused managed media server, the following changes occur:

- The managed media server's status changes from Paused to Online in the **Media Servers Status** column.
- An icon representing a managed media server in online state appears.

To resume a paused managed media server

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Select the managed media server you want to resume.
- 3 In the task pane, under **Media Server Tasks**, click **Paused**.

How paused storage devices appear on the Devices view in CASO

After pausing managed media server storage devices from the central administration server, the storage devices appear in the **Devices** view with both a Paused Managed Media Server status icon, and the word Paused.

However, when viewing the list of storage devices under the **Devices** view at the managed media server, the storage devices that were paused at the central administration server do not appear as Paused.

You must use F5 to refresh the **Devices** view on the managed media server to view the actual state.

See [“Media Servers view in CASO”](#) on page 1557.

Disabling communications in CASO

You can disable and enable managed media server communications from the central administration server. When you disable communications, the managed media server's communications status changes from Enabled to Disabled in the

Communications Status column. Its status change is also reflected in the Media Server Status column, which changes from Online to Unavailable. An icon representing a disabled managed media server appears.

To disable and enable communications

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Select the managed media server that you want to disable or enable.
- 3 In the task pane, under **Media Server Tasks**, click **Communication Enabled**.

Enabling communications in CASO

You can disable and enable managed media server communications from the central administration server. When you enable communications, the managed media server's communications status changes from Disabled to Enabled in the Communications Status column. Its status change is also reflected in the Media Server Status column, which changes from Unavailable to Online. Finally, an icon representing a fully functioning managed media server appears.

To enable communications

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Select the managed media server that you want to disable or enable.
- 3 In the task pane, under **Media Server Tasks**, click **Communication Enabled**.

Stopping Backup Exec services for CASO

You can stop and start the Backup Exec services on a managed media server from the central administration server.

To stop Backup Exec services from the central administration server

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Select the managed media server on which you want to start or stop services.
- 3 In the task pane, under **Media Server Tasks**, click **Backup Exec services**.
- 4 Click **Stop all services**.

Starting Backup Exec services for CASO

You can stop and start the Backup Exec services on a managed media server from the central administration server.

To start Backup Exec services from the central administration server

- 1 On the central administration server's navigation bar, click **Media Servers**.
- 2 Select the managed media server on which you want to start or stop services.
- 3 In the task pane, under **Media Server Tasks**, click **Backup Exec services**.
- 4 Click **Start all services**.

Connecting to a remote managed media server

To run administrative tasks on a managed media server from the central administration server, connect to the managed media server using the Backup Exec feature, Connect to Media Server.

To connect to a remote managed media server

- 1 Start Backup Exec on the central administration server.
- 2 From the navigation bar, click **Media Servers**.
- 3 Select the managed media server to which you want to connect.
- 4 In the task pane, under **Media Server Tasks**, click **Connect to media server**.
- 5 If the managed media server name does not appear, enter the computer name in the **Server** field.
- 6 Enter administrator or administrator equivalent logon information for the managed media server.
- 7 In the **Domain** field, enter the computer name of the managed media server.
- 8 Click **OK**.

Connection to the managed media server occurs.

After making a remote connection to a managed media server, the central administration server console is closed and the managed media server console becomes the active interface.

- 9 To return to the central administration server console, close the managed media server console and restart Backup Exec on the central administration server.

Viewing managed media server properties

You can view properties for managed media servers in the Media Servers view on the central administration server.

To view managed media server properties

- 1 Start Backup Exec on the central administration server.
- 2 From the navigation bar, click **Media Servers**.
- 3 Select the managed media server for which you want to view properties.
- 4 In the task pane, under **General Tasks**, click **Properties**.

Disaster Recovery in CASO

Use the Symantec Backup Exec Intelligent Disaster Recovery (IDR) option to protect both managed media servers and the central administration server in a CASO environment.

See [“About using IDR with the Central Admin Server Option”](#) on page 1845.

Before implementing the IDR option in a CASO environment, review the following:

- In a CASO environment, all disaster preparation files (*.dr files) created for each managed media server are centrally located on the central administration server.
- To create recovery media for any managed media server or central administration server, the IDR Preparation Wizard must be run at the central administration server. Or if running on remote administration environment, connect to the central administration server.
- If you want managed media servers to be protected using a bootable tape image, you must run the IDR Preparation Wizard at each of the managed media servers where a bootable tape device is installed.
- For CASO, two options appear on the Welcome screen when the IDR Preparation Wizard is run on a managed media server.

These options include the following:

- Yes, create the bootable tape image now.
See [“Creating a bootable tape image”](#) on page 1819.
- No, connect to a central administration server - If a bootable tape drive is not detected on a managed media server, only this option appears.
- You must locally back up and restore a central administration server.

Symantec Backup Exec Deduplication Option

This appendix includes the following topics:

- [About the Deduplication Option](#)
- [Requirements for the Deduplication Option](#)
- [About installing the Deduplication Option](#)
- [About OpenStorage devices](#)
- [About deduplication storage folders](#)
- [Sharing a deduplication device between multiple media servers](#)
- [About client-side deduplication](#)
- [About backup jobs for deduplication](#)
- [About optimized duplication](#)
- [About copying deduplicated data to tapes](#)
- [About using deduplication with encryption](#)
- [About restoring deduplicated data](#)
- [About disaster recovery of a deduplication storage folder](#)
- [About disaster recovery of OpenStorage devices](#)

About the Deduplication Option

The Backup Exec Deduplication Option supports a data-reduction strategy by optimizing storage and network bandwidth. The Deduplication Option supports integrated deduplication at the Backup Exec media server and on remote computers that have the Remote Agent for Windows Systems installed. It also allows data to be deduplicated and stored on intelligent disk devices from Symantec and other vendors.

Table P-1 Types of deduplication

Type of deduplication	Where deduplication occurs	Benefits
Media server deduplication	On the Backup Exec media server.	Reduces the size of backups, which reduces storage requirements.
Client-side deduplication	On the remote computer where the data is located. Note: The Remote Agent for Windows Systems is required on the remote computer to perform client-side deduplication.	Reduces network traffic because only unique data is sent across the network. It also reduces the backup window.
Appliance deduplication	On an intelligent disk device, such as Symantec PureDisk or a device from a third-party vendor.	Reduces the size of backups, which reduces storage requirements. It also reduces the backup window.

With a single Deduplication Option license key, you can use two types of deduplication devices.

Table P-2 Types of deduplication devices that work with the Deduplication Option

Type of device	Description
OpenStorage device	<p>Backup Exec uses Symantec's OpenStorage technology, which allows intelligent disk devices to integrate with Backup Exec. You can back up data to the Symantec PureDisk device and to storage devices from other vendors.</p> <p>You can find a list of compatible devices at the following URL:</p> <p>http://entsupport.symantec.com/umi/V-269-2</p> <p>See “About OpenStorage devices” on page 1574.</p>
Deduplication storage folder	<p>A deduplication storage folder provides integrated deduplication on the Backup Exec media server. A deduplication storage folder is a disk-based backup folder that is located on the Backup Exec media server. It is similar to a backup-to-disk folder.</p> <p>See “About deduplication storage folders” on page 1579.</p>

In addition to reducing storage requirements and network traffic, the Deduplication Option lets you do the following:

- Copy deduplicated data from an OpenStorage device or a deduplication storage folder to tape for long-term or off-site storage.
- Use optimized duplication, which lets you copy deduplicated data between OpenStorage devices from the same vendor and between deduplication storage folders.
- Use Symantec's Granular Recovery Technology (GRT) with jobs that use deduplication devices.
- Share OpenStorage devices and deduplication storage devices among multiple media servers when you use the Central Admin Server Option or the SAN Shared Storage Option.

See “[About installing the Deduplication Option](#)” on page 1573.

See “[Requirements for the Deduplication Option](#)” on page 1572.

See “[Sharing a deduplication device between multiple media servers](#)” on page 1589.

See [“About optimized duplication”](#) on page 1595.

See [“About copying deduplicated data to tapes”](#) on page 1597.

Deduplication methods for Backup Exec agents

Backup Exec supports the following deduplication methods:

- Client-side deduplication, either on an intelligent disk device or to a deduplication storage folder.
- Media server deduplication with a deduplication storage folder.
- Appliance deduplication on an OpenStorage device.

The following table lists the deduplication methods that are available for the Backup Exec agents.

Table P-3 Deduplication methods for Backup Exec agents

Agent	Client-side deduplication (file system/VSS)	Client-side deduplication (with Granular Recovery Technology enabled)	Media server deduplication (file system/VSS)	Media server deduplication (with Granular Recovery Technology enabled)	Appliance deduplication on an OpenStorage device
Remote Agent for Windows Systems	Yes	No	Yes	No	Yes
Agent for VMware	Yes Note: The Remote Agent for Windows Systems must be installed on the Guest virtual machine.	Yes Note: The Remote Agent for Windows Systems must be installed on the Guest virtual machine.	Yes	Yes	Yes

Table P-3 Deduplication methods for Backup Exec agents (*continued*)

Agent	Client-side deduplication (file system/VSS)	Client-side deduplication (with Granular Recovery Technology enabled)	Media server deduplication (file system/VSS)	Media server deduplication (with Granular Recovery Technology enabled)	Appliance deduplication on an OpenStorage device
Agent for Microsoft Hyper-V	Yes Note: The Remote Agent for Windows Systems must be installed on the Guest virtual machine.	Yes Note: The Remote Agent for Windows Systems must be installed on the Guest virtual machine.	Yes	Yes	Yes
Remote Agent for Linux or UNIX Servers	No	No	Yes	No	Yes
Agent for Enterprise Vault	Yes	No	Yes	No	No
Exchange Agent	Yes	Yes	Yes	Yes	Yes
SQL Agent	Yes	No	Yes	No	Yes
SharePoint Agent	Yes	Yes	Yes	Yes	Yes
Active Directory Agent	Yes	Yes	Yes	Yes	Yes
Oracle Agent for Linux Servers	No	No	Yes	No	Yes

Table P-3 Deduplication methods for Backup Exec agents (*continued*)

Agent	Client-side deduplication (file system/VSS)	Client-side deduplication (with Granular Recovery Technology enabled)	Media server deduplication (file system/VSS)	Media server deduplication (with Granular Recovery Technology enabled)	Appliance deduplication on an OpenStorage device
Oracle Agent	Yes	No	Yes	No	Yes
SAP Agent	Yes	No	Yes	No	Yes
Lotus Domino Agent	Yes	No	Yes	No	Yes
DB2 Agent	Yes	No	Yes	No	Yes
NetWare Agent	No	No	Yes	No	Yes
Remote Agent for Macintosh System	No	No	Yes	No	Yes

See [“How to restore individual items by using Granular Recovery Technology”](#) on page 321.

Requirements for the Deduplication Option

The requirements for the Deduplication Option vary depending on whether you use deduplication storage folders or OpenStorage devices. Before you install the Deduplication Option, you should determine what type of storage devices you want to use with it. Then, verify that your system meets the requirements for the storage devices you want to use.

Table P-4 Requirements for the Deduplication Option

Type of storage device	Requirements
Deduplication storage folders	<p>The following items are required:</p> <ul style="list-style-type: none"> ■ A 64-bit media server. ■ A media server with either one quad-core processor or two dual-core processors. ■ A dedicated volume to use as the location to store the deduplication storage folder. ■ 8 GB RAM, which supports up to 5 TB of deduplicated data. For more than 5 TB of data, use the following calculation to determine the required amount of RAM: $1.5 \text{ GB} \times N$, where N = the number of TBs of deduplicated data to be stored. Up to 16 TB of data is supported. <p>Example: Use the following equation to calculate the required amount of RAM for 10 TB of data:</p> $1.5 \times 10 = 15 \text{ GB RAM}$
OpenStorage devices	<p>To use a Symantec PureDisk device or a storage device from another vendor as an OpenStorage device, you must purchase the device and the appropriate OpenStorage connector from the device's vendor.</p> <p>You can use the Deduplication Option with OpenStorage devices on either a 32-bit media server or a 64-bit media server. The standard system requirements for Backup Exec apply to the Deduplication Option when you use OpenStorage devices.</p>

See [“About installing the Deduplication Option”](#) on page 1573.

About installing the Deduplication Option

The Deduplication Option is installed using the Backup Exec installation media. You install it locally as a separate, add-on component of Backup Exec. Before you attempt to install the Deduplication Option, verify that your system meets the requirements.

See [“Requirements for the Deduplication Option”](#) on page 1572.

See “[Installing additional Backup Exec options to the local media server](#)” on page 123.

About OpenStorage devices

OpenStorage is a Symantec technology that allows intelligent disk devices to integrate with Backup Exec.

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

Some intelligent disk devices can include multiple logical storage units. However, each logical storage unit is added as a single OpenStorage device. When you add an OpenStorage device, Backup Exec can automatically locate the logical storage units on that device.

Note: When you delete or erase the media from an OpenStorage device, it may take up to 48 hours for more space to become available. Backup Exec cannot always calculate the amount of space that will be made available.

After you add a deduplication storage device, it appears in the **Devices** view. However, an OpenStorage device does not appear in the **All Devices** device pool. An OpenStorage device cannot belong to any device pools to prevent a deduplication job from being sent to a non-deduplication device in a device pool if the OpenStorage device is busy.

If you use Backup Exec Central Admin Server Option or SAN Shared Storage Option, an OpenStorage device can be shared between multiple media servers. Sharing can be enabled when you add an OpenStorage device. You can select new media servers to share an OpenStorage device or remove the sharing ability for media servers at any time.

See “[Adding an OpenStorage device](#)” on page 1574.

See “[About sharing storage](#)” on page 444.

Adding an OpenStorage device

Follow these steps to add an intelligent disk device as an OpenStorage device.

See “[About OpenStorage devices](#)” on page 1574.

To add an OpenStorage device

- 1 On the navigation bar, click **Devices**.
- 2 In the task pane, under **Device Tasks**, click **Configure devices assistant**.

- 3 On the **Configure Devices Assistant** dialog box, under **Deduplication Option**, click **OpenStorage**.
- 4 If the **OpenStorage Configuration** dialog box appears, click **Add OpenStorage**.
This step does not apply if this is the first OpenStorage device. The **OpenStorage Configuration** dialog box appears only if an OpenStorage device already exists.
- 5 Complete the options on the **General** tab.
See "[General OpenStorage device options](#)" on page 1575.
- 6 Complete the options on the **Advanced** tab.
See "[Advanced OpenStorage device options](#)" on page 1576.
- 7 On the **Sharing** tab, select each media server that you want to use with this OpenStorage device.

The **Sharing** tab appears only if you have the Backup Exec Central Admin Server Option or the SAN Shared Storage Option installed. You must also have more than one media server.
- 8 Click **OK**.
- 9 Restart the Backup Exec services on the media servers that you selected in step 7.

See "[Starting and stopping Backup Exec services](#)" on page 168.

General OpenStorage device options

You can set the following options for an OpenStorage device.

See "[Adding an OpenStorage device](#)" on page 1574.

Table P-5 General OpenStorage device options

Item	Description
Name	Indicates the name of the device. If you do not provide a name, Backup Exec automatically creates a name after you select the server type.
Server	Indicates the fully-qualified name of the server on which the device exists.
Logon account	Indicates the name of the logon account that is required to access the device.
Server type	Indicates the type of OpenStorage device.

Table P-5 General OpenStorage device options (*continued*)

Item	Description
Logical storage unit	<p>Indicates the logical storage unit that you want to use. Backup Exec locates all of the logical storage units that are on the device and displays them in the list. You can select a logical storage unit from the list or type the name of one.</p> <p>Note: This option does not appear for Symantec PureDisk devices.</p>
Allow x concurrent operations for this device	<p>Indicates the number of jobs that you want to run at the same time on this device.</p>

Advanced OpenStorage device options

You can set the following options for an OpenStorage device.

See [“Adding an OpenStorage device”](#) on page 1574.

Table P-6 Advanced OpenStorage device options

Item	Description
Low space threshold	<p>Indicates the number at which Backup Exec suspends the job when the low space threshold is met.</p>

Table P-6 Advanced OpenStorage device options (*continued*)

Item	Description
Allow backup data to be sent directly to this device	<p>Enables a remote computer that is configured as a Remote Agent for deduplication to send data directly to the device. By using this option, the media server is bypassed, which leaves the media server free to perform other operations.</p> <p>If the OpenStorage device supports client-side deduplication, then this option enables Backup Exec to perform client-side deduplication. Note that client-side deduplication is CPU-intensive.</p> <p>If you enable this option, you must also do the following in backup jobs:</p> <ul style="list-style-type: none"> ■ Select resources on the remote computer as a backup selection. ■ Select the OpenStorage device as the destination for the backup job. ■ Select the Use client-side deduplication option. <p>This option appears in the Device and Media pane on the Backup Job Properties dialog box.</p> <p>See “About client-side deduplication” on page 1590.</p>
Split data stream every	Indicates the size at which you want Backup Exec to span to a new image. The default size is 50 GB.
Data stream chunk size	Indicates the size of a single write operation that Backup Exec issues. The default size varies based on the type of device that is being used.
Enable stream handler	Indicates whether stream handler is used. Backup Exec sets this option automatically when you select a server type. For some types of devices, this option does not appear at all. If Backup Exec does not set this option, contact the device's vendor for the recommended setting.

Viewing properties for OpenStorage devices

You can view all of the properties of an OpenStorage device and you can change some of the properties.

To view properties for OpenStorage devices

- 1 On the navigation bar, click **Devices**.
- 2 Select the device.
- 3 In the task pane, under **General Tasks**, click **Properties**.

See [“General OpenStorage Device Properties”](#) on page 1578.

See [“Advanced OpenStorage device options”](#) on page 1576.

General OpenStorage Device Properties

You can view all of the general properties of an OpenStorage device and you can change some of the properties.

See [“Viewing properties for OpenStorage devices”](#) on page 1578.

Table P-7 General OpenStorage Device Properties

Item	Description
Server	Indicates the fully-qualified name of the server on which the device exists.
Description	Indicates the description of the device.
Server type	Indicates the type of OpenStorage device.
Logon account	Indicates the name of the logon account that is required to access the device.
Paused	Lets you pause or resume the device.
Enabled	Lets you enable or disable the device.
Online	Indicates if the device is online. You cannot change this property.
Low Disk Space	Indicates if the device is low on disk space. You cannot change this property.
Allow x concurrent operations for this device	Indicates the maximum number of jobs that you want to run at the same time on this device.

Table P-7 General OpenStorage Device Properties (*continued*)

Item	Description
Total Capacity	Shows the total amount of storage space that is available on this device.
Used Capacity	Shows the total amount of storage space that is being used on this device.
Deduplication Ratio	Indicates the ratio of the amount of data before deduplication to the amount of data after deduplication.

About deduplication storage folders

A deduplication storage folder is a disk-based backup folder that you can use as a destination for backup jobs. When you use a deduplication storage folder, only unique data is stored.

Before you create a deduplication storage folder, you should review the requirements. Symantec recommends a dedicated volume and a large amount of RAM for a deduplication storage folder.

See [“Requirements for the Deduplication Option”](#) on page 1572.

When you create a deduplication storage folder, Backup Exec installs and configures a database that manages the deduplication process. You can store the deduplication storage folder and the database on the same volume or on separate volumes. However, storing the deduplication storage folder and the database on separate volumes improves the database's performance.

Note: You can create only one deduplication storage folder on a media server.

If you use Backup Exec Central Admin Server Option or SAN Shared Storage Option, a deduplication storage folder can be shared between multiple media servers. Sharing can be enabled when you add a deduplication storage folder. You can select new media servers to share a deduplication storage folder or remove the sharing ability for media servers at any time.

A deduplication storage folder can be created on a storage array. However, if a deduplication storage folder already exists on a media server, then another folder cannot be added to a storage array that is connected to that media server. To add a deduplication storage folder, you must use the Storage Array Configuration Wizard.

After you create a deduplication storage folder, it appears in the **Devices** view under the name of the Backup Exec media server. However, a deduplication storage folder does not appear in the **All Devices** device pool. A deduplication storage folder cannot belong to any device pools to prevent a deduplication job from being sent to a non-deduplication device in a device pool if the deduplication storage folder is busy.

You can pause, enable, disable, rename, refresh, and delete a deduplication storage folder. When you use Backup Exec's **Delete** option on a deduplication storage folder, the folder is removed from the Backup Exec Database. However, the folder and the files in it remain on the disk.

Note: When you delete media from a deduplication storage folder, it may take up to 48 hours for more space to become available. Backup Exec cannot always calculate the amount of space that will be made available.

See [“Adding a deduplication storage folder”](#) on page 1580.

See [“About sharing storage”](#) on page 444.

See [“About using the Storage Array Configuration Wizard”](#) on page 2021.

Adding a deduplication storage folder

A deduplication storage folder is a disk-based backup folder that you can use as a destination for backup jobs. You can add only one deduplication storage folder on a media server.

Note: To add a deduplication storage folder to a storage array, use the Storage Array Configuration Wizard.

See [“About using the Storage Array Configuration Wizard”](#) on page 2021.

See [“About deduplication storage folders”](#) on page 1579.

To add a deduplication storage folder

- 1 On the navigation bar, click **Devices**.
- 2 In the task pane, under **Device Tasks**, click **Configure devices assistant**.
- 3 On the **Configure Devices Assistant** dialog box, under **Deduplication Option**, click **Deduplication Storage Folder**.

- 4 If the **Deduplication Storage Folder Configuration** dialog box appears, click **Add a Deduplication Storage Folder**.

This step does not apply if this is the first deduplication storage folder. The **Deduplication Storage Folder Configuration** dialog box appears only if a deduplication storage folder already exists on another media server.

- 5 Complete the options on the **General** tab.
See [“General options for deduplication storage folders”](#) on page 1581.
- 6 Complete the options on the **Advanced** tab.
See [“Advanced options for deduplication storage folders”](#) on page 1582.
- 7 On the **Sharing** tab, select the media servers that you want to use with this deduplication storage folder.

The **Sharing** tab appears only if you have the Backup Exec Central Admin Server Option or the SAN Shared Storage Option installed. You must also have more than one media server.
- 8 Click **OK**.
- 9 Restart the Backup Exec services on the media servers that you selected in step 7.
See [“Starting and stopping Backup Exec services”](#) on page 168.

General options for deduplication storage folders

The following options are available for deduplication storage folders.

See [“Adding a deduplication storage folder”](#) on page 1580.

Table P-8 General options for deduplication storage folders

Item	Description
Name	Indicates the unique name of the folder.
Storage path	Indicates the location of the folder on the computer. Symantec strongly recommends that you use a dedicated volume. Note: After you select a storage path, you cannot change it. Also, you cannot use the root directory. You must use a path.

Table P-8 General options for deduplication storage folders (*continued*)

Item	Description
Database path	<p>Indicates the location where you want to place the database that is installed when you create the deduplication storage folder. You may want to use a separate volume for the database path. This option is not required. However, using a separate path may improve the database's performance. If you do not select a different volume for the database path, then Backup Exec uses the same path that you entered in the Storage path field.</p> <p>Note: After you set a database path, you cannot change it.</p>
Allow x concurrent operations for this device	<p>Indicates the maximum number of jobs that you want to run at the same time on this device.</p>
Logon account	<p>Indicates the name of the logon account that is required to access this device.</p> <p>Note: When you specify a Backup Exec logon account for a deduplication storage folder, an additional user account is created for the deduplication components with the same user name and password. However, if you change the credentials for the Backup Exec logon account, the credentials for the additional user account are not changed automatically. You must use the <code>spausr.exe</code> utility to update the password for the additional user account.</p> <p>See “Changing the password for the logon account for a deduplication storage folder” on page 1588.</p>

Advanced options for deduplication storage folders

The following options are available for deduplication storage folders.

See [“Adding a deduplication storage folder”](#) on page 1580.

Table P-9 Advanced options for deduplication storage folders

Item	Description
Low space threshold	Indicates the number at which Backup Exec moves the job to a new device or suspends the job if this device is not in a device pool.
Allow remote agents to use client-side deduplication	<p>Enables a remote computer that is configured as a Remote Agent for deduplication to send data directly to the deduplication storage folder. After the data is deduplicated, then only unique data is sent directly to the deduplication storage folder. By using this option, the media server is bypassed, which leaves the media server free to perform other operations.</p> <p>If you enabled this option, you must also do the following in backup jobs:</p> <ul style="list-style-type: none"> ■ Select resources on the remote computer as a backup selection. ■ Select the deduplication storage folder as the destination for the backup job. ■ Select the Use client-side deduplication option. <p>This option appears in the Device and Media pane on the Backup Job Properties dialog box.</p> <p>See “About client-side deduplication” on page 1590.</p>
Data stream chunk size	Indicates the size of a single write operation that Backup Exec issues. The default size varies based on the type of device being used.
Log Level	Indicates the type of information you want to include in the diagnostic logs for this device. The choices range from critical errors only to all types of messages.
Log Retention	Indicates the number of days you want to keep the logs for this device.

Deleting a deduplication storage folder

When you use Backup Exec's **Delete** option on a deduplication storage folder, the folder is removed from the Backup Exec Database. However, the folder and the files in it remain on the disk.

To delete a deduplication storage folder

- 1 On the navigation bar, click **Devices**.
- 2 Select the deduplication storage folder.
- 3 On the task pane, under **General Tasks**, click **Delete**.
- 4 On the **Delete Deduplication Storage Folder** dialog box, click **Yes**.

See [“About deduplication storage folders”](#) on page 1579.

Viewing properties of a deduplication storage folder

You can view all of the properties of a deduplication storage folder and you can change some of the properties.

To view properties of a deduplication storage folder

- 1 On the navigation bar, click **Devices**.
- 2 Select the deduplication storage folder.
- 3 In the task pane, under **General Tasks**, click **Properties**.

See [“General Deduplication Storage Folder Properties”](#) on page 1584.

See [“Advanced options for deduplication storage folders”](#) on page 1582.

General Deduplication Storage Folder Properties

You can view all of the general properties of a deduplication storage folder and you can change some of the properties.

See [“Viewing properties of a deduplication storage folder”](#) on page 1584.

Table P-10 General Deduplication Storage Folder Properties

Item	Description
Name	Indicates the name that was entered when the deduplication storage folder was configured. It may be a user-defined name or the default name that Backup Exec entered. You can change the name at any time.

Table P-10 General Deduplication Storage Folder Properties (*continued*)

Item	Description
Server	Indicates the name of the computer on which the deduplication storage folder was created.
Storage path	Indicates the location of the folder on the computer. Symantec strongly recommends that you use a dedicated volume.
Database path	Indicates the location of the database that was installed when you created the deduplication storage folder. You cannot change the location of the database.
Paused	Lets you pause or resume the device.
Enabled	Lets you enable or disable the device.
Online	Shows whether the device is online or offline. If a check does not appear in the check box, then the device is offline. You cannot change this property.
Low Disk Space	Indicates that the device is low on disk space.
Allow x concurrent operations for this device	Indicates the maximum number of jobs that you want to run at the same time on this device.
Total Capacity	Shows the total amount of storage space that is available on this device.
Used Capacity	Shows the total amount of storage space that is being used on this device.
Deduplication Ratio	Indicates the ratio of the amount of data before deduplication to the amount of data after deduplication.
Hardware Name	Indicates the name of the storage array on which the deduplication folder is located. This property appears only if the deduplication storage folder was added to a storage array with the Storage Array Configuration Wizard.

Table P-10 General Deduplication Storage Folder Properties (*continued*)

Item	Description
Hardware Status	<p>Indicates the status of the storage array hardware. This property appears only if the deduplication storage folder was added to a storage array with the Storage Array Configuration Wizard</p> <p>The values for the hardware status are as follows:</p> <ul style="list-style-type: none">■ OK The storage array is online.■ Offline The storage array and the virtual disks on the storage array are offline. Backup Exec cannot access them. <p>Refer to the vendor documentation and management software that are supplied with the storage array. After the storage array is brought online, the virtual disks are automatically brought online.</p>

Table P-10 General Deduplication Storage Folder Properties (*continued*)

Item	Description
Hardware Health	<p>Displays one of the following values to indicate the health of the hardware:</p> <ul style="list-style-type: none">■ OK The storage array is online.■ Warning The storage array may fail or produce errors, but it is currently operational. The storage array and the virtual disks on the storage array are offline. Backup Exec cannot access them.■ Critical The storage array has failed. The storage array and the virtual disks on the storage array are offline. Backup Exec cannot access them. <p>Refer to the vendor documentation and management software that are supplied with the storage array. After the storage array is brought online, the virtual disks are automatically brought online.</p> <p>This property appears only if the deduplication storage folder was added to a storage array with the Storage Array Configuration Wizard</p>

Table P-10 General Deduplication Storage Folder Properties (*continued*)

Item	Description
Disk Classification	<p>Describes the RAID level for the storage array.</p> <p>Disk classifications are as follows:</p> <ul style="list-style-type: none"> ■ Simple (RAID 0) A single physical disk, no striping, or parity. No redundancy/ ■ Span A set of multiple physical disks that are concatenated together. No striping or parity. No redundancy. ■ Stripe A set of multiple physical disk extents with data striped across the physical disks. No redundancy. ■ Mirror (RAID 1) A pair or multiple pairs of physical disks with the same data written to each physical disk of the pair. Provides for data redundancy/ ■ Stripe with parity (RAID 5 or RAID 6) Three or more physical disks with data striped across the physical disks, with one disk's worth of space that is used for parity. Provides for data redundancy. ■ Unknown <p>This property appears only if the deduplication storage folder was added to a storage array with the Storage Array Configuration Wizard.</p>
Logon account	<p>Indicates the logon account that is being used to access the device. You can change the logon account at any time.</p>

Changing the password for the logon account for a deduplication storage folder

When you specify a Backup Exec logon account for a deduplication storage folder, an additional user account is created for the deduplication components with the same user name and password. However, if you change the credentials for the

Backup Exec logon account, the credentials for the additional user account are not changed automatically. You must use the `spausser.exe` utility to update the password for the additional user account. This account is known as the "User 1" account when you use the `spausser.exe` utility to view a list of user names that are associated with the deduplication storage folder.

To change the password for the logon account for a deduplication storage folder

- 1 On the Backup Exec **Network** menu, click **Logon Accounts**.
- 2 Select the Backup Exec logon account that you want to change, and then click **Edit**.
- 3 Click **Change Password**.
- 4 Type the new password.
- 5 Click **OK**.
- 6 At a command prompt, type the following command:

```
spausser.exe -c -u <UserName>
```

The user name is case-sensitive . If you do not know the user name, type the following command to find the user name that is associated with "User 1":

```
spausser.exe -l
```

You will be prompted for the old password and a new password. Be sure that the new password is the same as the password that you used in step 4.

See [“About deduplication storage folders”](#) on page 1579.

Sharing a deduplication device between multiple media servers

If you use the Backup Exec Central Admin Server Option or the SAN Shared Storage Option, you can select which media servers can share a deduplication storage folder, an OpenStorage device, or a Remote Agent for Deduplication. When you add a deduplication storage folder, an OpenStorage device, or a Remote Agent for Deduplication, the media server that you used to add the device is automatically selected for sharing.

Note: To share a deduplication storage folder, you must add it as an OpenStorage device on all media servers that you want to access the folder, except for the media server that was used to create it.

See [“About sharing storage”](#) on page 444.

To share a deduplication device between multiple media servers

- 1 On the navigation bar, click **Devices**.
- 2 In the **Devices** view, right-click the deduplication storage folder, the OpenStorage device, or the Remote Agent for Deduplication that you want media servers to access.
- 3 Select **Manage sharing**.
The **Manage sharing** option appears only if you have the Backup Exec Central Admin Server Option or the SAN Shared Storage Option installed. You must also have more than one media server.
- 4 Select the deduplication device that you want to share.
- 5 Under **Media Servers**, select the media servers that you want to use with the deduplication device.
- 6 Click **OK**.
- 7 Restart the Backup Exec services on the media servers that you selected in step 5.

About client-side deduplication

Client-side deduplication enables a remote computer that is configured as a Remote Agent for deduplication to send data directly to an OpenStorage device or a deduplication storage folder. By using client-side deduplication, the media server is bypassed, which leaves the media server free to perform other operations. If your deduplication device supports client-side deduplication, a remote computer deduplicates data and then sends only the unique data directly to a deduplication storage folder or an OpenStorage device.

Note: Client-side deduplication may increase the CPU utilization on the remote computer if your deduplication device supports client-side deduplication.

When you create a backup job with client-side deduplication, keep in mind the following items:

- The backup job can include resources from only one remote computer.
- The Remote Agent for Windows Systems must be installed and running on the remote computer.
- The remote computer must be configured as a Remote Agent for Deduplication.
- The remote computer must be pingable.

- The remote computer cannot be a Backup Exec media server.
- A deduplication storage folder or an OpenStorage device must be selected in the **Device and Media** view for the backup job.
- The option **Use client-side deduplication** must be selected in the **Device and Media** view for the backup job. This option is selected by default when you select a deduplication storage folder or an OpenStorage device as the destination for a backup job.
- The Backup Exec service account enables remote computers to directly access a Symantec PureDisk device. If you want to perform backups that are enabled for Granular Recovery Technology, the Backup Exec service account must be valid for any remote computer that directly accesses the Symantec PureDisk device. You must verify that the remote computers are in the same domain, or that the remote computers have a domain trust relationship with the domain in which the media server resides.

If you do not configure the remote computer to use client-side deduplication, then the data from the remote computer is sent to the media server to be deduplicated. Then, the deduplicated data is backed up to the deduplication storage folder or the OpenStorage device. This process increases the CPU utilization on the media server. However, this process is useful if you are backing up older remote computers.

See [“Configuring client-side deduplication”](#) on page 1591.

See [“Configuring a Remote Agent for Deduplication”](#) on page 1592.

Configuring client-side deduplication

Client-side deduplication enables a remote computer that is configured as a Remote Agent for Deduplication to send data directly to an OpenStorage device or a deduplication storage folder.

See [“About client-side deduplication”](#) on page 1590.

Table P-11 How to configure client-side deduplication

Action	Notes	For more information
Configure an OpenStorage device or a deduplication storage folder to be directly accessed.	On the Advanced tab, select Allow backup data to be sent directly to this device .	See “Adding a deduplication storage folder” on page 1580. See “Adding an OpenStorage device” on page 1574.

Table P-11 How to configure client-side deduplication (*continued*)

Action	Notes	For more information
Configure a Remote Agent for Deduplication.	<p>Set up the remote computer to directly access OpenStorage devices and deduplication storage folders for backups.</p> <p>Note: The remote computer must have the Remote Agent for Windows Systems installed on it.</p>	See “Configuring a Remote Agent for Deduplication” on page 1592.
Create a backup job.	<p>In the Selections view, select resources from the remote computer that is set up as a Remote Agent for Deduplication.</p> <p>In the Device and Media view, select an OpenStorage device or a deduplication storage folder as the destination device. Then, verify that Use client-side deduplication is selected.</p>	See “Device and media options for backup jobs and templates” on page 339.

Configuring a Remote Agent for Deduplication

You can set up a remote computer to allow backup data to be sent directly to an OpenStorage device or a deduplication storage folder. The remote computer must have the Remote Agent for Windows Systems installed on it. For direct access to a third-party OpenStorage device, the vendor plug-in for the device must also be installed. The appropriate plug-in for a Symantec PureDisk device is included in the Remote Agent for Windows Systems, so no additional plug-in is required.

Note: The option to configure a Remote Agent for Deduplication appears only if you have already configured an OpenStorage device or a deduplication storage folder.

See [“About client-side deduplication”](#) on page 1590.

See [“Configuring client-side deduplication ”](#) on page 1591.

To configure a Remote Agent for Deduplication

- 1 On the navigation bar, click **Devices**.
- 2 In the task pane, under **Device Tasks**, click **Configure devices assistant**.
- 3 On the **Configure Devices Assistant** dialog box, click **Remote Agent Deduplication**.
- 4 If the **Remote Agent Deduplication Configuration** dialog box appears, click **Add a Remote Agent for Deduplication**.

This step does not apply if this is the first Remote Agent for Deduplication. The **Remote Agent Deduplication Configuration** dialog box appears only if a Remote Agent for Deduplication already exists.

- 5 Complete the options to set up the remote agent.
See [“General options for a Remote Agent for Deduplication”](#) on page 1593.
- 6 On the **Sharing** tab, select each media server that you want to share this Remote Agent for Deduplication.

The **Sharing** tab appears only if you have the Backup Exec Central Admin Server Option or the SAN Shared Storage Option installed. You must also have more than one media server.

- 7 Click **OK**.
- 8 Restart the Backup Exec services on the media servers that you selected in step 6.
See [“Starting and stopping Backup Exec services”](#) on page 168.

General options for a Remote Agent for Deduplication

The following options are available for a Remote Agent for Deduplication.

See [“Configuring a Remote Agent for Deduplication”](#) on page 1592.

Table P-12 General options for a Remote Agent for Deduplication

Item	Description
Server	<p>Indicates the name of the computer that you want to add as a Remote Agent for Deduplication.</p> <p>Note: The naming format that you use to enter the computer name must also be used to select the remote computer for backup. For example, if you use the IP address here, you must also use the IP address for the backup selection. Otherwise, client-side deduplication does not occur.</p>
Port	<p>Indicates the port to use for communications between the media server and the remote computer.</p>
Description	<p>Displays a description that you choose.</p>
Enable ICMP ping operations for Backup Exec to detect the server	<p>Lets the media server use ICMP ping to locate the remote computer.</p>
Logon account	<p>Indicates the logon account that is required to access the remote computer.</p>

Viewing properties of a Remote Agent for Deduplication

You can view all of the properties of a Remote Agent for Deduplication and you can change some of the properties.

To view properties of a Remote Agent for Deduplication

- 1 On the navigation bar, click **Devices**.
- 2 Select the Remote Agent for Deduplication.
- 3 In the task pane, under **General Tasks**, click **Properties**.

See [“Remote Agent for Deduplication properties”](#) on page 1594.

Remote Agent for Deduplication properties

You can view all of the properties of a Remote Agent for Deduplication and you can change some of the properties.

See [“Viewing properties of a Remote Agent for Deduplication”](#) on page 1594.

Table P-13 Remote Agent for Deduplication properties

Item	Description
Server	Indicates the name of the computer that you want to add as a Remote Agent for Deduplication.
Port	Indicates the port to use for communications between the media server and the remote computer.
Description	Displays a user-defined description of the remote agent.
Enable ICMP ping operations for Backup Exec to detect the server	Lets the media server use ICMP ping to locate the remote computer.
Logon account	Indicates the logon account that is required to access the remote computer.

About backup jobs for deduplication

You set up a backup job for deduplication in the same way as you set up a regular backup job. When you select either an OpenStorage device or a deduplication storage folder as the destination device, then deduplication occurs when the job runs. Optionally, if you want the remote agent to send data directly to the device, you can select the option for client-side deduplication. No other additional options are necessary to create a backup job for deduplication.

See [“About client-side deduplication”](#) on page 1590.

About optimized duplication

Backup Exec supports optimized duplication, which enables deduplicated data to be copied directly from one OpenStorage device to another OpenStorage device from the same vendor. For example, you can copy data from one Symantec PureDisk device to another Symantec PureDisk device. If you use the Central Admin Server Option, you can also copy data from a deduplication storage folder on a managed media server to a deduplication storage folder on another managed media server. The data is copied over the network, thereby avoiding the Backup Exec media server. Because the data is deduplicated, only unique data is copied between the devices.

Note: Optimized duplication is not available for backup sets that are enabled for Granular Recovery Technology.

Optimized duplication is available for OpenStorage devices from selected vendors. You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

To copy data between OpenStorage devices, you must create a duplicate backup job. The destination device for the duplicate job must be the same type of device from the same vendor as the device that was used in the source backup job. You can restore data from either device.

See “[Setting up optimized duplication](#)” on page 1596.

Setting up optimized duplication

Optimized duplication enables deduplicated data to be copied directly from one OpenStorage device to another OpenStorage device from the same vendor. You can also copy data from one deduplication storage folder to another deduplication storage folder.

See “[About optimized duplication](#)” on page 1595.

You set up a duplicate backup job to perform optimized duplication.

Table P-14 How to set up optimized duplication

Step	For more information
Create a backup job that uses an OpenStorage device or a deduplication storage folder as the destination.	See “ Creating a backup job by setting job properties ” on page 332.
Create a duplicate backup job and select the appropriate OpenStorage device or deduplication storage folder as the destination. Note: The destination device for the duplicate job must be the same type of device from the same vendor as the device that was used in the source backup job.	See “ Duplicating backed up data ” on page 369.

About copying deduplicated data to tapes

Backup Exec lets you copy deduplicated data from an OpenStorage device to tape for long-term or off-site storage. When data is copied to tape, it is rehydrated. In other words, the files are reassembled into their original form and are not deduplicated.

To copy deduplicated data to tapes, you must create a duplicate backup job that copies the backup sets from the OpenStorage device to a tape device.

See [“Duplicating backed up data”](#) on page 369.

About using deduplication with encryption

You should not use the Backup Exec encryption options for backup jobs that deduplicate data. Data cannot be deduplicated when it is encrypted.

See [“About encryption”](#) on page 415.

About restoring deduplicated data

You set up a restore job to restore deduplicated data in the same way as you set up a regular restore job. No additional settings are required.

See [“About disaster recovery of a deduplication storage folder”](#) on page 1597.

See [“About disaster recovery of OpenStorage devices”](#) on page 1599.

About disaster recovery of a deduplication storage folder

A deduplication storage folder is stored on the Backup Exec media server. If your media server experiences a disaster, then the data from the deduplication storage folder is lost. Therefore, you should take steps to prepare for recovery from a system failure. To prepare for a disaster, Backup Exec lets you take a snapshot of a deduplication storage folder. The snapshot includes the folder, the contents of the folder, and the associated database for the folder. You can store the snapshot on tape, which you can then use to recover your deduplication storage folder after a disaster.

When you restore data from the snapshot, the following processes occur:

- Backup Exec stops the deduplication services if they are running. The deduplication services are separate from the Backup Exec services, so the Backup Exec services are not affected.

- Backup Exec deletes any files that are present in the deduplication storage folder and in the associated database.
- The deduplication storage folder is restored to its original location, along with the contents of the folder and the associated database.
- The deduplication services are restarted.

See [“Preparing for disaster recovery of a deduplication storage folder”](#) on page 1598.

Preparing for disaster recovery of a deduplication storage folder

To prepare for a disaster, Backup Exec lets you take a snapshot of a deduplication storage folder. The snapshot includes the folder, the contents of the folder, and the database for the folder.

See [“About disaster recovery of a deduplication storage folder”](#) on page 1597.

See [“Performing disaster recovery of a deduplication storage folder”](#) on page 1598.

To prepare for disaster recovery of a deduplication storage folder

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Select **New Backup Job**.
- 3 In the backup selection list, under **Shadow Copy Components**, expand **User Data**, and then select **Backup Exec Deduplication Storage**.
- 4 In the task pane, under **Destination**, click **Device and Media**.
- 5 In the **Device** field, select a tape device.
- 6 Complete any additional options that you want to use.
- 7 Run the job.

Performing disaster recovery of a deduplication storage folder

If a disaster occurs on the media server where a deduplication storage folder is stored, you can recover the folder and its associated database after you recover the server. If you use the Backup Exec Intelligent Disaster Recovery Option (IDR Option) to recover the media server, you must recover the deduplication storage folder using the following procedure. The IDR Option does not recover the deduplication storage folder or its associated database during the recovery of the media server.

When you restore a deduplication storage folder, the original folder and its associated database are deleted and then replaced by the restored folder and database.

Note: You cannot redirect the restore of a deduplication storage folder. A deduplication storage folder must be restored to its original location.

See [“About disaster recovery of a deduplication storage folder”](#) on page 1597.

See [“Preparing for disaster recovery of a deduplication storage folder”](#) on page 1598.

To perform a disaster recovery of a deduplication storage folder

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Select **New Restore Job**.
- 3 In the restore selection list, select **Shadow Copy Components**.
- 4 Complete any additional options that you want to use for this restore job.
- 5 Run the restore job.
The job will complete with exceptions.
- 6 When the restore job completes, restart the Backup Exec services.

About disaster recovery of OpenStorage devices

The following disaster recovery scenarios are possible for OpenStorage devices:

- The device fails.
- The media server that uses the device fails.

If the device fails, you should consult the documentation from the device's vendor. If the media server fails and you need to reinstall Backup Exec on the media server, you must reconfigure the device, and inventory and catalog the media from it after the media server is recovered.

See [“About disaster recovery of a deduplication storage folder”](#) on page 1597.

Symantec Backup Exec Desktop and Laptop Option

This appendix includes the following topics:

- [About the Desktop and Laptop Option](#)
- [About the components of DLO](#)
- [Before you install DLO](#)
- [System requirements for the DLO Administration Console](#)
- [About installing the Backup Exec Desktop and Laptop Option](#)
- [About setting a recovery password](#)
- [Checking data integrity](#)
- [Changing DLO service credentials](#)
- [About administrator accounts in DLO](#)
- [About automated permissions management in DLO](#)
- [About limited restore in DLO](#)
- [Using a list of individual accounts to manage DLO permissions](#)
- [Using domain groups to manage DLO permissions](#)
- [About default DLO settings](#)
- [Changing default DLO profile settings](#)
- [Changing default DLO backup selection settings](#)

- Changing default DLO global settings
- Configuring DLO to use a specific port for database access
- About using Backup Exec Retrieve with DLO
- About updating DLO
- Starting the DLO Administration Console from Backup Exec
- About the DLO Overview view
- Connecting to DLO on a different Backup Exec Media Server
- How to configure DLO
- About DLO profiles
- About backup selections in DLO
- About Delta File Transfer
- About DLO Storage Locations
- About Automated User Assignments
- About configuring global exclude filters in DLO
- About managing Desktop Agent users
- Modifying computer properties
- Enabling or disabling a desktop computer
- Deleting a desktop computer from DLO
- Backing up a desktop from the DLO Administration Console
- Restoring files and folders from the DLO Administration Console
- Searching for files and folders to restore with DLO
- About DLO emergency restore and recovery passwords
- Computer History pane options and Job History pane options
- About monitoring alerts on the DLO Administration Console
- About configuring notification methods for DLO alerts
- About configuring recipients for notification in DLO
- About DLO reports

- [About maintaining the DLO database](#)
- [About clustering the Desktop and Laptop Option](#)
- [About the DLO command syntax](#)
- [About the Desktop Agent](#)
- [Desktop Agent terminology](#)
- [Features and benefits of the Desktop Agent](#)
- [System requirements for the Desktop Agent](#)
- [Installing the Desktop Agent](#)
- [How to configure the Desktop Agent](#)
- [About the Desktop Agent Console](#)
- [About using the Desktop Agent to back up your data](#)
- [About modifying Desktop Agent settings](#)
- [About synchronizing desktop user data](#)
- [About the status of the Desktop Agent](#)
- [About suspending or canceling a job](#)
- [Viewing usage details](#)
- [Restoring files by using the Desktop Agent](#)
- [About using Backup Exec Retrieve to restore files](#)
- [About monitoring job history in the Desktop Agent](#)
- [About using DLO with other products](#)
- [Troubleshooting the DLO Administration Console](#)
- [Troubleshooting the Desktop Agent](#)
- [Accessibility and DLO](#)

About the Desktop and Laptop Option

The Backup Exec Desktop and Laptop Option (DLO) provides automated file protection for desktops and laptops (collectively referred to as desktops). Protection is provided whether the computer is connected to the network or offline. When

the desktop is not connected to the network, the files are backed up to a user data folder on the desktop. When the computer reconnects to the network, the files are backed up from the local desktop user data folder to the designated network user data folder.

Users who have multiple computers can synchronize the data between their computers so the most up-to-date file versions are available on all of their computers.

Note: This product is intended to provide file-level protection for desktop user data and is not intended to provide a full system backup.

About the components of DLO

DLO is comprised of the following components:

Table Q-1 Components of DLO

Component	Description
DLO Administration Console	<p>The DLO Administration Console is part of Backup Exec and runs on the Backup Exec media server. The DLO Administration Console runs in a separate window that you access from Backup Exec.</p> <p>From the DLO Administration Console, the Administrator can do the following:</p> <ul style="list-style-type: none"> ■ Create profiles, which control the desktop user’s level of interaction with the Desktop Agent. In addition, you can use profiles to define the types of files to back up, and to set the backup schedule. ■ Create network user data folders, which are locations where data is stored. ■ Create Automated User Assignments, which determine the storage location and profile to which users are assigned. ■ Add users to DLO manually. ■ View history log files, receive alerts, and restore files to a desktop.
DLO database	<p>The DLO database is part of Backup Exec and runs on the Backup Exec media server.</p>
DLO Maintenance Service	<p>The maintenance server is installed by default when DLO is installed. Only one maintenance server is required. However, in large installations it may be more efficient to have one maintenance server for each Storage Location host (File Server).</p>

Table Q-1 Components of DLO (*continued*)

Component	Description
Desktop Agent	<p>The Desktop Agent resides on the desktops and laptops that you want to protect. The Desktop Agent may run in the background, automatically protecting files.</p> <p>Alternatively, desktop users with full access to the Desktop Agent interface can do the following:</p> <ul style="list-style-type: none"> ■ Schedule backups ■ Select which types of files to back up ■ Restore files ■ Synchronize file versions between different computers ■ View the status of backups

See [“How to configure DLO”](#) on page 1640.

See [“About DLO profiles”](#) on page 1641.

See [“About backup selections in DLO”](#) on page 1658.

See [“About DLO Storage Locations”](#) on page 1676.

See [“About Automated User Assignments”](#) on page 1683.

Before you install DLO

Before you install DLO, you should consider the following:

Table Q-2 Pre-installation considerations

Item	Description
Domains and Active Directory	The media server and DLO Storage Locations must be in a Windows Domain or Active Directory. Computers that run the Desktop Agent can be outside of a Windows Domain or Active Directory. However, they must authenticate with the domain or directory to access the media server or Storage Locations.

Table Q-2 Pre-installation considerations (*continued*)

Item	Description
Server loading	<p>DLO can be treated as a network file server. The ideal server for DLO has a fast network connection and a fast set of disks. The CPU is not as critical as these other factors for the DLO file server.</p> <p>The number of Desktop Agents that can successfully back up to one DLO server depends on many factors. However, Symantec recommends less than 400 clients per server when the server runs WIndows 2000 Advanced Server. If more than 400 clients are attached, file operations may begin to fail when Paged Pool memory runs out.</p>
Authentication	<p>DLO Administration Console</p> <p>Any user who has full admin rights on the media server where DLO is installed can manage the DLO Administration Console.</p> <p>The user's account must be a domain account.</p> <p>In addition, the account must have rights to do the following on any remote server used for Storage Locations or network user data folders:</p> <ul style="list-style-type: none"> ■ Create network shares ■ Manage permissions of network shares and directories <p>You can use a domain administrator account, or grant a standard domain account with local administrative rights to the servers hosting the DLO resources.</p> <p>See “About administrator accounts in DLO” on page 1618.</p> <p>Desktop Agent</p> <p>DLO requires domain accounts. Every Desktop Agent user must log on to DLO using a domain account. If you have users who log on using local accounts, they can still use DLO, but they must have domain credentials to authenticate with DLO.</p>

Table Q-2 Pre-installation considerations (*continued*)

Item	Description
Database Selection	<p>By default DLO installs its own instance of SQL Express 2005.</p> <p>Note: If you use an existing database instance, named pipes must be enabled. If DLO installs its own SQL Express 2005 instance, named pipes are enabled automatically.</p> <p>You can also manually configure DLO to use an existing local or remote instance of the following:</p> <ul style="list-style-type: none"> ■ SQL Express 2005 ■ SQL Server 2005 ■ MSDE 2000 <p>The following are pros for SQL Express 2005:</p> <ul style="list-style-type: none"> ■ Free ■ Unless you back up more than 1000 Desktop Agents per media server, SQL Express should be sufficient for most needs. <p>The following are cons for SQL Express 2005:</p> <ul style="list-style-type: none"> ■ The database is limited to a single processor, resulting in slower I/O to the database under load. ■ A maximum table size of 4 GB for SQL Express, although DLO is unlikely to reach this limit. <p>The following are pros for SQL Server:</p> <ul style="list-style-type: none"> ■ Allows reasonable scalability beyond 1000 Desktop Agents. ■ Database tools are included with SQL Server. <p>The following are cons for SQL Server:</p> <ul style="list-style-type: none"> ■ Cost, however you do not need to purchase a SQL Server client license for each Desktop Agent. <p>The following are pros for MSDE:</p> <ul style="list-style-type: none"> ■ Free ■ Unless you back up more than 1000 Desktop Agents per media server, MSDE should be sufficient for most needs. <p>The following are cons for MSDE:</p> <ul style="list-style-type: none"> ■ The number of concurrent connections to the database is limited, resulting in slower I/O to the database under load. ■ A maximum table size of 2 GB for MSDE, although DLO is unlikely to reach this limit.

Table Q-2 Pre-installation considerations (*continued*)

Item	Description
Time Synchronization	<p>All computers running the DLO Administration Console or the Desktop Agent should be set to a common time. You can configure the Windows Time Synchronization service on the network.</p> <p>See the Microsoft Web site for additional information.</p>
Firewalls	<p>DLO is designed to work in firewall environments. In order for DLO to function properly in a firewall environment, network file shares must be visible after establishing a remote connection such as VPN. If file sharing is not allowed, DLO does not transfer files to or from the network user data folder. Desktop computer files are still protected to the desktop user data folder, and are transferred when the network user data folder is accessible.</p> <p>You can push-install DLO through a firewall to a computer that runs Windows XP Professional with Service Pack 2. Before you begin the push-install, enable the group policy called "Allow remote administration exception" on the destination computer. See your Windows documentation for more information.</p> <p>You can also push-install DLO through a firewall to a Windows Server 2008 computer. Before you begin the push-install, enable Print and File Sharing, along with Windows Management Instrumentation (WMI) on the destination computer's Windows Firewall Exception list. See your Windows documentation for more information.</p>
MDAC Support for the DLO Administration Console	<p>DLO supports versions 2.7 and 2.8 of MDAC. However, MDAC 2.8 is the default and is installed during the installation if it is not already installed. When MDAC 2.8 is installed, a restart may be required, and a computer administrator must complete the installation process. If a non-administrator logs on first after the MDAC 2.8 installation, the process generates a number of errors.</p> <p>To force the use of MDAC 2.7, you can add the following CmdLine value in the setup.ini file in the DLO installation set:</p> <p>REQUIREDMDACVERSION="2.7".</p> <p>Example:</p> <p>CmdLine=REQUIREDMDACVERSION="2.7" /!*v %TEMP%\DLOConsoleInstall.log</p>

Table Q-2 Pre-installation considerations (*continued*)

Item	Description
MDAC Support for the Desktop Agent	<p>The Desktop Agent supports versions 2.7 and 2.8 of MDAC. However, MDAC 2.8 is the default.</p> <p>To force the use of MDAC 2.7, you can add the following CmdLine value in the setup.ini file in the Desktop Agent installation set:</p> <p>REQUIREDMDACVERSION="2.7".</p> <p>Example:</p> <pre>CmdLine=/qf DEFAULTMEDIASERVER="SERVERNAME" DLOBBINSTANCENAME="BKUPEXEC" LAUNCHCLIENT="1" REQUIREDMDACVERSION="2.7" /!*v "%TEMP%\DLOAgentInstall.log"</pre>

System requirements for the DLO Administration Console

The following are the minimum system requirements for running this version of the DLO Administration Console.

Table Q-3 Minimum system requirements

Item	Description
Operating System	<p>The Administration Console runs on the following operating systems:</p> <ul style="list-style-type: none"> ■ Microsoft Windows Server 2003 server family ■ Microsoft Windows Server 2003 x64 Editions ■ Microsoft Windows Server 2003 R2 ■ Microsoft Windows XP Service Pack 2 or later ■ Microsoft Windows XP Professional x64 Edition ■ Microsoft Windows Vista ■ Microsoft Windows Server 2008 ■ Microsoft Windows Server 2008 R2 ■ Microsoft Windows 7 (to support remote administration) <p>You cannot install DLO on computers that run the Windows Server Core option.</p>
Internet Browser	Internet Explorer 5.01 or later; however, version 5.5 is recommended
Processor	Pentium system

Table Q-3 Minimum system requirements (*continued*)

Item	Description
Memory	Required: 256 MB RAM Recommended: 512 MB or more for better performance Recommended: 512 MB or more for better performance
Disk Space	150 MB hard disk space is required after Microsoft Windows is installed (typical installation)
Other Hardware	The following hardware is recommended: <ul style="list-style-type: none">■ Network interface card■ CD-ROM drive■ Printer that is supported by Windows (optional)■ Mouse (recommended)

Windows Domains and Active Directory are supported. Other authentication schemes, such as Novell E-Directory and NIS+, are not supported.

About installing the Backup Exec Desktop and Laptop Option

The DLO Administration Console is installed as a separate add-on component of Backup Exec. The Backup Exec media server and any Storage Locations must be in a Windows Domain or Active Directory. Novell E-Directory, NIS+ and other non-Windows Domain or Active Directory authentication schemes are not supported.

After you install the Administration Console, you can install the Desktop Agent or direct the desktop users on how to install it. Computers that run the Desktop Agent can be outside of a Windows Domain or Active Directory. However, they must authenticate with the domain or directory to access the media server or Storage Locations.

When you install DLO from the Backup Exec installation media, the Desktop Agent install set is created on the Backup Exec media server. The install set is located in a directory that is shared and available by a UNC path.

If you install DLO after Backup Exec has been clustered, you must run the cluster configuration wizard again. Backup Exec can then determine that DLO is present and can reconfigure the Backup Exec group to account for it.

See [“How to deploy the Desktop Agent”](#) on page 1611.

See [“Installing the Desktop Agent”](#) on page 1758.

See [“Before you install DLO”](#) on page 1605.

See [“Installing a custom installation of Backup Exec”](#) on page 119.

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 153.

How to deploy the Desktop Agent

You can deploy the Desktop Agent from the Desktop Agent install share to the desktop computer in the following ways.

Table Q-4 Desktop Agent deployment methods

Deployment Methods	Description
Push-install	<p>Push-install the Desktop Agent by using the Backup Exec installer.</p> <p>See “Push-installing the Remote Agent and Advanced Open File Option to remote computers” on page 134.</p> <p>After a push-install of the Desktop Agent, it may take up to two minutes for the Desktop Agent to start on the desktop computer.</p> <p>To push-install DLO to a computer that runs Windows XP with the Windows Firewall enabled, File and Print Sharing must be enabled on the desktop computer. File and Print Sharing is configured in the Windows Firewall Exceptions tab.</p>
email	<p>Send a hypertext link to the install files or include the install files as an attachment.</p>
Web page	<p>Place the install files on your company’s intranet.</p>
Logon scripts	<p>Create a file that includes commands for installing the Desktop Agent. Then assign the script to the User Properties for the employees who need to use DLO. The commands are executed automatically when the user logs on to the network. For more information about logon scripts, refer to your Microsoft Windows documentation.</p>
Microsoft Systems Management Server (SMS)	<p>Use this automated system to distribute the Desktop Agent install set to the desktop computers, which then initiate the installation. For more information about SMS, refer to your Microsoft documentation.</p>

Table Q-4 Desktop Agent deployment methods (*continued*)

Deployment Methods	Description
CD-ROM	To distribute the Desktop Agent installation files on a CD-ROM, place the contents of the \\media server\DLO Agent share on the CD-ROM. Users can then run setup.exe from the CD-ROM. The installed Desktop Agent is correctly associated with the media server.

Customizing the Desktop Agent installation

The Desktop Agent installation can be customized to meet specific needs. For example, it can run silently with no user interface displayed, or it can display either a basic or complete user interface. To customize the installation, you can modify the Setup.ini file in the DLO Agent setup directory.

For example, for a silent installation, edit CmdLine in the Setup.ini file as follows:

Original:

```
CmdLine=/qf DEFAULTMEDIASERVER="Desktop3" LAUNCHCLIENT="1" /l*v  
"%TEMP%\DLOAgentInstall.log"
```

Modified:

```
CmdLine=/qn DEFAULTMEDIASERVER="Desktop3" LAUNCHCLIENT="1" /l*v  
"%TEMP%\DLOAgentInstall.log"
```

For an installation with a basic interface but no option to cancel the installation, edit CmdLine in the Setup.ini file as follows:

Original:

```
CmdLine=/qf DEFAULTMEDIASERVER="Desktop3" LAUNCHCLIENT="1" /l*v  
"%TEMP%\DLOAgentInstall.log"
```

Modified:

```
CmdLine=/qb! DEFAULTMEDIASERVER="Desktop3" LAUNCHCLIENT="1" /l*v  
"%TEMP%\DLOAgentInstall.log"
```

To customize the Desktop Agent installation

- 1 In the Desktop Agent setup directory, open the Setup.ini file for editing.
- 2 Modify the value that begins CmdLine= /qf:

Desktop Agent installation interface	<p>Modify the /qf term to change the interface that the Desktop Agent user sees during installation of the Desktop Agent.</p> <ul style="list-style-type: none">■ /qf The full user interface appears. A cancel option is provided.■ /qb A basic progress dialog box appears. A cancel option is provided.■ /qb! A basic user interface appears. A cancel option is not provided.■ /qn The installation will be silent. <p>For a completely silent install, run the following command after you modify the Setup.ini file:</p> <pre>"setup.exe /s"</pre>
Set the Default Media Server	<p>DEFAULTMEDIASERVER specifies the media server to which you want the Desktop Agent to attach after installation.</p>
Launch the Desktop Agent	<p>The LAUNCHCLIENT option specifies whether the Desktop Agent should be launched immediately following installation.</p> <p>To launch immediately, set LAUNCHCLIENT="1"</p> <p>To prevent immediate launch, set LAUNCHCLIENT="0"</p>
Specify the Required MDAC Version	<p>DLO uses MDAC 2.8 by default. Force the use of MDAC 2.7 by adding the following:</p> <pre>REQUIREDMDACVERSION=2.7</pre>
Suppress Reboot	<p>To suppress a restart, even if one is required, add the following:</p> <pre>REBOOT=ReallySuppress</pre>

Logging Options

Logging options can be modified by changing the l*v variable.

l*v "%TEMP%\DLOAgentInstall.log"

Turns on verbose logging and creates a log file at the specified location.

For additional Windows Installer logging options, see <http://support.microsoft.com/kb/314852/EN-US/>.

- 3 Save and close the Setup.ini file.

Preparing for a manual push-deployment of the Desktop Agent

Before you try a manual push deployment of the Desktop Agent, you should perform the following steps. These steps are not necessary when you use the Backup Exec push installer.

Table Q-5 How to prepare for a manual push-deployment of the Desktop Agent

Step	Action	Notes
Step 1	Locate the following files: <ul style="list-style-type: none">■ *.mst■ *.cab■ DLOBuildInfo.ini■ *.msi	The files should be located in the \\<servername>\DLOAgent directory.

Table Q-5 How to prepare for a manual push-deployment of the Desktop Agent
(continued)

Step	Action	Notes
Step 2	<p>Run the msixec command using the value in setup.ini from the cmdline key as a base:</p> <pre>/qf DEFAULTMEDIASERVER="<From setup.INI File>" DLOBINSTANCENAME="<FromSetup.INI File>" LAUNCHCLIENT="1" REQUIREDMDACVERSION="2.8" TRANSFORMS="1033.mst" /!v "%TEMP%\DLOAgentInstall.log"</pre>	<p>For a silent installation, replace /qf with /qn. To install without user interaction, but with a display of the installation progress, replace /qf with /qb.</p> <p>If MDAC 2.7 is used, you must replace REQUIREDMDACVERSION="2.8" with REQUIREDMDACVERSION="2.7". No other values are valid. The installation fails if the MDAC version on target system is less than the REQUIREDMDACVERSION value.</p> <p>The specification of the TRANSFORMS property is required. The property affects the installer user interface and the start menu shortcuts. The DLO Agent is installed with support for all eight languages, regardless of which transform is chosen.</p> <p>See “Values for the TRANSFORM property of the msixec command” on page 1615.</p> <p>MSI 3.1 is required on the target systems. The MSI 3.1 installer is included in the following directory:</p> <pre>\\<servername>\DLOAgent\ WindowsInstaller-KB893803-v2-x86.exe</pre>

Values for the TRANSFORM property of the msixec command

TRANSFORMS should be set to one of the .mst files, according to the language that the desktop user uses.

See [“Preparing for a manual push-deployment of the Desktop Agent”](#) on page 1614.

Table Q-6 Values for the TRANSFORM property of the msixec command

Value	Language
1031.mst	German

Table Q-6 Values for the TRANSFORM property of the msiexec command
(continued)

Value	Language
1033.mst	English
1034.mst	Spanish
1036.mst	French
1040.mst	Italian
1041.mst	Japanese
1042.mst	Korean
2052.mst	Chinese (PRC) (Simplified)
1028.mst	Chinese (Traditional)

About setting a recovery password

When the DLO Administration Console opens for the first time, the Recovery Password Wizard opens. You must set a recovery password to enable DLO to run. If you upgraded from a previous version of DLO and previously set a recovery password, DLO uses the existing password.

The recovery password enables you to retrieve encrypted data that would otherwise be lost if the DLO database is damaged or corrupted.

After this recovery password is set, you can change it only by using the DLO command line interface tools.

See [“Checking data integrity”](#) on page 1616.

See [“About the -SetRecoveryPwd Command”](#) on page 1752.

See [“About the -EmergencyRestore Command”](#) on page 1752.

Checking data integrity

The Data Integrity Scanner simplifies the process of scanning network user data from previous DLO backups to detect unrestorable backup data. It verifies that all data is encrypted using the most recent user key. It also ensures that all data has the correct recovery key for emergency restoration.

When Desktop Agents are upgraded, they automatically perform a data integrity check. When the Administration Console is opened, it identifies Desktop Agents

that have not been checked for integrity. If any are found, you are prompted to scan them.

To check data integrity

- 1 On the **Tools** menu, click **Wizards > Data Integrity Scanner**.
- 2 Click **Next**.
- 3 To set options for quarantining data and scanning computers, click **Advanced Options**.
See [“Data Integrity Scanner Options”](#) on page 1617.
- 4 Click **Start**.
- 5 Review the scan results.
- 6 Click **Next**.
- 7 Click **Finish**.
- 8 If the scan identified data that was encrypted with outdated keys but you did not choose to quarantine the data, repeat this procedure and set the option to quarantine data that is encrypted with outdated keys.

Data Integrity Scanner Options

The **Data Integrity Scanner** simplifies the process of scanning network user data from previous DLO backups to detect unrestoreable backup data.

See [“Checking data integrity”](#) on page 1616.

Table Q-7 Data Integrity Scanner Options

Item	Description
Permanently remove previously quarantined data	Deletes all previously quarantined data.
Quarantine data encrypted with outdated keys	Quarantines all files with outdated keys. If this option is not checked, data is scanned without being quarantined. After data is quarantined, the Desktop Agent backs up a new version of the file with the correct encryption key.
Include computers that have already been validated	Forces all data to be rescanned, even if it has previously been validated.
Verbose output	Lets you receive detailed information from the scan.

Changing DLO service credentials

When DLO is installed, you must specify account credentials to be used to run the DLO Administration Service. This account is used to create Storage Locations and network user data folders. The account must have rights to create shares on any computers where backup data is to be stored. Using a Domain Administrator account is recommended. To create Storage Locations in another Domain, there must be appropriate trust relationships in effect.

To change DLO service credentials

- 1 On the **Tools** menu, click **Manage Service Credentials**.
- 2 Click **Change DLO Service Account Information**.
- 3 Enter account credentials.

See [“Service Account Information options”](#) on page 1618.

Service Account Information options

You can change the account credentials that are used to run the DLO Administration Service.

See [“Changing DLO service credentials”](#) on page 1618.

Table Q-8 Service Account Information options

Item	Description
Change DLO service account information	Lets you change the DLO service account information.
User name	Indicates the user name for the account to be used.
Domain name	Indicates the domain for this account.
Password	Indicates the password for this account.
Confirm password	Confirms the password.

About administrator accounts in DLO

Any user who has full administrator rights on the media server can manage the DLO Administration Console. The user’s account must be a domain account.

In addition, the account must have rights to do the following on any remote server used for Storage Locations or network user data folders:

- Create network shares
- Manage permissions of network shares and directories

A domain administrator account usually has the required rights.

When you search for files to restore or view history logs, the DLO Administration Console uses the credentials of the currently logged in user to access the resources. If the user does not have the correct permissions to access a resource, DLO prompts the user to enter credentials. If credentials are input, they are used to access the folder, but are not saved.

See [“About automated permissions management in DLO”](#) on page 1622.

See [“About limited restore in DLO”](#) on page 1622.

See [“Adding an administrator account”](#) on page 1620.

See [“Editing an administrator account”](#) on page 1621.

See [“Removing an administrator account”](#) on page 1621.

See [“Using a list of individual accounts to manage DLO permissions”](#) on page 1623.

See [“Using domain groups to manage DLO permissions”](#) on page 1623.

Administrator Account Management options

You can add a new administrator, remove an existing administrator, or change the settings for an existing administrator.

Table Q-9 Administrator Account Management options

Item	Description
User name	Shows the name of the user who has administrative rights.
Description	Shows the description of the user who has administrative rights.
Restore Rights	Shows whether the user has full restore rights or limited restore rights.
Add	Lets you add a new administrator.
Remove	Lets you remove an existing administrator.
Edit	Lets you change the settings for an existing administrator.

Table Q-9 Administrator Account Management options *(continued)*

Item	Description
Permissions	Lets you use domain accounts to manage administrators.

Adding an administrator account

Any user who has full administrator rights on the media server can manage the DLO Administration Console. The user’s account must be a domain account.

To add an administrator account

- 1 On **Network** menu, click **Administrator Accounts**.
- 2 Click **Add**.
- 3 Enter the user name of the user to which you want to give administrative rights.
- 4 Type a description and any applicable notes.
- 5 If you want to provide this DLO administrator with full restore privileges, including the ability to restore desktop user data to an alternate location, check **Grant administrator full restore privileges..**

Add Administrator Account options

When you add an administrator, the user name is required. All other fields are optional.

See “[Adding an administrator account](#)” on page 1620.

Table Q-10 Add Administrator Account options

Item	Description
User Name	Indicates the name of the user you want to give administrative rights to. Use the format DomainName\UserName
Description	Shows a description for this administrator account.
Notes	Provides any relevant notes regarding the administrator account.

Table Q-10 Add Administrator Account options (*continued*)

Item	Description
Grant administrator full restore privileges	<p>Provides this DLO administrator with full restore privileges, including the ability to restore desktop user data to an alternate location.</p> <p>If you allow someone other than the desktop user who owns the data to restore files to an alternate location, you can compromise data security.</p>

Editing an administrator account

You can change the description, add notes, or change the restore privileges for an administrator account.

To edit an administrator account

- 1 On **Network** menu, click **Administrator Accounts**.
- 2 Select the account that you want to edit.
- 3 Click **Edit**.
- 4 Update the description and the notes as needed.
- 5 Do one of the following:
 - If you want to provide this DLO administrator with full restore privileges, including the ability to restore desktop user data to an alternate location, check **Grant administrator full restore privileges**.
 - If you want to provide this DLO administrator with limited restore privileges, uncheck **Grant administrator full restore privileges**. Limited restore privileges do not include the ability to restore a desktop user's files to an alternate location.
- 6 Click **OK**.

Removing an administrator account

Follow these steps to remove an administrator account from DLO.

To remove an administrator account

- 1 On **Network** menu, click **Administrator Accounts**.
- 2 Select the account that you want to remove.

- 3 Click **Remove**.
- 4 Click **Yes** to confirm that you want to remove the administrator.

About automated permissions management in DLO

DLO can automatically manage permissions for accessing network user data folders. An administrator on the media server can create and configure DLO administrator accounts for users. You can use DLO Administrator accounts to avoid having to add users to the administrators group on the media server.

DLO administrator accounts can be managed in the following ways:

- Grant administrative access to individual users
This option is the default configuration for DLO account management. If you use a list of individuals, you can specify which individuals have full restore rights, and which have limited restore rights.
See [“Using a list of individual accounts to manage DLO permissions”](#) on page 1623.
- Use domain groups to manage DLO administrators
If you specify domain groups, you can grant full restore rights to one group and grant limited restore rights to another group. The domain groups must already exist or must be created by a domain administrator. For DLO, we recommend using the groups DLOFullAdmin and DLOLimitedAdmin. The full administrator group is used to grant administrators read access to user’s data. The limited administrator group only supplies list access, thus protecting the user’s data from unauthorized access.
When accessing a network user data folder, the DLO console automatically checks the folder to ensure it can read the files and data within. If the console is unable to access the folder, DLO uses the specified domain administrator group to set permissions on the files and folders it needs to access. By making these files and folders a member of the specified DLO administrator group, all DLO administrators are automatically granted permissions to access the necessary resources.
See [“Using domain groups to manage DLO permissions”](#) on page 1623.

About limited restore in DLO

The purpose of the limited restore feature is to prevent restoration of data to an alternate location by unauthorized users. By default, DLO administrators cannot restore a desktop user’s files to an alternate location, providing an additional level of data security. A DLO administrator can be granted full restore privileges, which allows the administrator to restore data to an alternate location. When a DLO

administrator has limited restoration rights, there may be other administrative functions that they are not able to perform.

Using a list of individual accounts to manage DLO permissions

An administrator on the media server can create and configure DLO administrator accounts for individual users. Accounts can be individually configured to specify full or limited restore rights.

See [“About administrator accounts in DLO”](#) on page 1618.

Alternatively, DLO can be configured to use domain groups for permissions management.

See [“Using domain groups to manage DLO permissions”](#) on page 1623.

To configure DLO to use a list of individual accounts for permissions management

- 1 On **Network** menu, click **Administrator Accounts**.
- 2 Click **Permissions**.
- 3 Uncheck the **Use domain groups to manage access to network user data folders** check box.

When this check box is checked, domain groups are listed on the **Administrator Account Management** dialog box. When this check box is unchecked, individual user accounts are listed. If you change from one type of account management to another, the previous settings are retained for future use. For example, if you have a list of individual DLO administrators and then change your configuration to use domain groups, the list of individual accounts is saved. The list of individual accounts is used again if you uncheck the check box.

- 4 Click **OK**.
- 5 Click **OK** twice.

Using domain groups to manage DLO permissions

An administrator on the media server can create and configure DLO administrator accounts for users using the DLO Administrator Account Management dialog. One method of managing DLO administrative access is to use domain groups to specify who has rights to administer DLO. Two groups can be specified. The first group is granted full restore privileges. The second group has limited restore privileges.

See [“About administrator accounts in DLO”](#) on page 1618.

Alternatively, DLO can be configured to use a list of accounts for permissions management.

See [“Using a list of individual accounts to manage DLO permissions”](#) on page 1623.

To configure DLO to use domain groups for permissions management

- 1 On the **Network** menu, click **Administrator Accounts**.
- 2 Click **Permissions**.
- 3 Check the **Use domain groups to manage access to network user data folders** check box.

When this check box is checked, domain groups are listed on the Administrator Account Management dialog box. When this check box is unchecked, individual user accounts are listed. If you change from one type of account management to another, the previous settings are retained for future use. For example, if you have a list of individual DLO administrators and then change your configuration to use domain groups, the list of individual accounts is saved. The list of individual accounts is used again if you uncheck the check box.

- 4 Select the appropriate options as follows:

For DLO administrators with full restore privileges, use the domain group

To grant full restore privileges to DLO administrators in a specified domain group, enter or browse to a fully qualified domain group.

Example: Enterprise\DLOFullAdmins

Full restore privileges include the ability to restore a desktop user's files to an alternate location.

For DLO administrators with limited restore privileges, use the domain group

To grant limited restore privileges to DLO administrators in a specified domain group, enter or browse to a fully qualified domain group.

Example: Enterprise\DLOLimitedAdmins

Limited restore privileges do not include the ability to restore a desktop user's files to an alternate location.

- 5 Click **OK** twice.

Permissions options

You can use domain groups to manage DLO permissions.

See “Using domain groups to manage DLO permissions” on page 1623.

Table Q-11 Permissions options

Item	Description
Use domain groups to manage access to network users data folders	Lets you use domain groups to specify who has rights to administer DLO.
For DLO administrators with full restore privileges, use the domain group	<p>Lets you enter or browse to a fully qualified domain group for the DLO administrators with full restore privileges..</p> <p>Example: Enterprise\DLOFullAdmins</p> <p>Full restore privileges include the ability to restore a desktop user’s files to an alternate location.</p>
For DLO administrators with limited restore privileges, use the domain group	<p>Lets you enter or browse to a fully qualified domain group for the DLO administrators with limited restore privileges..</p> <p>Example: Enterprise\DLOLimitedAdmins</p> <p>Limited restore privileges do not include the ability to restore a desktop user’s files to an alternate location.</p>

About default DLO settings

When you start DLO for the first time, defaults are already configured. You can adjust the defaults to meet the needs of your environment. Default settings are available for profiles, backup selections, and Global Settings.

You can change default settings for profiles, backup selections, and Global Settings.

Note: Changes to Global Settings take place immediately and apply globally to all Desktop Agents. Changes to the default profile and backup selection settings apply only to new profiles and backup selections and do not affect those that already exist.

Changing default DLO profile settings

The default DLO profile settings can be modified.

To change default profile settings

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the task pane, under Tool Tasks, click **Options**.

- 3 In the properties pane, under New Profile Defaults, select any of the following options:
 - **General**
See [“General options for a profile”](#) on page 1642.
 - **User Settings**
See [“User Settings options for a profile”](#) on page 1647.
 - **Schedule**
See [“Schedule options for a profile”](#) on page 1652.
 - **Options**
See [“Options for a profile”](#) on page 1654.
- 4 Change the options as needed.

Changing default DLO backup selection settings

The default DLO backup selection settings can be modified.

To change default backup selection settings

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the task pane, under Tool Tasks, click **Options**.
- 3 In the Properties pane, under New Backup Selection Defaults, click **Revisions**.
- 4 Set backup selection revision options.
See [“Revision Control options for DLO backup selections”](#) on page 1664.
- 5 In the Properties pane, under New Backup Selection Defaults, click **Options**.
- 6 Set the backup selection options.
See [“Options for a DLO backup selection”](#) on page 1666.

Changing default DLO global settings

The default DLO global settings can be modified.

Note: These settings apply immediately to all Desktop Agents.

To change default global settings

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the task pane, under Tool Tasks, click **Options**.

- 3 In the **Properties** pane, under **Global Settings**, click **Options**.
- 4 Set global options.
See “[Global Settings options](#)” on page 1627.
- 5 In the **Properties** pane, under **Global Settings**, click **Desktop Agent Intervals**.
- 6 Set the Desktop Agent Interval defaults.
See “[Desktop Agent Interval options](#)” on page 1628.
- 7 In the **Properties** pane, under **Global Settings**, click **User Activity Settings**.
- 8 Set the User Activity Settings.
See “[User Activity Settings options](#)” on page 1630.
- 9 In the **Properties** pane, under **Global Settings**, click **LiveUpdate**.
- 10 Set the LiveUpdate defaults.
See “[LiveUpdate options](#)” on page 1630.

Global Settings options

You can set the default options to limit backup activity, determine the output method for reports, and determine thresholds for the Desktop Agent.

Table Q-12 Global Settings options

Item	Description
All Desktop Agents	Prevents all Desktop Agents from backing up data.
Incremental backups of Outlook PST files	Prevents all users from performing incremental backups of Microsoft Outlook PST files. See “ About using DLO to back up Outlook PST files incrementally ” on page 1769.
Incremental backups of Lotus Notes email files	Prevents all users from performing incremental backups of Lotus Notes files. See “ Configuring the Desktop Agent for incremental backup of Lotus Notes files ” on page 1771.
Generate reports in PDF format	Creates the reports in pdf format if the Adobe Acrobat Reader is installed. If the reader is not installed, the reports appear in HTML format.
Generate reports in HTML format	Creates the reports in HTML format.

Table Q-12 Global Settings options (*continued*)

Item	Description
Time to delay Desktop Agent startup after user logs in	Indicates the number of seconds to delay the start of the Desktop Agent after the user logs on . The Desktop Agent start is only delayed if this option is checked and the Desktop Agent is started from the startup menu.
Desktop Agent low disk error threshold	Indicates the percentage of available disk space below which the Desktop Agent stops writing to the desktop user data folder.
Desktop Agent low disk warning threshold	Indicates the percentage of available disk space at which the Desktop Agent issues a low disk warning.
Desktop Agent low quota warning threshold	Indicates the percentage of the desktop user data folder storage limit below which a warning is generated. For example, if the desktop user data folder is limited to 30 MB and the low quota warning threshold is set at 10%, a warning is generated when less than 3MB are available.

Desktop Agent Interval options

You can set the default options that determine how long the Desktop Agent takes to perform certain activities.

Table Q-13 Desktop Agent Interval options

Item	Description
How long to wait before retrying the backup of a previously busy file	Indicates the number of minutes DLO waits before it retries the backup of previously busy file. If the wait time is reduced, Desktop Agent computers spend more CPU time and disk I/O trying to back up files if they are busy. If the time is set higher, files are backed up less frequently. The recommended default is 5 minutes.
How long to wait before retrying the backup of a previously failed file	Indicates the number of minutes to wait before retrying the backup of a file that previously failed to back up. If the wait time is reduced, computers spend more CPU time and disk I/O trying to back up the files that previously failed to back up. If the time is increased, files are backed up less frequently. The recommended default is 60 minutes.

Table Q-13 Desktop Agent Interval options (*continued*)

Item	Description
How long to retain backups of files that have been removed from backup selections	<p>Indicates the number of days to retain backups of the files that have been removed from backup selections.</p> <p>A longer retention time causes the files to be left on the server for a longer time. A shorter retention time provides more space in the backup folders. However, it also reduces the time during which users can restore the files that were removed from the backup selections. The recommended default is 30 days.</p>
Minimum time between history updates	<p>Indicates the number of minutes to wait between history updates.</p> <p>If there is a lot of activity, a reduced time between updates causes the computers to spend more CPU time and disk I/O to update history. A higher wait time reduces the frequency of history updates. The recommended default is 15 minutes.</p>
Minimum time between postings of the same alert	<p>Indicates the number of hours to wait between postings of the same alert.</p> <p>When there is a recurring alert, it shows up only once during the specified time interval. If the time is set too low, the alert log can fill up with multiple postings of the same alert. The recommended default is 24 hours.</p>
Minimum time between closing a job log and starting a new one	<p>Indicates the number of minutes to wait between closing a job log and starting a new one.</p> <p>When the time between job logs is reduced, more job logs appear. The recommended default is 30 minutes.</p>
Minimum time between maintenance cycles	<p>Indicates the number of minutes to wait between maintenance cycles.</p> <p>A lower time between maintenance cycles means more CPU time and disk I/O is spent conducting maintenance. Maintenance cycles remove obsolete files and folders. The recommended default is 1440 minutes, which is 24 hours.</p>
Minimum time between checking for changes to Lotus Notes email files	<p>Indicates the number of seconds between checks for changes to Lotus Notes files.</p> <p>A lower time results in more CPU time and disk I/O is used to determine if Lotus Notes files have changed. The recommended default is 30 seconds.</p>

Table Q-13 Desktop Agent Interval options *(continued)*

Item	Description
Time during which Desktop Agents randomly respond to restart requests	<p>Indicates the number of minutes during which the Desktop Agents will randomly respond to restart requests.</p> <p>When a large number of Desktop Agents are restarted, the Desktop Agents are restarted randomly over a specified period of time. This method prevents the potential for overloading DLO by starting a large number of Desktop Agents at the same time.</p> <p>The recommended default is 30 minutes.</p>

User Activity Settings options

You can set the default options that determine how DLO handles user activity.

Table Q-14 User Activity Settings options

Item	Description
Enable user activity restrictions	<p>Determines how DLO performs tasks when users interact with their desktop computers. User activity is based on typing and mouse movement.</p>
Limit network bandwidth usage to	<p>Indicates the maximum network bandwidth that DLO can use when the user interacts with the desktop computer.</p>
Restrictions will be removed when there has been no user activity for x seconds	<p>Indicates the number of seconds of user inactivity after which DLO will no longer restrict jobs.</p>
Maximum scanner items per second	<p>Indicates the maximum number of items that you want to be processed per second during a file system scan.</p> <p>File system scans occur on the following occasions:</p> <ul style="list-style-type: none"> ■ During the first backup of a desktop computer, ■ After an abnormal system shutdown ■ When the change journal is truncated. <p>This setting reduces the impact of the scan on the desktop computer while the user is active.</p>

LiveUpdate options

You can set the default options that determine how DLO interacts with Symantec LiveUpdate.

Table Q-15 LiveUpdate options

Item	Description
Enable Desktop Agent scheduled automatic updates	Turns on scheduled automatic updates.
When checking for updates	Indicates how you want DLO to check for updates. The following options are available: <ul style="list-style-type: none"> ■ Automatically download and install all available Desktop Agent updates ■ Only notify Desktop Agents of available updates (updates are not installed or downloaded)
Frequency	Indicates how often you want DLO to check for updates. The following options are available: <ul style="list-style-type: none"> ■ Once ■ Daily ■ Weekly ■ Monthly
Interval	Indicates the time to check for updates. The specific options that are available vary with the frequency selected.

Configuring DLO to use a specific port for database access

You might want to configure DLO to use a specific port for database access. For example, if a fixed port is already used for the SQL Server, you may need to configure DLO to use the same port to access the DLO database.

To configure DLO for alternate database access through a specific port

- 1 Select a unique port number for the DLO database and then use `svrnetcn.exe` to set the new port number.
- 2 On the computers that run the DLO Administration Console from outside the firewall, create the following registry key as a DWORD value if it does not exist and set the `DBUseTCP` flag to 1:

```
HKLM\SOFTWARE\Symantec\DLO\3.0\AdminConsole\DBUseTCP
```

- 3 On the computers that run the Desktop Agent from outside the firewall, create the following registry key as a DWORD value if it does not exist and set the DBUseTCP flag to 1:

HKCU\Software\Symantec\DLO\3.0\Client\DBUseTCP or
HKLM\SOFTWARE\Symantec\DLO\3.0\Client\DBUseTCP
- 4 Set the DBTcpPort on the computers that you modified in steps 2 and 3 to the port number you set in step 1.
- 5 Restart the modified computers.

About using Backup Exec Retrieve with DLO

DLO can be configured to integrate with Backup Exec Retrieve, a feature of the Symantec Continuous Protection Server (CPS). Backup Exec Retrieve allows desktop users to view, search, and restore files directly to their workstation using a Web browser.

Backup Exec Retrieve is optimized for Microsoft Internet Explorer (6.0 or later). Backup Exec Retrieve also works with other Web browsers, although screen layouts may vary.

Backup Exec Retrieve is protected by Windows-level security. Desktop users are prompted for Windows domain logon information. This information is used to restrict the files that you can view and retrieve.

If you protect a server, users of Backup Exec Retrieve can only access files and folders originally written to a share. That is, unless a folder is shared from a file server, you cannot see or retrieve the files.

To enable DLO integration with CPS, the CPS Continuous Management Service (CMS) must be installed on the same computer as the Backup Exec Media Server. Additionally, the administrator must install a Continuous Protection Agent (CPA) on each DLO file server. The CPA installation can be performed by manually running setup on each computer, or using the push-install mechanism in the CPS Administration Console. See the Symantec Backup Exec Continuous Protection Server Administrator's Guide for more information.

Desktop users can search and restore files that were backed up with DLO.

See [“About using Backup Exec Retrieve to restore files”](#) on page 1790.

About updating DLO

Symantec provides updates in the following ways:

- Periodic product updates are delivered by Symantec LiveUpdate. Some updates are not installed automatically on Desktop Agents. However, security updates are installed automatically on Desktop Agents.
- Significant product upgrades are delivered on the Symantec Web site or on installation media.

Updating the DLO Administration Console

The default installation directory for Backup Exec DLO is:

C:\Program Files\Symantec\Backup Exec\DLO

If DLO is upgraded from a previous version that was installed in a different location, the installation is moved to this new location.

To update the DLO Administration Console

- 1 Install the DLO Administration Console.
See [“About installing the Backup Exec Desktop and Laptop Option”](#) on page 1610.
- 2 Start the DLO Administration Console, and then set a recovery password.
See [“About setting a recovery password”](#) on page 1616.
- 3 If you are updating from DLO version 9.1, run the Data Integrity Scanner.
See [“Checking data integrity”](#) on page 1616.

Updating the Desktop Agent

When the media server is updated, either through a full install, hotfix or Service Pack release, the Desktop Agents need to be updated in one of the following ways:

Table Q-16 How to update the Desktop Agent

Method	Description
From the Desktop Agent	Run the setup.exe file from the computer on which the Desktop Agent is installed. The setup.exe file is located in the following directory: \\<media server>\DLOAgent\update_13.0\Setup.exe.
From the Backup Exec Administration Console	Use the Backup Exec push-install feature.

Table Q-16 How to update the Desktop Agent (*continued*)

Method	Description
From the DLO Administration Console	Use the publish command in the DLO Command Line Interface.

See [“About the -Update Command”](#) on page 1749.

See [“About the DLO command syntax”](#) on page 1740.

To update Desktop Agents from the DLO Administration Console using the Command Line Interface

- 1 Update the media server as directed in the update documentation.
- 2 From the command line on the media server, change to the DLO installation directory.

The default installation directory is:

C:\Program Files\symantec\Backup Exec\DLO

- 3 Run `DLOCommandu.exe` with the update option to add the configuration file and make note of the ID number that is returned when this command is run:

```
DLOCommandu -update -add  
-f..\agents\dlo\update_13.0\DLOAgentUpdate_BE.ini
```

If the configuration file has been moved or renamed, you must specify the full path and file name in the command.

Sample output:

```
ID=3  
Name= 13.0 Update  
Description=Updates Backup Exec DLO Desktop Agent to 13.0  
Version=3.1 Build=3.XX.XX  
srcPath=\\MediaServerName\DLOAgent\update_13.0  
cmdPath=%DOWNLOADDIR%  
cmdName=AutomatedAgentUpgrade.exe  
cmdArgs=-s
```

- 4 Run `DLOCommandu.exe` with the publish command to make the update available to Desktop Agent users.

```
DLOCommandu -update -publish -UI y -U UserName  
DLOCommandu -update -publish -UI y -P ProfileName
```

Where *y* is the ID number returned when the 'add' command was run in step 3. Using an asterisk in place of *UserName* or *ProfileName* publishes the update to all users.

When this command is executed, it returns a list of all users that are targeted for update. Users will be updated the next time the Desktop Agent application is started.

Performing a silent upgrade of the Desktop Agent

Desktop Agents can be upgraded silently. During a silent upgrade, the user is not prompted to download and start the upgrade. However, the user is prompted to confirm the upgrade.

To run the Desktop Agent upgrade silently

- 1 From the Desktop Agent upgrade folder, open the `DLOAgentUpdate_BE.ini` file for editing.
- 2 Set `PromptUser=0`.

- 3 Save and Close the file.
- 4 Run the upgrade.

About upgrading DLO to Windows Vista

DLO includes the following changes to support Windows Vista:

- The Documents folder no longer includes subfolders for Music, Pictures, or Videos. Therefore, if you select the Documents folder and the Include subfolders option, the data in the Music, Pictures, and Videos folders is not backed up. You must select each folder that you want to back up.
- All backed-up data is stored in the AppData folder.
- File revisions and delta file copies begin anew. Previous data remains in the user data folders according to the settings in the deleted file retention policy.
- You can synchronize data between two computers that run Windows Vista. You cannot synchronize data between a computer that runs Windows Vista and a computer that runs an earlier operating system, such as Windows XP.

Upgrading From NetBackup Professional to DLO

The NetBackup Professional (NBUP) to Desktop Agent upgrade is only available for NBUP customers running version 3.51.20 or later. If you are not running 3.51.20, consider upgrading your NBUP server and clients before upgrading to the Desktop Agent.

This mechanism installs the Desktop Agent onto desktop computers that are currently running the NBUP client. You can remove the NBUP client when installing the Desktop Agent or leave the NBUP client installed and run both applications concurrently. The two options appear as separate upgrades in the NBUP Console, so you can remove NBUP from some profiles and continue to run NBUP for other profiles.

The upgrade from NBUP to DLO requires the following additional components that are distributed with the Desktop Agent install set:

- A DLO Client (Remove NBUP).VPK file. It contains instructions and an executable to upgrade the system to DLO and remove NBUP at the same time.
- A DLO Client (Leave NBUP).VPK file. It contains instructions and an executable to upgrade the system to DLO and leave NBUP installed but increment the version number so that it appears NBUP was upgraded.

DLO supports versions 2.7 and 2.8 of Microsoft Data Access Components (MDAC). However, MDAC 2.8 is the default. If a non-administrator logs on first after the MDAC 2.8 installation, the process generates a number of errors. To avoid these

errors when MDAC 2.7 is already installed, you can force the use of MDAC 2.7. Modify the package.ini file in the Upgrades folder. In the package.ini file, add REQUIREDMDACVERSION="2.7" to the DefaultRuleXML line.

Example:

```
DefaultRuleXML=<MSIPropertiesAppend>TRANSFORMS="%%%LANG_FILE%%%"  
REBOOT=ReallySuppress LAUNCHCLIENT="0" REQUIREDMDACVERSION="2.7"  
</MSIPropertiesAppend><MSIPropertiesFile>setup.ini</MSIPropertiesFile>
```

To upgrade from NetBackup Professional to DLO

- 1 Contact Technical Support to receive the NBUP to Desktop Agent upgrade. The required files are DLOAgent_LeaveNBP.vpk and DLOAgent_RemoveNBP.vpk.
- 2 From the NBUP server, or any computer with the NBUP console installed, launch the appropriate file; DLO Client (leave NBUP).vpk or DLO Client (remove NBUP).vpk. The file uploads the upgrade package to the NBUP server. Repeat this process for the other vpk file to make both the leave and remove NBUP options available for selection in various profiles.
- 3 Create a folder entitled DLOAgent in C:\Program Files\Veritas NetBackup Professional\Upgrades, or in the appropriate location if you installed NBUP in a location other than the default.
- 4 Copy the entire contents of the DLOAgent share on the media server into the DLOAgent folder on the NBUP server.
- 5 Launch the NBUP Console.
- 6 Open the profile properties.
- 7 On the Upgrades tab, select the appropriate upgrade (leave NetBackup Professional or remove NetBackup Professional) and enable it by checking the **Enable this upgrade** check box. Select the other options you want for this upgrade.
- 8 Repeat step 1 through step 7 for each NBUP Profile you want to upgrade to DLO.
- 9 Follow the standard procedure for upgrading NBUP ("Check for upgrade now" in the console or refresh the client). See the NetBackup Professional Administrator's Guide for additional information.

If the Desktop Agent installation is successful, the NBUP version number in the NBUP administration console changes to one of the following numbers:

- 9.1.0.0 for the computers that still have NetBackup Professional installed
- 0.0.0.1 for the computers on which NetBackup Professional was removed

Starting the DLO Administration Console from Backup Exec

The DLO Administration Console is launched from the Job Setup view in Backup Exec. From the DLO Administration Console, you can configure DLO and manage desktop backup and restore operations.

To start the DLO Administration Console from Backup Exec

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under Backup Strategy Tasks, click **Configure desktop and laptop backups**.

About the DLO Overview view

The DLO Overview view provides two options; the Getting Started view and a System Summary.

The Getting Started view provides convenient links to help you set up and manage DLO.

From this view, you can perform the following tasks or access the help that is associated with them:

- Deploy the Desktop Agent
- Set preferences and default settings
- Manage alerts and notifications
- Run reports
- Create a profile
- Create a Storage Location
- Create an Automated User Assignment
- Add users

The following information is available in the System Summary overview:

Table Q-17 System Summary options

Item	Description
Last Backup Result	<p>Summarizes the completion status of the last operation that was performed on each computer that is protected by DLO.</p> <p>Totals are provided for the number of computers that completed the last job in the following categories:</p> <ul style="list-style-type: none"> ■ With Errors - The last operation was completed, but errors were generated. ■ With Warnings - The last operation was completed, but warnings were generated. ■ Canceled - The job was canceled or refreshed by the user during the job. ■ Successful - The job was successfully completed without warnings or errors, and it was not canceled or refreshed by the user during the job. <p>Errors take precedence over warnings, so if there are both errors and warnings, the last backup result is With Errors.</p>
Pending Jobs	Lists restore jobs that were requested by the DLO administrator and that have not yet run.
Active Alerts	Lists the alerts that DLO administrators have not cleared by the and the alert grooming process have not yet removed. Alert grooming is managed from the Backup Exec Administration Console.
Server Status	<p>Lists the status of each DLO server.</p> <p>Server status can be one of the following:</p> <ul style="list-style-type: none"> ■ Running ■ Stopped
Server Load	Lists the number of desktops that are being protected by DLO and the total number of installed Desktop Agent users. These numbers may not be the same if some users protect multiple computers with DLO. Both online and offline users are counted.

Connecting to DLO on a different Backup Exec Media Server

To connect to DLO on a different Backup Exec media server, the user account needs to have full administrator rights to the media server and it must also be a domain account.

See [“About administrator accounts in DLO”](#) on page 1618.

To connect to DLO on a Backup Exec media server

- 1 On the DLO Network menu, click **Connect to Media Server**.
- 2 Select the appropriate options.
See [“Connect to Media Server options for DLO”](#) on page 1640.
- 3 Click **OK**.

Connect to Media Server options for DLO

On the **Connect to Media Server** dialog box, you enter the credentials that are required to connect to DLO on a different Backup Exec media server.

See [“Connecting to DLO on a different Backup Exec Media Server”](#) on page 1639.

Table Q-18 **Connect to Media Server** options for DLO

Item	Description
Server	Indicates the name of the media server you want to connect to, or select a media server from the drop-down menu.
User name	Indicates the user name for an account with administrator access to the media server.
Password	Indicates the password for this account.
Domain	Indicates the domain for this account.

How to configure DLO

Before desktop users can back up data, you must set up the following options in the following order:

- Create a profile. A profile determines what files are backed up, when the files are backed up, and the level of interaction the desktop user has with the Desktop Agent.
See [“About DLO profiles”](#) on page 1641.
- Determine where you want to store user data on the network. DLO requires an individual user data folder on the network for each desktop user.
See [“About DLO Storage Locations”](#) on page 1676.
See [“About managing Desktop Agent users”](#) on page 1696.

- Create an Automated User Assignment to automatically assign a Storage Location and profile to new users, or configure new users manually. See “[About Automated User Assignments](#)” on page 1683.

You can set up DLO by using the Desktop and Laptop Configuration Wizard, or by setting options manually. The DLO configuration wizard provides a series of wizards that help you set up DLO in the correct order.

Starting the Configuration Wizard

The DLO configuration wizard provides a series of wizards that help you set up DLO in the correct order.

To start the Configuration Wizard

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the task pane, under Getting Started, click **DLO Configuration using wizard**.
- 3 If you want the Configuration Wizard to appear each time the DLO Administration Console is started, click **Always show this wizard at startup**.

About DLO profiles

Profiles are used to customize settings for specific groups of similar users. For example, a group of highly technical users may require the ability to modify the backup selections and schedules. Less experienced users may require a fully-automated backup service.

In a profile, you can set the following items:

- Backup file and folder selections
- Desktop and network user data folder storage limits
- Backup schedules
- The desktop user’s level of interaction with the Desktop Agent
- Logging options
- Network bandwidth usage

You cannot modify settings for individual Desktop Agent users from the DLO Administration Console unless an individual user is the only user assigned to a profile. However, you can grant permission to Desktop Agent users to modify their own settings.

Creating a new DLO profile

New profiles can be created to meet the specific needs of desktop users, and to support the existing IT environment.

To create a new DLO profile

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the Settings pane, click **Profiles**.
- 3 In the task pane, under Settings Tasks, click **New Profile**.
- 4 On the General tab, enter a name for the profile, and set bandwidth settings, storage limits, and the desktop user data folder path.

See [“General options for a profile”](#) on page 1642.

- 5 On the Backup Selections tab, select the backup selections that you want to apply to users of this profile.
- 6 On the User Settings tab, configure the level of interaction that desktop users can have with the Desktop Agent.

If an individual user changes the user settings, the changes apply only to that user.

See [“User Settings options for a profile”](#) on page 1647.

- 7 On the Schedule tab, select the frequency with which you want data to be backed up..

See [“Schedule options for a profile”](#) on page 1652.

- 8 On the Options tab, select logging options and mail options.

See [“Options for a profile”](#) on page 1654.

- 9 On the Connection Policies tab, click **Add** to add a new connection policy

See [“Add/Edit Connection Policy options”](#) on page 1656.

- 10 Click **OK**.

General options for a profile

On the **General** tab, you can enter a name for the profile, and set bandwidth settings, storage limits, and the desktop user data folder path.

See [“Creating a new DLO profile”](#) on page 1642.

Table Q-19 General options for a profile

Item	Description
Profile Name	Indicates the name of the new profile that you want to create. The profile name cannot contain any of the following characters: \ "# \$ % ^ & * () = + / { } [] ' "
Description	Indicates a description for the profile.
Enable Profile	Enables or disables the profile. Profiles are enabled by default.
Limit bandwidth (KB/sec)	<p>Controls the rate at which data is sent to the network user data folder.</p> <p>Limiting the bandwidth is a means to manage the trade-off between backup speed and the impact of backups on the local computer, network, and server. The default limit is meant to be a conservative setting to minimize the impact of backups.</p> <p>However, many factors can affect the setting, such as the following:</p> <ul style="list-style-type: none"> ■ Network speed ■ Connection type ■ The amount of data that is backed up ■ The total number of computers that back up to DLO. <p>If computer performance is not impacted, but DLO data transfer is slow, a higher bandwidth setting may be more suitable. If computer performance is noticeably impacted during backups, a lower value reduces the impact of backups on computer performance. However, backups take longer to complete.</p> <p>Data transfer is only limited when data is written to the network user data folder. Data transfer is not limited when data is written to the desktop user data folder. Data transfer is not limited during the incremental backup of Outlook PST files or Lotus Notes NSF files.</p>

Table Q-19 General options for a profile (*continued*)

Item	Description
<p>Yield bandwidth to other programs</p>	<p>Enables DLO to reduce data transfer over the network when other applications on the desktop computer transfer data. DLO automatically resumes normal data transfer rates when other applications are not using this resource.</p> <p>The yield bandwidth option monitors network traffic on the desktop computer. If DLO uses more than 90% of the total current traffic, DLO is not throttled.</p> <p>DLO throttles itself to use only the otherwise unused portion of the connection when the following conditions are met:</p> <ul style="list-style-type: none"> ■ DLO traffic drops below 90% of the total network traffic on the desktop ■ Total traffic is over 60% of the maximum traffic that is seen on the connection <p>For example, if there was 70% total usage, DLO throttles itself to 30% of maximum.</p> <p>Selecting this option can improve system performance when other network-intensive applications are running at the same time. Data transfer is only limited when data is written to the network user data folder. Data transfer is not limited when data is written to the desktop user data folder.</p>
<p>Limit network user data folder to (MB)</p>	<p>Limits the disk space available on the network to store DLO backup files and type the amount of space you want to use for storage.</p>

Table Q-19 General options for a profile (*continued*)

Item	Description
Enable desktop user data folder	<p>Enables the use of the desktop user data folder. When Enable desktop user data folder is checked, the following actions occur:</p> <ul style="list-style-type: none"> ■ Files are copied to the desktop user data folder first. ■ Files are then copied to the network user data folder from the desktop user data folder. <p>These actions occur even when DLO is configured to keep zero revisions in the desktop user data folder.</p> <p>When Enable desktop user data folder is unchecked, files are copied straight to the network user data folder from the original location.</p> <p>Advantages to enabling the desktop user data folder:</p> <ul style="list-style-type: none"> ■ Offline protection is provided because revisions can be stored locally as well as on the network. ■ Because files are more quickly saved to the local computer than to the network, the time a file is held open for backup is reduced. <p>Advantages to disabling the desktop user data folder:</p> <ul style="list-style-type: none"> ■ If local revisions are not required, this option prevents backup files from being stored in the desktop user data folder. No revisions are saved in the desktop user data folder even if backup selections specify that a certain number of revisions should be stored locally. ■ Works well for desktop users with very limited disk space. ■ DLO creates empty place holders in the desktop user data folder, even if the folder is disabled or the number of revisions is set to zero. The place holders can be seen in the Desktop User data folder, but contain no data. They indicate which files and folders have been backed up and saved to the network user data folder.

Table Q-19 General options for a profile (*continued*)

Item	Description
<p>Limit desktop user data folder to</p>	<p>Limits the disk space that is available to store DLO backup files.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> ■ A percentage of the total disk space (%) Indicates the percentage of the total disk space that you want to allocate to storage of DLO backup files in the desktop user data folder. ■ A size (MB) Indicates the maximum amount of disk space, in megabytes, that you want to allocation to storage of DLO backup files in the desktop user data folder. <p>By limiting available disk space for the desktop user data folder, you can prevent overloading of the desktop's hard drive. However, backups can fail to run if the desktop user data folder space limit is set too low.</p>
<p>Override the default desktop user data folder path</p>	<p>Prevents the desktop user data folder from residing in the default location. You can type the path to the folder that you want all new Desktop Agent users that are assigned to this profile to use.</p> <p>The folder must already exist on the desktop before the new user logs on for the first time after being assigned to this profile. The Desktop Agent does not create the folder. If the folder does not exist before the user logs on for the first time, DLO uses the default folder for backups.</p> <p>The default folder is located at the following path on the computers that run Windows XP and earlier operating systems:</p> <p>\\Documents and Settings\<<user_name>\Local Settings\Application Data\Symantec</p> <p>The default folder is located at the following path on the computers that run Windows Vista:</p> <p>\\Users\<<user_name>\AppData\Symantec</p>

Backup Selection options for a profile

You can add, modify, and delete backup selections for a profile from this dialog box. When a new backup selection is created, it is available for selection in all profiles. Changes that are made to a backup selection in one profile affect all other

profiles that use the backup selection. Similarly, when a backup selection is deleted, the change affects all profiles that use the backup selection.

See [“Creating a new DLO profile”](#) on page 1642.

User Settings options for a profile

On the **User Settings** tab, you can configure the level of interaction that desktop users can have with the Desktop Agent.

See [“Creating a new DLO profile”](#) on page 1642.

Table Q-20 **User Settings** options for a profile

Item	Description
<p>Desktop Agent display settings</p>	<p>Determines the desktop user's level of interaction with the Desktop Agent.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> ■ Display the complete interface Enables desktop users to access all Desktop Agent options. ■ Display only the status Enables desktop users to view the status of backup jobs. With this option, desktop users cannot change settings for the Desktop Agent or access any options other than the status. Desktop users can right-click the system tray icon to open the status view or exit the program. ■ Display only the system tray icon Displays the Desktop Agent icon in the system tray in the lower right corner of the screen. Desktop users can right-click the system tray icon to exit the program. ■ Do not display anything Runs the Desktop Agent in the background. The desktop user cannot view the Desktop Agent.
<p>Restore data</p>	<p>Enables the users in this profile to restore their backed up files.</p> <p>See "Restoring files by using the Desktop Agent" on page 1786.</p>

Table Q-20 User Settings options for a profile (*continued*)

Item	Description
Add user-defined backup selections	<p>Enables the users in this profile to create and modify backup selections. This option does not allow users to modify backup selections that are made by the DLO administrator in the profile.</p> <p>Users can add a backup selection to back up a folder that is excluded from the profile backup selections. The only way to prevent users in a profile from backing up a specific folder is to uncheck this option.</p> <p>See “About backup selections in DLO” on page 1658.</p> <p>See “Modifying backup selections in the Desktop Agent's standard view” on page 1767.</p> <p>See “Modifying backup selections in the Desktop Agent's advanced view” on page 1768.</p>
Modify profile backup selections	<p>Enables the users in this profile to modify backup selections that were created by the DLO administrator for the profile.</p> <p>See “About backup selections in DLO” on page 1658.</p> <p>See “Modifying backup selections in the Desktop Agent's advanced view” on page 1768.</p>
Customize backup selection revision policy settings	<p>Enables the users in this profile to modify the revision policy settings. Users cannot change these settings if this option is not checked.</p>
Change backup selection encryption settings	<p>Lets the users in this profile enable or disable encryption of backup files.</p>

Table Q-20 **User Settings** options for a profile (*continued*)

Item	Description
Change backup selection compression settings	Lets the users in this profile enable or disable compression of backup files.
Customize profile logging settings	Enables the users in this profile to customize profile logging settings. See “Setting customized options on the Desktop Agent” on page 1775.
Customize profile email settings	Enables the users in this profile to customize mail settings in the profile. See “Setting customized options on the Desktop Agent” on page 1775.
Move local user data folder	Enables the users in this profile to move the local user data folder to a new location. See “Moving the desktop user data folder” on page 1777.
Change groom policy settings	Enables the users in this profile to customize grooming settings. See “Setting customized options on the Desktop Agent” on page 1775.
Synchronize files	Enables the users in this profile to synchronize data across all of their computers that run the Desktop Agent. See “About synchronizing desktop user data” on page 1778.
Customize local disk quota	Enables the users in this profile to limit the amount of disk space that can be used to store backup files in the desktop user data folder. See “Setting customized options on the Desktop Agent” on page 1775.

Table Q-20 User Settings options for a profile (*continued*)

Item	Description
Modify backup schedule	Enables the users in this profile to modify the schedule on which their files are backed up. See “Changing schedule options for a DLO backup job” on page 1773.
Customize connection policies	Enables the users in this profile to customize connected based policies. See “Customizing connection policies” on page 1777.
Cancel scheduled or manual jobs	Enables the users in this profile to cancel both scheduled jobs and manually initiated jobs. Scheduled jobs run again at the next scheduled time. Manual jobs must be restarted manually. See “About suspending or canceling a job” on page 1783.
Suspend jobs	Enables the users in this profile to suspend jobs for a specified amount of time. See “About suspending or canceling a job” on page 1783.
Disable Desktop Agent	Enables the users in this profile to disable the Desktop Agent from the tray icon. The users can also enable the Desktop Agent.
Work Offline	Enables the users in this profile to set the Desktop Agent to work offline. See “Changing your connection status” on page 1762.

Table Q-20 **User Settings** options for a profile (*continued*)

Item	Description
Save encrypted passwords used by DLO	Enables users to automatically authenticate to the media server or storage location in the event of an authentication failure. An authentication error may occur when the desktop user logs on using a local or cross-domain account. If you do not select this option, DLO prompts for a password in the event of an authentication failure.
Suppress errors and warnings	Prevents error and warning message boxes from being displayed when a user is not interacting with the Desktop Agent.
When user goes offline, automatically go back online after	Indicates the amount of time you want the Desktop Agent to wait before going online again after the Desktop Agent has placed it offline.
When user suspends a job or disables the Desktop Agent, automatically resume or enable after	Indicates the amount of time you want the Desktop Agent to wait before resuming a job or enabling the Desktop Agent after the user suspends a job or disables the Desktop Agent.

Schedule options for a profile

On the **Schedule** tab, you determine when jobs run for users with this profile.

See [“Creating a new DLO profile”](#) on page 1642.

Table Q-21 Schedule options for a profile

Item	Description
Whenever a file changes	<p>Backs up files whenever they change.</p> <p>On NTFS drives, backups occur automatically whenever a file changes. For FAT drives, you must enter a backup interval in the Back up changed files every field.</p>
According to a schedule	<p>Backs up files according to a customized schedule.</p>
Edit schedule	<p>Lets you configure the backup schedule.</p> <p>See "Backup Schedule options" on page 1654.</p>
When initiated by the user	<p>Enables desktop users to determine when to back up their files.</p>
Do nothing	<p>Lets the user proceed with a logout, restart, or shutdown even when there are files that require backup.</p> <p>If a job is already running, the user is prompted to log out, restart, or shut down when the job is complete.</p>
Prompt user to run job	<p>Prompts the user to run a backup job before proceeding with the logout, restart, or shutdown.</p> <p>If a job is already running, the user is prompted to cancel the job before continuing with the logout, restart or shutdown.</p>
Run job immediately	<p>Backs up waiting files without prompting before proceeding with a logout, restart, or shutdown.</p> <p>If a job is already running, the user is prompted to cancel the job before continuing with the logout, restart or shutdown.</p>
Run job as scheduled	<p>Proceeds with a logout, restart, or shutdown, and to back up files according to the schedule.</p> <p>If a job is already running, the user is prompted to cancel the job before continuing with the logout, restart or shutdown.</p>
Run job at next login	<p>Proceeds with a logout, restart, or shutdown without prompting, and to run a job at the next logon.</p> <p>If a job is already running, the user is prompted to cancel the job before continuing with the logout, restart or shutdown.</p>

Backup Schedule options

You can set up the days and times to run backup jobs.

Table Q-22 Backup Schedule options

Item	Description
Run on these days	Indicates the days on which you want to back up files.
Run once at	Runs a single backup on the days you selected at the time specified.
Run every	Runs the backup jobs at the specified time interval on the days you selected.
From	If you selected Run every, indicates the beginning of the time interval over which you want backups to begin.
Until	If you selected Run every, indicates the end of the time interval over which you want backups to begin. This field specifies the end of the time period within which backups can begin. If a backup is in progress at this time, it runs to completion.
Start backup jobs over a period of	Staggers start times for backup jobs. Rather than starting all backup jobs at exactly the time indicated, DLO distributes the start times over the specified interval to better distribute the demands on the server and network.

Options for a profile

On the **Options** tab, you can select logging options and mail options.

See [“Creating a new DLO profile”](#) on page 1642.

Table Q-23 Options for a profile

Item	Description
Keep log files for a minimum of (days)	Indicates the minimum number of days to keep log files. Log files are not deleted until they are at least as old as specified. Log files are not deleted until their combined size exceeds the setting for the combined size of all log files.

Table Q-23 Options for a profile (continued)

Item	Description
After minimum number of days, delete oldest log files when combined size exceeds (MB)	<p>Indicates the maximum combined size of all log files to be retained before the oldest log files are deleted.</p> <p>You may have more than the specified number of MB of log files stored if none of the log files is as old as specified in the Keep log files for a minimum of (days) setting.</p>
Log groom messages	Creates logs for grooming operations.
Log information messages for backup	Creates logs for all backup operations.
Log warning messages	Creates logs for all operations that generate warnings.
Enable message level incremental backups of Outlook PST files	<p>Enables incremental backups of Microsoft Outlook Personal Folder (PST) files. Incremental backups must be enabled to allow PST files to be backed up while they are open.</p> <p>If this option is not checked, PST files that are configured in Outlook are fully backed up each time the PST file is saved. In general, PST files are saved when Outlook is closed.</p> <p>When Outlook PST files are backed up incrementally, only one revision is maintained regardless of the number of revisions that are set in the backup selection.</p> <p>Microsoft Outlook must be your default mail application for DLO to perform incremental backups of PST files.</p> <p>Synchronized files cannot be backed up incrementally.</p> <p>See “About using DLO to back up Outlook PST files incrementally” on page 1769.</p>
Ignore PST files which have not been configured in Outlook	Excludes PST files that are not registered with the Microsoft Outlook client on a Desktop Agent computer.
Enable message level incremental backups of Lotus Notes email files	<p>Enables incremental backups of Lotus Notes email files. Additional configuration may be necessary.</p> <p>See “Configuring the Desktop Agent for incremental backup of Lotus Notes files” on page 1771.</p> <p>When Lotus Notes NSF files are backed up incrementally, only one revision is maintained regardless of the number of revisions that are set in the backup selection.</p>

Add/Edit Connection Policy options

You can limit or disable backups based on the connection type.

Table Q-24 Add/Edit Connection Policy options

Item	Description
Connection Type	<p>Indicates the type of connection for which you want to limit or disable backups.</p> <p>You can choose one of the following connection types:</p> <ul style="list-style-type: none"> ■ Dialup Limits or disables backups when using a dial-up connection . ■ IP address range Limits or disables backups for a specific IP address range. Specify whether you want the connection policy to apply to computers that are or are not in the IP address range you specify. Select IPv6 or IPv4 and enter the IP address range for the connection policy. IPv6 addresses are only supported on Windows XP and later operating systems. IPv6 addresses are not enforced for Desktop Agents running on Windows 2000. An additional connection policy using IPv4 addresses may be desired for Desktop Agents on Windows 2000 computers. ■ Active Directory Limits or disables backups using Active Directory. Select Configure to configure the Active Directory settings. See “Customizing connection policies” on page 1777.
Disable network backup	Prevents users from backing up to the network user data folder. Backups continue to the desktop user data folder.
Disable network backup for files greater than	Prevents users from backing up files larger than a specified size based on the connection type. Enter a files size in KB.
Limit network bandwidth usage to	Restricts the usage of network bandwidth to the specified value. The value must be entered in KB/sec format.
Enforce policy according to scheduled window	Causes the connection policy to apply only during the specified period of time.
Schedule	Lets you set the time during which the policy will be in effect. Schedules can be set to run weekly or for a specific date range.

Schedule options for a profile's connection policy

You can set the time when you want a connection policy to be in effect.

Table Q-25 Schedule options for a profile's connection policy

Item	Description
Occurs	Indicates the frequency of the schedule.
Starts at/on	Indicates the time and day when the schedule begins.
Ends at/on	Indicates the time and day when the schedule ends.

Copying a DLO profile

You can use an existing profile as the basis for a new profile. For example, if an existing profile contains many of the same settings that you want to use for a new profile. The copy can then be modified as required to meet the needs of a new group of desktop users.

To copy a profile

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the Settings pane, click **Profiles**.
- 3 Right-click the profile you want to copy.
- 4 Click **Copy**.
- 5 Type a name for the new profile.
- 6 Type a description of the new profile.
- 7 Click **OK**.

Modifying a DLO profile

Profiles can be modified as required to meet the changing needs of user groups.

Note: Modifications to a profile can cause users of that profile to cancel jobs, load settings, restart backup engines and scan their backup selection tree.

To modify a profile

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the Settings pane, click **Profiles**.
- 3 In the Results pane, select the profile you want to modify.
- 4 In the task pane, under General Tasks, click **Properties**.
- 5 Modify the properties as needed.
- 6 Click **OK**.

About backup selections in DLO

Backup selections specify which files and folders you want to back up on desktops. An administrator can create a backup selection that applies to all users in a profile. In this case, the backup selection is called a profile backup selection. In addition, desktop users who have sufficient rights can create and modify backup selections.

Within each backup selection you can do the following:

- Specify the path to be backed up.
- Choose to include or exclude subfolders, file types, or specific folders.
- Set the number of revisions that are retained for each file in the backup selection.
- Set the frequency with which revisions are saved.
- Set instructions on how long to retain backup files.
- Configure the backup selection to transfer only the changed portions of files.
- Compress or encrypt the files for transmission and storage.
- Specify how long to retain backup files after the source files are deleted.

Caution: Symantec strongly recommends that you consider disk space when choosing backup selections for desktops and laptops. A large number of local copies may cause the Desktop Agent user's computer to run out of disk space. For example, you may want to avoid selecting entire drives for backup or synchronization.

See [“About default backup selections in DLO”](#) on page 1659.

See [“Removing default DLO backup selections from a profile”](#) on page 1660.

See [“Adding a DLO backup selection to a profile”](#) on page 1660.

See [“General options for DLO backup selections”](#) on page 1661.

See [“Including and excluding files or folders from a DLO backup selection”](#) on page 1662.

See [“About revision control in DLO”](#) on page 1663.

See [“Setting options for a DLO backup selection”](#) on page 1666.

See [“How to use DLO macros in backup selections”](#) on page 1667.

See [“Modifying a DLO backup selection”](#) on page 1670.

See [“Deleting DLO backup selections”](#) on page 1670.

About default backup selections in DLO

DLO is configured to back up commonly used files and folders by default. You can add additional backup selections or cancel the use of default backup selections.

The following items are backed up by default:

Table Q-26 Default backup selections

Backup selection	Description
My Documents	All files in My Documents (Documents in Windows Vista)
My Favorites	Internet Explorer Favorites
Outlook PST Files	PST files in the default location
My Desktop	All files on the Desktop
Notes Files (Multi-user)	Lotus Notes data for multiple user install
Notes Archive (Multi-user)	Lotus Notes archive for multiple user install
Notes Files (Single-user)	Lotus Notes data for single user install
Notes Archive (Single-user)	Lotus Notes archive for single user install
My Music	All files in My Music (Music in Windows Vista)
My Pictures	All files in My Pictures (Pictures in Windows Vista)
My Videos	All files in My Videos (Videos in Windows Vista)

The default backup selections assume that applications use default paths. If custom paths were used during installation or modified thereafter, you must customize the backup selections to insure they work properly.

See [“Modifying a DLO backup selection”](#) on page 1670.

Removing default DLO backup selections from a profile

Default profile backup selections are appropriate for most DLO installations. In some cases, you may want to remove or replace default backup selections.

To remove default backup selections from a profile

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the Settings pane, click **Profiles**.
- 3 In the Results pane, select the profile you want to modify.
- 4 In the task pane, under General Tasks, click **Properties**.
- 5 On the Backup Selections tab, uncheck the backup selections that you do not want to use.
- 6 Click **OK**.

Adding a DLO backup selection to a profile

When a new backup selection is created for a profile, that profile backup selection is available for selection in all other profiles.

To add a backup selection to a profile

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the Settings pane, click **Profiles**.
- 3 In the Results pane, select the profile for which you want to add a backup selection.
- 4 In the task pane, under General Tasks, click **Properties**.
- 5 On the Backup Selections tab, click **Add**.

If you customize NTFS permissions or directory attributes such as compression or encryption for backed up files or folders, these settings are not backed up. You must reapply these settings after you restore the files. If you use a password for a Microsoft Outlook PST file, you must reset the password after you restore a PST file.

- 6 Read the message that appears, and then click **OK**.
- 7 Do any of the following to customize the backup selection properties:

- On the **General** tab, set general backup selection properties including the name, description, and folder to be backed up.
See [“General options for DLO backup selections”](#) on page 1661.
- On the **Include/Exclude** tab, include or exclude specific files from this backup selection.
See [“Including and excluding files or folders from a DLO backup selection”](#) on page 1662.
- On the **Revision Control** tab, set revision control for this backup selection.
See [“Revision Control options for DLO backup selections”](#) on page 1664.
- On the **Options** tab, set Delta File Transfer, encryption, and compression options for this backup selection.
See [“Options for a DLO backup selection”](#) on page 1666.

8 Click **OK** twice.

General options for DLO backup selections

When a backup selection is created, the name, description and path to be backed up are specified in the backup selection general dialog box. When the backup selection is created, the name, description and backup path can be modified as needed.

See [“Adding a DLO backup selection to a profile”](#) on page 1660.

Table Q-27 General options for DLO backup selections

Item	Description
Name	Shows a descriptive name for the backup selection.
Description	Shows a clear description of the backup selection. This description may include, for example, the folder selected, the group of users it was created for, or the purpose for creating the backup selection.
Type a folder name	Lets you add a specific folder to the backup selection. Type the path to the folder, including the folder name. For example, to add a folder named My Data on drive C, type C:\MyData. You can use macros to define the folders you want to back up. See “How to use DLO macros in backup selections” on page 1667.

Table Q-27 General options for DLO backup selections (*continued*)

Item	Description
Select a pre-defined folder	Lets you choose a predefined folder from the list provided. You can use macros to enter predefined folders. See “How to use DLO macros in backup selections” on page 1667.
Include subfolders	Lets you back up all subfolders in the specified directory. This option is selected by default. On the computers that run Microsoft Windows Vista, this option does not include the Music, Pictures, or Videos folders in the backup selection. On the computers that run previous versions of Microsoft Windows, this option does include My Music, My Pictures, and My Videos folders in the backup selections.

Including and excluding files or folders from a DLO backup selection

Each backup selection can be configured to either include all files and folders, or to include or exclude specific files and folders. In addition, specific file types or folders can be specified for inclusion or exclusion using wildcards.

Files and folders can also be excluded from all backups for all users using global exclude filters. Several file types are excluded by default. These global excludes can be viewed or modified in the Global Excludes dialog box.

See [“About configuring global exclude filters in DLO”](#) on page 1687.

To include or exclude files or folders from a backup selection

- 1 Open the Backup Selection dialog box.
See [“Adding a DLO backup selection to a profile”](#) on page 1660.
- 2 On the Include/Exclude tab, select one of the following options:

- | | |
|--|---|
| Include all file types | Select this option to include all file types in this backup selection. |
| Include and exclude only the items listed below | Select this option to include or exclude only specific files or file types. When this option is selected, a wildcard include is added to back up all files not specifically excluded. |

- 3 To add a filter to the Include/Exclude list, verify that you selected Include and exclude only the items listed below, and click **Add Include** or **Add Exclude**.
- 4 If you selected Add Exclude, you will be notified that all previously backed up files matching this exclude will be deleted from this backup selection. Click **Yes** to continue or **No** to cancel.
- 5 Select the appropriate options.
See [“Add Global Exclude Filter options”](#) on page 1689.
- 6 Click **OK**.

Include/Exclude options for DLO backup selections

Each backup selection can be configured to either include all files and folders, or to include or exclude specific files and folders.

See [“Including and excluding files or folders from a DLO backup selection”](#) on page 1662.

Table Q-28 Include/Exclude options for DLO backup selections

Item	Description
Include all file types	Includes all file types in this backup selection.
Include and exclude only the items listed below	Lets you include or exclude only specific files or file types. When this option is selected, a wildcard include is added to back up all files not specifically excluded.

About revision control in DLO

Revisions are versions of a file at a specific point in time. You configure revision settings when you create a backup selection. When a file is changed and backed up, DLO stores a new revision. DLO stores and maintains a specific number of revisions for all files in a backup selection. Because backup selections are configured separately, the number of revisions that are retained in each backup selection can vary. When the number of revisions is exceeded, DLO removes the oldest revision.

You can limit the number of revisions that are retained in a given period of time. If you back up a document frequently while you work on it, all of your revisions could potentially be a few minutes apart. By specifying that you want to retain only two revisions every 24 hours, at least 120 minutes apart, you can retain older revisions for a longer period of time. While some intermediate versions are not

retained, it does support situations in which returning to an older revision is needed.

Another consideration in determining the number of revisions to retain is the amount of storage space that is required to store the data. The amount of space that is required for backups can be estimated by multiplying the number of revisions that are retained by the amount of data protected.

For example, if you retain three revisions of each file and have 10 MB of data to back up, approximately 30 MB of disk space are required if file sizes remain consistent between revisions.

Although compression can improve the space utilization, it varies significantly with file type and other factors. Typical compression ratios are approximately 2:1, so in the example, the maximum disk space usage might be reduced to approximately 15 MB.

About file grooming in DLO

The Desktop Agent grooms revisions based on backup selection settings. Revisions are groomed as new revisions are created. The oldest revision is deleted when a new revision is created that exceeds the limit.

Maintenance grooming is the process of removing backups of deleted files. It occurs at most one time every 24 hours. Maintenance grooming occurs during the first backup that runs after 24 hours have passed since the last maintenance grooming.

Revision Control options for DLO backup selections

For each backup selection, you can specify the following settings:

- The number of revisions that are retained in the desktop and network user data folders.
- The amount of time between revisions.

See [“Adding a DLO backup selection to a profile”](#) on page 1660.

Table Q-29 Revision Control options for DLO backup selections

Item	Description
Keep x revisions in the desktop user data folder	<p>Indicates the number of revisions to keep in the desktop user data folder for each file in the backup selection.</p> <p>When Outlook PST files or Lotus Notes NSF files are backed up incrementally, only one revision is maintained regardless of the number of revisions that are set in the backup selection.</p>
Limit to	<p>Limits the number of revisions that are retained in a given amount of time.</p> <p>You can specify the following:</p> <ul style="list-style-type: none"> ■ The number of versions to retain. ■ The time period during which you want to retain the versions. ■ The minimum amount of time that must elapse between backups in this backup selection. <p>The oldest revision is deleted when a new revision is created that exceeds one of these limits.</p>
Keep x revisions in the network user data folder	<p>Indicates the number of revisions to keep in the network user data folder for each file in the backup selection.</p>
Limit to	<p>Limits the number of revisions that are retained in a given amount of time.</p> <p>You can specify the following:</p> <ul style="list-style-type: none"> ■ The number of versions to retain. ■ The time period during which you want to retain the versions. ■ The minimum amount of time that must elapse between backups in this backup selection. <p>The oldest revision is deleted when a new revision is created that exceeds one of these limits.</p>
Discard all revisions in the desktop user data folder older than	<p>Indicates the number of days after which all revisions in the desktop user data folder are deleted.</p> <p>The most recent revision is not discarded.</p>
Discard all revisions in the network user data folder older than	<p>Indicates the number of days after which all revisions in the network user data folder are deleted.</p> <p>The most recent revision is not discarded.</p>

Setting options for a DLO backup selection

DLO backup selections can be further customized for Delta File Transfer, compression, and encryption. In addition, you can specify how long to keep backup files after the original source files are deleted.

To set options for a DLO backup selection

- 1 Open the Backup Selection dialog box.
 See [“Adding a DLO backup selection to a profile”](#) on page 1660.
- 2 On the Options tab, select the appropriate options
 See [“Options for a DLO backup selection”](#) on page 1666.
- 3 Click **OK** to save the backup selection.

Options for a DLO backup selection

DLO backup selections can be further customized for Delta File Transfer, compression, and encryption. In addition, you can specify how long to keep backup files after the original source files are deleted.

See [“Setting options for a DLO backup selection”](#) on page 1666.

Table Q-30 Options for a DLO backup selection

Item	Description
Delta File Transfer	<p>Indicates that each time a file is backed up, only the part of the file that has changed is transferred and stored in the network user data folder. In addition, Delta file transfer uses compression. Enabling this option requires that you have installed and configured a maintenance server.</p> <p>See “Adding a new Maintenance Server” on page 1672.</p>
Compression	<p>Enables files in this backup selection to be compressed for data transfer over the network and for storage in the desktop and network user data folders.</p> <p>This setting applies to files that were created after this feature is activated. Previously stored files are not compressed.</p> <p>Delta File Transfer also uses compression.</p>

Table Q-30 Options for a DLO backup selection (*continued*)

Item	Description
Encryption	<p>Encrypts files for transfer and to store files from this backup selection in an encrypted format in the network user data folder.</p> <p>This setting applies to files that were transmitted and stored after this feature is activated. Previously stored files are not encrypted.</p> <p>The Advanced Encryption Standard (AES) and a 128-bit key length are used. In the desktop user data folder, versions are stored unencrypted. In the network user data folder, versions are stored encrypted. Transfer over the network is encrypted.</p>
Desktop user data folder after	Indicates the number of days for DLO to wait between when a source file is deleted from the desktop and when all versions of that file are deleted from the desktop user data folder.
Network user data folder after	Indicates the number of days for DLO to wait between when a source file is deleted from the desktop and when all versions of that file are deleted from the network user data folder.

How to use DLO macros in backup selections

You can type macros into the Type a folder name field of the backup selection dialog box to automatically back up specific folders.

See [“General options for DLO backup selections”](#) on page 1661.

The following macros are supported:

Table Q-31 Folder macros for use with backup selections

Backup Selection Macro	Folders backed up
%LOCALFIXEDDRIVES%	<p>All local fixed drives.</p> <p>Note: DLO is not designed to back up removable media. Trying to back up a floppy disk or CDROM may result in errors.</p>

Table Q-31 Folder macros for use with backup selections (*continued*)

Backup Selection Macro	Folders backed up
%MACHINENAME%	Represents the desktop user's computer name. Example: C:\documents\%machine name% represents C:\documents\UsersMachineName.
%CURRENTUSERNAME%	Represents the user name of the currently logged-on user. Example: If the local administrator is logged on to the computer, C:\documents\%current username% represents 'C:\documents\Administrator'
%CURRENTUSERPROFILE%	All files and folders in the C:\Documents and Settings\current user profile directory (for Windows XP) or the C:\Users\current user profile directory (for Windows Vista).
%CURRENTUSERMYDOCS%	The My Documents directory for the user who is logged on.
%CURRENTUSERFAVORITES%	The Favorites directory for the user who is logged on.
%CURRENTUSERPRINTHOOD%	The Printers directory for the user who is logged on.
%CURRENTUSERNETHOOD%	The Network Locations directory for the user who is logged on.
%CURRENTUSERDESKTOP%	The Desktop directory for the user who is logged on.
%CURRENTUSERRECENT%	The Recent Files directory for the user who is logged on.
%PROGRAMFILES%	The Windows Program Files directory. Example: %PROGRAMFILES%\lotus\notes\data\archives
%LOCALAPPDATA%	The Windows local application data directory: <ul style="list-style-type: none"> ■ On Windows XP: Documents and Settings\<user_name>\Local Settings\Application Data</user_name> ■ On Windows Vista: Users\<user_name>\AppData\Local</user_name>

The following additional pre-defined folder macros are available for selection in the backup selection dialog box:

Table Q-32 Macros for pre-defined folders in the backup selection dialog box

Folder Name	Pre-Defined Folder Macro	Folders Backed Up
My Documents	%CURRENTUSERMYDOCS%	The My Documents directory for the user who is logged on.
My Music	%CURRENTUSERMYMUSIC%	The My Music folder for the user who is logged on.
My Pictures	%CURRENTUSERMYPICTURES%	The My Pictures folder for the user who is logged on.
My Videos	%CURRENTUSERMYVIDEO%	The My Videos folder for the user who is logged on.
Desktop	%CURRENTUSERDESKTOP%	The Desktop directory for the user who is logged on.
Favorites	%CURRENTUSERFAVORITES%	The Favorites directory for the user who is logged on.
PrintHood	%CURRENTUSERPRINTHOOD%	The Printers directory for the user who is logged on.
NetHood	%CURRENTUSERNETHOOD%	The Network Locations directory for the user who is logged on.
Recent	%CURRENTUSERRECENT%	The Recent Files directory for the user who is logged on.
All local fixed drives	%LOCALFIXEDDRIVES%	All local fixed drives.

When you enter a path that uses a macro, a backslash is automatically added immediately following the macro. For example, if you type %LOCALFIXEDDRIVES%\Documents, an extra backslash is added and it appears as "x:\\Documents" in the Desktop Agent backup selection advanced view. It does not show at all in the Desktop Agent backup selection Standard view. The correct way to type this macro is%LOCALFIXEDDRIVES%Documents. This macro properly resolves to x:\Documents.

Modifying a DLO backup selection

Profile backup selections can be modified from the DLO Administration Console.

To modify a backup selection

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the Settings pane, click **Profiles**.
- 3 In the results pane, select the profile you want to modify.
- 4 In the task pane, under General Tasks, click **Properties**.
- 5 On the Backup Selections tab, select the backup selection you want to modify, and then click **Modify**.
- 6 Click **OK** to indicate that you read the message stating that modifying this backup selection changed all profiles that use this selection.
- 7 Change the backup selection as needed.

See [“General options for DLO backup selections”](#) on page 1661.

See [“Including and excluding files or folders from a DLO backup selection”](#) on page 1662.

See [“Revision Control options for DLO backup selections”](#) on page 1664.

See [“Setting options for a DLO backup selection”](#) on page 1666.

- 8 Click **OK** twice.

Deleting DLO backup selections

Before you can delete a backup selection, you must be sure that it is not in use by any profiles. When you delete a backup selection from one profile, DLO deletes it from every profile.

When you delete a backup selection, the backup versions are deleted in the same manner as when source files are deleted. They are groomed after the number of days specified in the backup selection.

To delete a backup selection

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the Settings pane, click **Profiles**.
- 3 In the results pane, select the profile that contains the backup selection you want to delete.
- 4 In the task pane, under General Tasks, click **Properties**.
- 5 On the Backup Selections tab, select the backup selection you want to delete.

- 6 Click **Delete**.
- 7 Click **Yes**.

About Delta File Transfer

The Delta File Transfer feature enables incremental transfer and storage of backup data. When this option is enabled, the initial backup requires transfer of the entire file. Subsequent backups require only the transfer of the parts of the file that have changed, reducing the bandwidth required and improving backup speed.

Delta File Transfer is not limited to certain programs or file types. However, you can exclude certain file types. Default excludes are configured for Delta File Transfer because these file types do not benefit from this technology. These file types are already highly compressed.

See [“About configuring global exclude filters in DLO”](#) on page 1687.

Delta File Transfer is only used to transfer and store backup files in the network user data folder. Backup files that are stored in the Desktop User Data Folder are not stored using deltas. When a Desktop Agent user works offline, the local revisions are stored in their entirety in the desktop user data folder. When the user works online again, Delta File Transfer is used to transfer data to the network user data folder.

Requirements for Delta File Transfer

Delta File Transfer requires the use of the DLO maintenance server. The maintenance server manages the deletion of previous delta revisions from storage locations. The maintenance server is only required when the Delta File Transfer option is enabled, but it is installed by default when DLO is installed. If the media server is also the Storage Location host, then no additional steps are required to configure the maintenance server.

Only one maintenance server is required. However, in large installations it may be more efficient to have one maintenance server for each Storage Location host (File Server).

Maintenance Server technical information and tips

The Desktop Agent uses Windows RPC over named pipes to communicate with the maintenance server. For the maintenance server to function, named pipe traffic must not be blocked at any point between the DLO Client and the maintenance server.

The rolloff operation for delta revisions can require significant bandwidth. For this reason, the maintenance server should be installed on the computer that hosts the Storage Location.

However, there are situations where the maintenance server cannot be installed on the same computer as the Storage Location server. For example, the maintenance server cannot be installed on a NAS device. In this case, the maintenance server should be installed on a computer with a high bandwidth connection to the Storage Location.

A maintenance server can manage one or more Storage Locations. A maintenance server always manages the Storage Locations that are located on same computer as the maintenance server. The maintenance server can be configured to manage additional Storage Locations hosts from the DLO Administration Console. The maintenance server uses delegation to access remote Storage Locations.

See [“Configuring a maintenance server for delegation”](#) on page 1673.

How to enable Delta File Transfer for a backup selection

Delta File Transfer is off by default. However, you can enable it for a given backup selection.

See [“About Delta File Transfer”](#) on page 1671.

In addition, if a maintenance server manages file servers that are on a target other than itself, the maintenance server must be configured for delegation.

See [“Configuring a maintenance server for delegation”](#) on page 1673.

Delta File Transfer can also be selected as the default compression type. If the default compression setting is changed to Delta, all new backup selections use Delta compression by default.

See [“About default DLO settings”](#) on page 1625.

Adding a new Maintenance Server

A default maintenance server is installed with DLO. You can also install a stand-alone maintenance server from the installer.

After you install a new maintenance server, you must add the maintenance server to DLO. After you add the maintenance server to DLO, you can then specify the files servers that it should manage.

See [“Reassigning a file server”](#) on page 1675.

To add a new Maintenance Server

- 1 Verify that the new maintenance server has been installed.
- 2 On the DLO Administration Console, on the DLO navigation bar, click **Setup**.
- 3 In the task pane, under Manage Tasks, click **Maintenance servers**.
- 4 Click **Add**.
- 5 Navigate to the computer where the maintenance server is installed, and then select the computer.
- 6 Click **OK**.

Configuring a maintenance server for delegation

You can configure a maintenance server to manage Storage Locations that are hosted by a different computer. Then, you must configure the maintenance server to access these locations on behalf of desktop users that run the Desktop Agent. This configuration is managed by using Active Directory.

For detailed information on delegating Active Directory administration, see the following Microsoft Web site:

<http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/directory/activedirectory/actdid1.mspx>

Table Q-33 Configuring a maintenance server for delegation

Step	Description
Step 1	Verify that the following conditions are met: <ul style="list-style-type: none"> ■ Domains are Windows 2000 or later. NT 4 domains are not supported. ■ Both the Desktop Agent user's account and the maintenance service's account must be in the same forest. ■ Desktop Agent user and computer accounts must be in mutually trusted domains. ■ Desktop and server operating systems must be Windows 2000 or later
Step 2	Confirm that the desktop user account is configured for delegation. See " Confirming that the desktop user's account is configured for delegation " on page 1674.

Table Q-33 Configuring a maintenance server for delegation (*continued*)

Step	Description
Step 3	Confirm that the server process account is trusted for delegation. See “Confirming that the server process account is trusted for delegation” on page 1674.

Confirming that the desktop user's account is configured for delegation

The following procedure is part of the process to configure a maintenance server for delegation.

See [“Configuring a maintenance server for delegation”](#) on page 1673.

To confirm that the desktop user's account is configured for delegation

- 1 Log on to the domain controller by using a domain administrator account.
- 2 On the taskbar, click **Start > Programs > Administrative Tools > Active Directory Users and Computers**.
- 3 Under the domain, click the **Users** folder.
- 4 Right-click the user account to be delegated, and then click **Properties**.
- 5 On the Account tab, in the Account options list, verify that the following option is not selected:
Account is sensitive and cannot be delegated
- 6 Click **OK**.

Confirming that the server process account is trusted for delegation

The following procedure is part of the process to configure a maintenance server for delegation.

See [“Configuring a maintenance server for delegation”](#) on page 1673.

To confirm that the server process account is trusted for delegation

- 1 Log on to the domain controller by using a domain administrator account.
- 2 On the taskbar, click **Start > Programs > Administrative Tools > Active Directory Users and Computers**.
- 3 Right-click the Computers folder, and then click **Properties**.
- 4 Right-click the computer on which the maintenance server runs, and then click **Properties**.

- 5 On the General page, click **Trust computer for delegation**.
- 6 Click **OK**.

Changing the default maintenance server

When DLO is installed, a maintenance server is installed and set as the default maintenance server. New storage locations are automatically assigned to the default maintenance server when they are created. If you want new storage locations to be assigned to a different maintenance server by default, you must change this setting.

To change the default maintenance server

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the task pane, under Manage Tasks, click **Maintenance servers**.
- 3 In the Maintenance Servers list, select the maintenance server that you want to set as the default.
- 4 Click **OK**.

Reassigning a file server

You can reassign a file server to another maintenance server that DLO recognizes. For example, when you create a new storage location, it is automatically assigned to the default maintenance server. You may want to reassign it to a different maintenance server.

To reassign a file server

- 1 Verify that the new maintenance server has been installed and configured.
- 2 On the DLO navigation bar, click **Setup**.
- 3 In the task pane, under Manage Tasks, click **Maintenance servers**.
- 4 Select the maintenance server that currently manages the file server.
- 5 Click **Edit**.
- 6 Select the file server you want to reassign.
- 7 Click **Reassign**.
- 8 Select the new maintenance server from the drop-down menu.
- 9 Click **OK** three times.

About DLO Storage Locations

Storage Locations are locations on network computers where network user data folders are automatically created.

The Desktop and Laptop Option stores each user’s data in the following places:

Table Q-34 Locations where user data is stored

Location	Description
Desktop user data folder on the user’s computer	Provides protection and restore capabilities even when the computer is disconnected from the network.
Network user data folder on the network	Provides an additional level of protection, and enables the files to be backed up to secondary media when the server is backed up.

When a user is automatically added to DLO using an Automated User Assignment, a network user data folder is created in a Storage Location as specified in the Automated User Assignment. If network shares already exist for desktop users, they can be specified as network user data folders when users are manually added to DLO. If existing network shares are used as network user data folders, Storage Locations are not used.

DLO supports the use of hidden shares (for example; "Share\$") as Storage Locations on NTFS volumes or as network user data folders for FAT32 volumes. However, you cannot create them with the DLO Administration Console. They must be created and configured manually.

See [“How to use hidden shares as Storage Locations”](#) on page 1677.

Supported Storage Location configurations

The following table summarizes supported configurations for DLO Storage Locations:

Table Q-35 Storage Location configuration support

Description	Supported	Not Supported
All media server platforms	X	
Windows 2000 NAS/SAK NAS devices	X	
Local media server direct-attached storage	X	
SAN	X	

Table Q-35 Storage Location configuration support (*continued*)

Description	Supported	Not Supported
Windows-networking accessible NAS Devices (Quantum, Network Appliance, etc.)	X	
FAT, FAT32, and NTFS partitions are supported as Storage locations, although FAT and FAT32 are not recommended. NTFS is the preferred file system for Storage Locations	X	
NetWare 3.1x, 4.x, or E-Directory Storage Locations		X
UNIX file systems or SAMBA shares on UNIX systems		X

How to use hidden shares as Storage Locations

DLO supports the use of hidden shares (for example; "Share\$") as Storage Locations on NTFS volumes or as network user data folders for FAT32 volumes. However, you must create and configure them manually. They cannot be created with the DLO Administration Console. Hidden shares cannot be used for FAT based Storage Locations.

The following permissions should be used:

Table Q-36 Permission settings for hidden shares

Drive Type	User or Group	Permissions
Share permissions on NTFS volumes	Administrator	Allow Full Control, Change, Read
	Everyone	Allow Full Control, Change, Read
Security permissions on NTFS volumes	Administrator	Full control
	Everyone	Allow Read & Execute Allow List Folder Contents Allow Read

Table Q-36 Permission settings for hidden shares (*continued*)

Drive Type	User or Group	Permissions
	Special security permissions or advanced settings	Allow Traverse Folder/Execute File Allow List Folder/Read Data Allow Read Attributes Allow Read Extended Attributes Allow Read Permissions
Advanced security permissions on NTFS volumes	Administrator	Allow Full Control
	Everyone	Allow Traverse Folder / Execute File Allow List Folder / Read Data Allow Read Attributes Allow Read Extended Attributes Allow Read Permissions
Share permissions on FAT volumes	Administrator	Allow Full Control, Change, Read
	Owner	Allow Full Control, Change, Read
	Full Admin Group	Allow Full Control, Change, Read
	Limited Admin Group	Allow Read

Creating DLO Storage Locations

A Storage Location should be used by only one media server. When multiple media servers use the same Storage Location, if the Storage Location is deleted from one media server, then the other media server can no longer access it.

Storage Locations must be in a Windows Domain or Active Directory. Computers running the Desktop Agent can be outside of a Windows Domain or Active Directory. However, the computers must authenticate with the domain or directory

to access the media server or Storage Locations. Users are prompted to provide domain credentials when the Desktop Agent is launched.

If your original files reside on an NTFS volume, then the desktop user data folder and the network user data folder should also be NTFS. If your original files are on NTFS and either the desktop user data folder or network user data folder are on a FAT or FAT32 volume, you may see duplicate entries in the Restore and Restore Search screens. If duplicates do appear, you can select either file to restore.

After you create a Storage Location, you cannot change it. However, you can delete a Storage Location if no users or Automated User Assignments are assigned to them. You can move users to new Storage Locations.

See [“Moving Desktop Agent users to a new network user data folder”](#) on page 1701.

If you receive errors when creating Storage Locations, verify that the logon account for the service named MSSQL\$BKUPEXCDLO has sufficient rights to create directories and change permissions for the Storage Locations. Use the Windows Service Control Panel to change the logon account for the MSSQL\$BKUPEXCDLO instance. You can avoid these problems if you specify a domain account when you install Backup Exec.

To create DLO Storage Locations

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the task pane, under Settings Tasks, click **New Storage Location**.
- 3 Select the appropriate options.
See [“New Storage Location options”](#) on page 1679.
- 4 Click **OK**.

New Storage Location options

When you create a new Storage Location, you must enter the following information.

See [“Creating DLO Storage Locations”](#) on page 1678.

Table Q-37 **New Storage Location options**

Item	Description
Computer name	Indicates the name of the computer on which you want to create the Storage Location.
Path	Indicates the location on the computer where you want to create the Storage Location. Storage Locations should be in the same domain as the media server or in a domain that trusts the media server’s domain.

Table Q-37 **New Storage Location options** (*continued*)

Item	Description
Storage Location name	Indicates the name for the new Storage Location. The name cannot contain any of the following characters: \'@#\$\$%^&*()=+ /{}[]
Summary	<p>Lists the location and format of network user data folders that will be created for new users who are assigned to this Storage Location. Network user data folders are automatically created in the Storage Location.</p> <p>DLO uses the %USERDOMAIN% and %USERNAME% variables to determine the actual folder path for each user who is assigned to a Storage Location. DLO uses the user's domain and user name to create a unique network user data folder name for that user. If the user is logged on with credentials that do not allow access to the Storage Location, the user is prompted to enter alternate domain credentials.</p> <p>The network administrator can access this folder, but cannot configure the variables.</p>

Configuring a remote Windows share or NAS device for DLO Storage Locations

You can create DLO Storage Locations on remote Windows shares or network-attached storage devices.

In addition, you can configure Storage Locations so that the DLO administration service does not run as an administrator-level user, but the DLO administration groups are assigned the appropriate permission levels on a pre-existing share.

See [“Configuring a remote Windows share or NAS device for DLO Storage Locations using non-administrator case”](#) on page 1681.

Table Q-38 **Configuring a remote Windows share or NAS device for DLO Storage Locations**

Step	Description
Step 1	Validate that DLO 5.1 MP1 or later is installed.
Step 2	Ensure that the account credentials that are used for DLO services have full administrator rights to the remote storage location or NAS device.

Table Q-38 Configuring a remote Windows share or NAS device for DLO Storage Locations (*continued*)

Step	Description
Step 3	Make sure that the volume desired to be used for DLO has been assigned a drive letter on the remote storage location or NAS device. See hardware vendor documentation on share creation and naming.
Setp 4	Create a new Storage Location. Use the browse feature to indicate the location on the computer where the Storage Location will be created. This step insures that the path and the DLO service account are valid.

Configuring a remote Windows share or NAS device for DLO Storage Locations using non-administrator case

You can create DLO Storage Locations on remote Windows shares or network-attached storage devices.

In addition, you can configure Storage Locations so that the DLO administration service does not run as an administrator-level user, but the DLO administration groups are assigned the appropriate permission levels on a pre-existing share.

To configure Storage Locations using non-administrator case

- 1 Configure DLO to use existing domain groups to automatically manage access to network user data folders.

Check the Automatically grant DLO Administrators access to network user data folders check box and provide the required domain groups. Provide two groups: a group for full-DLO administrators and a group for limited-DLO administrators.

See [“About administrator accounts in DLO”](#) on page 1618.

- 2 From the Administrator Account Management dialog box, add the appropriate domain user accounts to the account manager. If the user will have full administrator rights, check the "Grant administrator full restore privileges" checkbox in the Add Administrator Account dialog. In addition to other users, be sure to grant the DLO Administration Service full restore privileges.
- 3 Create a folder on the remote storage location using an administrator, or administrator equivalent user.

- 4 Share the new folder. Ensure that 'Everyone' has full-access to the share.
- 5 Modify the folder's security permissions such that the full-DLO administrator group has full-control of the folder and that the limited-DLO administrator group has modify-control of the folder.
- 6 Using the DLO console, create a new Storage Location. Specify the computer name, drive and path, and share name for the folder just created.

Do not use the browse buttons at any point during the storage location creation as they will cause the process to fail.
- 7 Once the required fields are completed, click **OK**

Deleting DLO Storage Locations from a remote Windows share or NAS device

You cannot delete a Storage Location if it was created manually and the DLO Admin Service does not have full administrator rights to the server hosting the DLO Storage Location.

Table Q-39 Deleting DLO Storage Locations from a remote Windows share or NAS device

Step	Description
Step 1	Move or delete all users in the Storage Location.
Step 2	Manually remove the Storage Location share and folder from the server.
Step 3	Delete the Storage Location from the DLO Admin Console.

Deleting DLO Storage Locations

Before you can delete DLO Storage Locations, you must delete or reassign users and Automated User Assignments that use the Storage Location. The Storage Location that is associated with a user or Automated User Assignment is listed when you select Users or Automated User Assignments from the Setup view.

See [“Modifying Automated User Assignments”](#) on page 1686.

See [“Deleting Automated User Assignments”](#) on page 1687.

When a Storage Location is created using an existing share on a remote computer and DLO does not have full computer rights, the Storage Location cannot be deleted from the Administration Console. To remove the Storage Location, first delete

the Storage Location share and then delete the Storage Location from the Administration Console.

To delete Storage Locations

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, expand the file servers list by clicking the '+' next to File Servers.
- 3 In the selection pane, click the File Server on which the Storage Location resides.
- 4 In the results pane, click the storage location you want to delete.
- 5 In the task pane, under General Tasks, click Delete.
- 6 Click Yes.

About Automated User Assignments

Automated User Assignments are instructions that are applied when the Desktop Agent is first run on a desktop. The Automated User Assignment assigns a profile and network user data folder to each user who is automatically configured by DLO. These settings can be changed from the DLO Administration Console at a later time if necessary.

If a user is added manually to DLO, the DLO administrator selects a Storage Location and a profile. The Automated User Assignment is not used.

See [“About managing Desktop Agent users”](#) on page 1696.

Automated User Assignments are assigned to desktop users based either on their domain and group, or using Active Directory settings. Because users may match the criteria for more than one Automated User Assignment, the Automated User Assignments are prioritized. When the Desktop Agent runs for the first time, the user's domain and group credentials are checked against those of the Automated User Assignment. The Desktop Agent checks the credentials starting with the highest priority assignment. When a match is made, the share and profile that are specified in that Automated User Assignment are assigned to the new user.

If you modify Automated User Assignments, users who have already been configured are not affected. Only new users who are configured with the Automated User Assignment use the new settings.

See [“Changing the priority of Automated User Assignments”](#) on page 1686.

Creating Automated User Assignments

Automated User Assignments are assigned to Desktop Agent users based either on domain and group settings or Active Directory settings. The Automated User Assignment determines which Storage Location and Profile are assigned to the user.

To create a new Automated User Assignment

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Automated User Assignments**.
- 3 In the task pane, under Settings Tasks, click **New user assignment**.
- 4 Select the appropriate options.
 See [“New Automated User Assignment options”](#) on page 1684.
- 5 If you chose to use Active Directory to configure the User Assignment in step 4, configure the Active Directory settings, and then click **OK**
 See [“Active Directory Object options”](#) on page 1685.
- 6 Click **OK**.

New Automated User Assignment options

When you create a new Automated User Assignment, you must set up the following options.

See [“Creating Automated User Assignments”](#) on page 1684.

Table Q-40 **New Automated User Assignment options**

Item	Description
User assignment name	Indicates the name for the Automated User Assignment. The Automated User Assignment name cannot contain the following characters: \ "@#\$%^&*()=+ /{}[]'
Assign using Domain and Group	Matches Desktop Agent users to Automated User Assignments that are based on their domain and group.
Domain	Indicates the domain to which this Automated User Assignment applies.
Group	Indicates the group to which this Automated User Assignment applies.
Assign Using Active Directory	Matches Desktop Agent users to Automated User Assignments that are based on Active Directory settings.

Table Q-40 New Automated User Assignment options (*continued*)

Item	Description
Configure	Configures the User Assignment using Active Directory.
Storage Location	Indicates the Storage Location to be assigned to the users in the selected domain and group.
Profile	Indicates the profile to be assigned to the users in the selected domain and group.

Active Directory Object options

If you chose to use Active Directory to configure an Automated User Assignment, you must complete the following options.

See [“Creating Automated User Assignments”](#) on page 1684.

Table Q-41 Active Directory Object options

Item	Description
Object	For Automated User Assignments, the only option is User.
In LDAP Directory	Indicates the LDAP directory. When selecting Active Directory user accounts, you must select the specific directory that holds the user accounts. Be sure not to select the user groups directory. Browse to or type the exact path of the specific user accounts directory for which you are creating this rule.
Only the objects in this directory	Applies the connection policy to all objects in the specified directory.
Only the objects in this directory that match the criteria below	Applies the connection policy only to those objects in the specified directory that match the criteria entered.
Attributes	Indicates the attribute for which you want to apply the connection policy.
Condition	Indicates the condition to use to match the attribute and the value. Available options include =, <, <>, and >.
Value	Indicates the user-defined criteria that is used to determine matches. Wildcards can be used to specify the value.

Modifying Automated User Assignments

Modifying an Automated User Assignment affects only the users who are added to the assignment after it has been modified. Existing Desktop Agent users are unaffected.

Settings for existing Desktop Agent users can be modified from the Setup view of the DLO Administration Console.

See [“Changing the profile for a Desktop Agent user”](#) on page 1699.

To modify an Automated User Assignment

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Automated User Assignments**.
- 3 In the results pane, select the Automated User Assignment you want to modify.
- 4 In the task pane, under General Tasks, select **Properties**.
- 5 Modify the Automated User Assignment properties.

Changing the priority of Automated User Assignments

When you create an Automated User Assignment, DLO assigns a priority to it. The priority determines which Automated User Assignment is used when a user is a member of more than one domain and group. The most recently created Automated User Assignments have the lowest priority. You can change the priority of Automated User Assignments.

To change the priority of Automated User Assignments

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Automated User Assignments**.
- 3 In the results pane, select the Automated User Assignment for which you want to change the priority.
- 4 In the task pane, under Settings Tasks, select **Move priority up** or **Move priority down**.

Viewing Automated User Assignment properties

You can view the properties of Automated User Assignments.

To view Automated User Assignments

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Automated User Assignments**.

- 3 In the results pane, select an Automated User Assignment.
- 4 In the task pane, under General Tasks, select **Properties**.

Deleting Automated User Assignments

You can delete Automated User Assignments when you no longer need them.

To delete an Automated User Assignment

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Automated User Assignments**.
- 3 Click the Automated User Assignment to be deleted.
- 4 In the task pane, under General Tasks, click **Delete**.
- 5 Click **Yes**.

About configuring global exclude filters in DLO

DLO global exclude options enable you to do the following:

- Specify the attributes of files that you want to exclude from all backups.
- Specify files that you do not want to compress, encrypt, or back up with Delta File Transfer.
- Exclude attachments to emails or specific email folders from backup.

Global excludes apply to both Profile backup selections and user-created backup selections for all Desktop Agent users who back up to the media server on which the excludes are configured.

Excluded files are listed in the following locations:

Table Q-42 Location of excluded files

Component	Location of excluded files
Desktop Agent	On the Include/Exclude tab in the advanced view.
DLO Administration Console	On the Include/Exclude tab for a profile's backup selection.

Items that are configured for the global exclude list are not available for selection on the selection list.

When you add a global exclude, all previous backups that match the global exclude are deleted when their retention period expires. You set the retention period when you set up backup selections.

See [“Adding backup selections in the Desktop Agent's advanced view”](#) on page 1768.

In the following circumstances, the backups are deleted immediately during the next maintenance cycle:

- The retention period is set to 1 day.
- The default value is used for the time between maintenance cycles. The default value is 24 hours. If you change this value, it affects file retention for all files. See [“Changing default DLO global settings”](#) on page 1626.

Information about how to configure global excludes is available.

See [“Specifying files and folders to exclude from all DLO backups”](#) on page 1688.

See [“Excluding email from all DLO backups”](#) on page 1690.

See [“Excluding files and folders from compression”](#) on page 1692.

See [“Excluding files and folders from encryption”](#) on page 1693.

See [“Excluding files and folders from Delta File Transfer”](#) on page 1694.

See [“About using DLO macros to define global excludes”](#) on page 1696.

Specifying files and folders to exclude from all DLO backups

File and Folder global excludes specify which files and folders, or file and folder types, to exclude from all backups for all users.

To specify files and folders to exclude from all backups

- 1 On the **Tools** menu, click **Global Excludes**.
- 2 On the **Files/Folders** tab, do any of the following:

To exclude all files greater than a specific size	Check the Exclude all files greater than check box, and then enter a size in KB
---	--

To Exclude all files modified before a specified date	Check the Exclude all files modified before check box, and then enter a date
---	---

To add a new Files/Folders global exclude	Click Add , and then enter the name of the file or folder that you want to exclude.
---	--

See [“Add Global Exclude Filter options”](#) on page 1689.

- 3 Click **OK**.

Global Excludes options

Global excludes specify which items are excluded from all backups for all users.

Table Q-43 Global Excludes options

Item	Description
Filter	Lists an existing filter.
Description	Lists a description of an existing filter.
Apply To	Indicates whether an existing filter applies to files or folders.
Exclude all files greater than	Lets you exclude all files that are larger than the size you select.
Exclude all files modified before	Lets you exclude all files that were changed before the selected date.
Add	Lets you enter the name of the file or folder that you want to exclude.
Edit	Lets you change the properties of a global exclude.
Delete	Lets you remove a global exclude.

Add Global Exclude Filter options

You can exclude specific files, folders, and types of files and folders.

See [“Specifying files and folders to exclude from all DLO backups”](#) on page 1688.

Table Q-44 Add Global Exclude Filter options

Item	Description
Filter	<p>Determines which files or folders are excluded. You can use a file name, wildcard, or macro for the files you want to exclude.</p> <p>Examples:</p> <p>Wildcard: *.tmp</p> <p>File name: pagefile.sys</p> <p>Macro: %WINDIR%</p> <p>When using wildcards, you must use the asterisk (*) wildcard. For example, *.tmp returns all results with the .tmp extension while .tmp returns only files explicitly named .tmp.</p>
Extensions	Lets you select the types of files to include or exclude.

Table Q-44 Add Global Exclude Filter options (*continued*)

Item	Description
Description	Indicates a description of the global exclude.
Apply to	Indicates whether this global exclude should apply to files, folders, or both files and folders.

Excluding email from all DLO backups

You can exclude email attachments and messages from all backup jobs for all users.

Note: Lotus Notes emails cannot be filtered by attachment size or type.

To exclude email from all backups

- 1 On the **Tools** menu, click **Global Excludes**.
- 2 On the **Email** tab, do one of the following:

To exclude email attachments that are larger than a specific size

Do the following in the order listed:

- Check the **Exclude all attachments greater than** check box.
- Enter a size in KB.

This feature does not apply to Lotus Notes emails.

To exclude email messages that were received before a specified date

Do the following in the order listed:

- Check the **Exclude all messages received before** check box.
- Enter a date.

To exclude certain types of email attachments or certain email folders

Do the following in the order listed:

- Click **Add**.
- Enter the type of attachment that you want to exclude or the name of the folder that you want to exclude.
 See [“Add Global Email Exclude Filter options”](#) on page 1691.
- Click **OK**.

- 3 Click **OK**.

Global Exclude options for Email

You can exclude email attachments and messages from all backup jobs for all users.

See [“Excluding email from all DLO backups”](#) on page 1690.

Table Q-45 Global Exclude options for Email

Item	Description
Filter	Shows the name of the item that is excluded from backups.
Description	Shows the user-defined description of the item that is excluded from backups.
Apply to	Indicates whether the excluded item applies to email attachments or to specific folders.
Exclude all attachments greater than	Lets you exclude email attachments that are larger than a specific size.
Exclude all messaged received before	Lets you exclude email messages that were received before a specific date.

Add Global Email Exclude Filter options

You can exclude email attachments and messages from all backup jobs for all users.

See [“Excluding email from all DLO backups”](#) on page 1690.

Table Q-46 Add Global Email Exclude Filter options

Item	Description
Attachment file type	<p>Determines which attachment file types are excluded from backup by the global exclude.</p> <p>Lotus Notes emails cannot be filtered by attachment type.</p> <p>Filters can be file names or wildcards.</p> <p>Examples:</p> <p>Wildcard: *.tmp</p> <p>File name: pagefile.sys</p> <p>When using wildcards, you must use the asterisk (*) wildcard. For example, *.tmp returns all results with the .tmp extension while .tmp returns only files explicitly named .tmp.</p>
Mail folder name	<p>Indicates the name of the mail folder you want to exclude from backup.</p>
Description	<p>Indicates a description of the global exclude.</p>

Excluding files and folders from compression

You can prevent specific files or folders from being compressed by using a global exclude. When you use a global exclude, the types of files or folders that you select remain uncompressed for all users.

To exclude files and folders from compression

- 1 On the Tools menu, click **Global Excludes**.
- 2 On the Compressed Files tab, do any of the following:
 - To exclude all files greater than a specific size from compression, check the **Exclude all files greater than** check box and enter a size in KB.
 - To add a new compressed file global exclude, click **Add**, and then enter the files or folders to exclude.

See [“Global Exclude options for compression”](#) on page 1692.

- 3 Click **OK**.

Global Exclude options for compression

You can prevent specific files or folders from being compressed by using a global exclude. When you use a global exclude, the types of files or folders that you select remain uncompressed for all users.

See [“Excluding files and folders from compression”](#) on page 1692.

Table Q-47 Global Exclude options for compression

Item	Description
Filter	Shows the name of the item that is excluded from compression.
Description	Shows the user-defined description of the item that is excluded from compression.
Apply to	Indicates whether the excluded item applies to files, folders, or both files and folders.
Exclude all files greater than	Lets you exclude files that are larger than a specific size.

Excluding files and folders from encryption

You can prevent specific files or folders, or types of files and folders from being encrypted.

To exclude files and folders from encryption

- 1 On the Tools menu, click **Global Excludes**.
- 2 On the Encrypted Files tab, do any of the following:
 - To exclude files greater than a specific size, check the **Exclude all files greater than** check box, and then enter a size in KB.
 - To add a new encrypted file global exclude, click **Add**, and then enter the files or folders to exclude.
See [“Global Exclude options for encryption”](#) on page 1693.
- 3 Click **OK**.

Global Exclude options for encryption

You can prevent specific files or folders, or types of files and folders from being encrypted.

See [“Excluding files and folders from encryption”](#) on page 1693.

Table Q-48 Global Exclude options for encryption

Item	Description
Filter	Shows the name of the item that is excluded from encryption.
Description	Shows the user-defined description of the item that is excluded from encryption.
Apply to	Indicates whether the excluded item applies to files, folders, or both files and folders.
Exclude all files greater than	Lets you exclude files that are larger than a specific size.

Excluding files and folders from Delta File Transfer

You can prevent files and folders from being included in Delta File Transfer. Some types of files are excluded by default because they do not benefit from Delta File Transfer.

Files and folders that are excluded from Delta File Transfer are compressed with standard compression. However, you can prevent files and folders from being compressed by setting a global exclude for compression.

See [“Excluding files and folders from compression”](#) on page 1692.

To exclude files and folders from Delta File Transfer

- 1 On the Tools menu, click **Global Excludes**.
- 2 On the Delta File Transfer tab, do any of the following:
 - To exclude files greater than a specific size from Delta File Transfer, check the **Exclude all files greater than** check box, and then enter a size in KB.
 - To exclude files smaller than a specific size from Delta File Transfer, check the **Exclude all files less than** check box, and then enter a size in KB.
 - To add a new Delta File Transfer global exclude, click **Add**, and then enter the files or folders to exclude.
 See [“Global Exclude options for Delta File Transfer”](#) on page 1694.
- 3 Click **OK**.
 See [“About Delta File Transfer”](#) on page 1671.

Global Exclude options for Delta File Transfer

You can prevent files and folders from being included in Delta File Transfer.

See [“Excluding files and folders from Delta File Transfer”](#) on page 1694.

Table Q-49 Global Exclude options for Delta File Transfer

Item	Description
Filter	Shows the name of the item that is excluded from Delta File Transfer.
Description	Shows the user-defined description of the item that is excluded from Delta File Transfer.
Apply to	Indicates whether the excluded item applies to files, folders, or both files and folders.
Exclude all files greater than	Lets you exclude files that are larger than a specific size.
Exclude all files less than	Lets you exclude files that are smaller a specific size.

About excluding files that are always open

On desktop computers running Windows XP/2000, the following folders and file types are generally always open and DLO is unable to back them up:

- C:\Windows\System32\Config
- registry hives and logs, including *.DAT.LOG, *.LOG and the files system, SECURITY, default, SAM, and software
- C:\Windows\System32\wbem
- *.EVT
- *.LOG (in particular, STI_Trace.log, WIADEBUG.LOG, WIASERVC.LOG)
- *.DAT (in particular, NTUSER.DAT, USRCLASS.DAT)

To prevent these files from always being listed in the pending files list on the Desktop Agent, add them to the Global Excludes list or the backup selection exclude list.

See [“About backup selections in DLO”](#) on page 1658.

See [“About configuring global exclude filters in DLO”](#) on page 1687.

About using DLO macros to define global excludes

The following macros are typically used for excluding files using the global exclude option, but can also be used in backup selections.

Table Q-50 Global exclude macros

Macro	Folder
%TEMP%	The temp directory for the user who is logged on.
%WINDIR%	The Windows directory. Example: C:\Windows or C:\Winnt
%WEBTEMP%	The Web cache for the user who is logged on.
%RECYCLED%	Recycle bins
%SYSTEM%	The Windows system directory. Example: C:\Windows\system or C:\Winnt\system

About managing Desktop Agent users

The DLO Administrator manages Desktop Agent users from the DLO Administration Console.

From this interface, you can do the following tasks for users or groups of users:

- Add them to DLO manually.
- Enable or disable them.
- Move them to a new network share.
- Assign them to a different profile.

Desktop Agent users are added to DLO either automatically using Automated User Assignments, or manually from the DLO Administration Console.

See [“About Automated User Assignments”](#) on page 1683.

See [“Adding a single desktop user to DLO”](#) on page 1697.

See [“Importing multiple desktop users who have existing network storage ”](#) on page 1699.

See [“Viewing a list of Desktop Agent users”](#) on page 1703.

See [“Changing the profile for a Desktop Agent user”](#) on page 1699.

See [“Enabling or disabling DLO access for a desktop user”](#) on page 1700.

See “[Deleting a user from DLO](#)” on page 1700.

See “[Moving Desktop Agent users to a new network user data folder](#)” on page 1701.

Manually creating new network user data folders

To use a network share as a network user data folder, the folder must have the appropriate security attributes.

To manually create network user data folders and set security attributes

- 1 Create or locate a network share on the computer where backup files will be stored.
- 2 Right-click the share, and then click **Properties**.
- 3 On the Sharing tab, verify that **Share this folder** is selected.
- 4 Click **Permissions**.
- 5 Select the following permissions for user Everyone: Full Control, Change, Read.
- 6 Click **OK**.
- 7 On the Security tab, click **Advanced**.
- 8 Verify that the **Inherit from parent the permission entries that apply to child objects** check box is not checked.
- 9 Add Administrator and Everyone and give them full control permissions.
- 10 In this share, create a data folder for each user who will use this Storage Location, or verify that a data folder already exists.
- 11 Right-click the data folder for a user.
- 12 Click **Properties**.
- 13 Click **Security**.
- 14 Verify that the **Inherit from parent the permission entries that apply to child objects** check box is not checked.
- 15 Add Administrator and the user who will be assigned to the user data folder to the share permission list.
- 16 Set full permission for Administrator and the user.

Adding a single desktop user to DLO

Desktop users can be configured manually rather than with Automated User Assignments. You can use existing network folders that are dedicated to storing

backup data for specific users. These network folders become the DLO network user data folders.

When you add a single desktop user to DLO, you specify the user data folders. However, you can use Storage Locations also.

After you add a desktop user manually, the settings that you assign are applied the first time the desktop user runs the Desktop Agent.

To add a single desktop user to DLO

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Users**.
- 3 In the task pane, under User Tasks, click **New user**.
- 4 Complete the appropriate options.

See [“New User options”](#) on page 1698.

New User options

You must complete the following options to add a new user.

See [“Adding a single desktop user to DLO”](#) on page 1697.

Table Q-51 **New User options**

Item	Description
Enable User	Enables this user to use the Desktop Agent. Clear this check box to prevent the user from using the Desktop Agent.
User	Indicates the user's name.
Profile	Indicates the profile that you want to assign to this user.
Network user data folder	Indicates where this desktop user's backup files will be stored. You must use an existing folder. And, the security attributes must be set for the folder according to your organization's needs. For example, determine which users can access the folder. A Storage Location is not required when an existing network share is used as the network user data folder.
Storage Location	Indicates the Storage Locations to use for this user. The network user data folder for the new user is placed in this Storage Location.

Importing multiple desktop users who have existing network storage

You can use a comma-separated values file to import a list of new users who already have an existing location on the network to store data. This feature cannot be used to import network user data folders for existing Desktop Agent users.

The file must be in the following format and have the following information for each user:

user name, domain, profile, user data folder

For example, JSmith,enterprise,Default,\\Server1\Userdata\jsmith

To import multiple desktop users from a file

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Users**.
- 3 In the task pane, under User Tasks, click **Import users using wizard**.
- 4 Follow the wizard prompts.

Changing the profile for a Desktop Agent user

You can change the properties for a Desktop Agent user.

To change the profile for a Desktop Agent user

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Users**.
Users are listed in the results pane.
- 3 Select the user you want to modify.
- 4 In the task pane, under General Tasks, select **Properties**.
- 5 Select a new profile for this user.

User Properties options

You can change the properties for a Desktop Agent user.

See [“Changing the profile for a Desktop Agent user”](#) on page 1699.

Table Q-52 User Properties options

Item	Descriptions
Enable User	Enables this user to use the Desktop Agent, or clear it to prevent the user from using the Desktop Agent.

Table Q-52 **User Properties options** *(continued)*

Item	Descriptions
User	Shows the name of the user. This field cannot be edited.
Profile	Indicates the profile to apply to this user.
Network user data folder	Shows the location where the user's backup files are to be stored. It cannot be modified. You can move a user to a new location. See "Moving Desktop Agent users to a new network user data folder" on page 1701.

Enabling or disabling DLO access for a desktop user

This option allows you to either allow or prevent a user from using the Desktop Agent.

To enable or disable DLO access for a desktop user

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Users**.
Users are listed in the results pane.
- 3 Select the user you want to modify.
- 4 In the task pane, under General Tasks, select **Properties**.
- 5 Do one of the following:
 - Clear the **Enable user** check box to prevent the desktop user from backing up data with the Desktop Agent.
 - Check the **Enable user** check box to allow the desktop user to back up data with the Desktop Agent.

Deleting a user from DLO

If you want to permanently remove a user from the DLO database, you can delete the user's entry from DLO. Before deleting the user from the DLO Administration Console database, you should uninstall the Desktop Agent from the user's desktop. Otherwise, the user is automatically re-added if the Desktop Agent is run by the user and a matching user assignment exists in DLO. If you cannot uninstall the Desktop Agent from the user's computer, disable the user.

See ["Enabling or disabling DLO access for a desktop user"](#) on page 1700.

To delete a user from the DLO database

- 1 Uninstall the Desktop Agent from the user's computer.
- 2 On the DLO navigation bar, click **Setup**.
- 3 In the selection pane, click **Users**.
- 4 Click the user or users you want to delete.
- 5 In the task pane, under General Tasks, click **Delete**.
- 6 To delete the data that is stored in the user data folder, check **Delete data stored in the user data folder**.

When the Delete data stored... option is selected, backup data is deleted from the network user data folder, but not from the desktop user data folder. When the Desktop Agent is uninstalled from the desktop computer, an option is provided to delete the desktop user data folder.

- 7 Click **Yes** or **Yes to All** to delete the user.

If you delete a user without first uninstalling the Desktop Agent from the user's desktop, the Desktop Agent on that user's desktop closes automatically.

Moving Desktop Agent users to a new network user data folder

When Desktop Agent users are moved to new network user data folders, the contents of each network user data folder are moved to a new directory. The new directories can be existing Storage Locations or other directories on the network.

When the network user data folder is moved to a UNC location, the permissions on the new location may need to be modified. The local administrator group and the owner of the files must have read and change permissions for the network user data folder. In addition, the Everyone group should be removed.

See "[Manually creating new network user data folders](#)" on page 1697.

When the transfer is complete, each affected Desktop Agent turns off, and then automatically restarts within a 30-minute window.

After the data is successfully moved, the data in the old network user data folders is deleted. Subsequent backups are stored in the new location for each user.

To move one or more Desktop Agent users to a new network user data folder

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Users**.
- 3 Select one or more user to be moved.
- 4 In the task pane, under User Tasks, click **Move network user data folder**.

- 5 Select the location for the new network user data folder.
- 6 Click **Start** to begin the data transfer.

Move User Data Folder options

You can move Desktop Agent users to a new network user data folder.

See [“Moving Desktop Agent users to a new network user data folder”](#) on page 1701.

Table Q-53 Move User Data Folder options

Item	Description
User	Lists the domain and the user name of the selected user or users.
From	Lists the current network user data folder location.
Move the user data folder to an existing Storage Location	Lets you choose an existing Storage Location from the drop-down list. A new network user data folder is created in the new Storage Location for each user who is moved.
Move the contents of the user data folder to an alternative location	Lets you specify a new Storage Location. Type the path in the box provided, or click Browse and navigate to the new location. A new network user data folder is created in the new Storage Location for each user who is moved.

Migrating a desktop user to a new computer

When a desktop user receives a new computer, DLO can be used to migrate user data to the new computer. DLO accomplishes this task by staging a user’s backed up data on the new computer using a restore process. When the user logs in, the data is restored to the same location it occupied on the original computer. The final restoration of data occurs automatically when the user logs in and does not require a connection to the media server.

To migrate a desktop user to a new computer

- 1 Restore the user data.
See [“Restoring files and folders from the DLO Administration Console”](#) on page 1707.
- 2 Select the option **Stage this user data on an alternate computer for a new DLO installation.**

Viewing a list of Desktop Agent users

You can view a list of the users who are configured to use the Desktop Agent.

To view a list of Desktop Agent users

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Users** to list users in the results pane.

Modifying computer properties

Computer properties can be viewed and modified from the DLO Administration Console. Computer properties are based on the profile to which the desktop computer owner is assigned. The desktop user can change the computer properties if that user has sufficient rights in the profile.

To view and modify computer properties

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Computers**.
- 3 Right-click the computer for which you want to modify properties, and click **Properties**.
- 4 Do any of the following:
 - On the Schedule tab, adjust the schedule as needed.
See [“Schedule options”](#) on page 1773.
 - On the Options tab, modify the logging and disk space usage options.
See [“Options for the Desktop Agent ”](#) on page 1776.
 - On the Backup Folders tab, view the backup folders for the computer.
 - On the Backup Selections tab, modify the backup selections for the computer.
See [“Adding a DLO backup selection to a profile”](#) on page 1660.
Profile backup selections are not listed, and can only be modified directly in the profile.
See [“Modifying a DLO backup selection”](#) on page 1670.
 - On the Synchronized Selections tab, view synchronized selections for a computer.
Synchronized selections can only be viewed from the Administration Console. They are configured on the Desktop Agent.
See [“About synchronizing desktop user data”](#) on page 1778.
 - On the Connection Policies tab, view and modify connection policies.

Profile defined connection policies can only be modified in the profile.
See “[Customizing connection policies](#)” on page 1777.

Enabling or disabling a desktop computer

When a computer is disabled, the Desktop Agent remains on the desktop computer. The Desktop Agent can be used to restore files and view history, but backups are disabled and the user cannot modify Desktop Agent settings.

To enable or disable a desktop computer

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Computers**.
- 3 In the results pane, select one or more computers to be enabled or disabled.
- 4 Right-click the selected computers, and then do one of the following:
 - Click **Enable** to enable the Desktop Agent to run on the selected computers.
 - Click **Disable** to prevent the Desktop Agent from running on the selected computers.

Deleting a desktop computer from DLO

Deleting a desktop computer from DLO removes the computer from the DLO database and deletes the backed up files. This feature is most commonly used for a desktop computer that is no longer in use. Deleting a computer does not disable the Desktop Agent software. If subsequent backups are performed by the Desktop Agent, the computer entry is added back to DLO. To prevent further backups from the computer, disable the computer rather than deleting it.

To delete a desktop computer from DLO

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Computers**.
- 3 In the results pane, select one or more computers to be deleted.
- 4 In the task pane, under General Tasks, click **Delete**.
- 5 When asked if you want to delete each selected computer and all backup files, click **Yes**.

Backing up a desktop from the DLO Administration Console

The DLO Administration Console can be used to run an immediate backup on one or more desktop computers. This practice allows the administrator to force a backup of a computer running in manual or scheduled mode.

To run an immediate backup on a desktop computer

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the selection pane, click **Computers**.
- 3 In the results pane, select one or more computers on which to run an immediate backup.
- 4 In the task pane, under Computer Tasks, click **Run backup now**.

Setting blackout windows

DLO can be configured to stop backups at specific times to selected file servers, or to file servers that are managed by a specific maintenance server. This feature is called a blackout window. When a blackout window is configured for a selected resource, backups to network user data folders are suspended during the specified period.

Blackout windows are specific to the resource for which they are created. To use the same schedule for two different resources, you must configure them separately.

To configure a blackout window for a network resource

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the task pane, under Tool Tasks, click **Blackout windows**.
- 3 From the File Server list, select a network resource for which you want to configure a blackout window.
- 4 Do one of the following:
 - To edit an existing schedule, select it from the drop-down menu.
 - To create a new schedule click **New**.
- 5 In **Occurs**, indicate whether you want the blackout window to occur weekly or on a specific date.
- 6 Enter the starting time and day and the ending time and day.
- 7 Click **OK**.

Blackout Window options

DLO can be configured to stop backups at specific times to selected file servers, or to file servers that are managed by a specific maintenance server.

See [“Setting blackout windows”](#) on page 1705.

Table Q-54 **Blackout Window** options

Item	Description
File Servers	Lists the file servers to which you can apply a blackout window.
Schedules	Lists the existing blackout windows.
New	Lets you create a new blackout window.
Delete	Lets you remove the blackout window that appears in Schedules .
Enable Schedule	Activates the schedule that appears in Schedules ..
Occurs	Indicates whether the blackout window occurs weekly or on a specific date.
Starts at	Indicates the start time for the blackout window. For a blackout window on a specific date, enter the date on which the blackout window is to start. For a weekly blackout window, select the day of the week on which the blackout window is to start.
Ends at	Indicates the end time for the blackout window. For a blackout window on a specific date, enter the date on which the blackout window is to end. For a weekly blackout window, select the day of the week on which the blackout window is to end.

Deleting a blackout window schedule

You can delete a schedule for a blackout window.

To delete a blackout window schedule

- 1 On the DLO navigation bar, click **Setup**.
- 2 In the task pane, under Tool Tasks, click **Blackout windows**.
- 3 Under Schedules, select the schedule to be deleted.

- 4 Click **Delete**.
- 5 Click **OK**.

Restoring files and folders from the DLO Administration Console

The administrator can restore files and folders to a desktop computer from the DLO Administration Console.

DLO does not restore a file to its original location if the file is in use by another application.

If DLO encounters a file that is in use, you can do one of the following tasks to restore the file:

- Schedule a time to restore the file. The file is restored after the computer restarts. You are not notified when the file is restored.
- Log on to the desktop computer with an administrative account. After you log on with an administrative account, run a restore job to overwrite the locked file and restore it.
- Close the file in the other application.
- Restore the file to an alternate location.

To restore files and folders from the DLO Administration Console

- 1 On the DLO navigation bar, click **Restore**.
- 2 In the Computer pane, click the desktop from which the data to be restored originated.
- 3 In the Backup Folder pane tree view, select the folder containing the files you want to restore.
- 4 To restore the entire folder, check the folder in the Backup Folder pane.
- 5 To restore specific files, check the files in the File Version pane.
- 6 If multiple versions exist for a file, select the file version you want to restore.

When a desktop user deletes an original file, the backup files are retained until the file grooming process deletes them. If an original file has been deleted, but backup files are still available, the icon for the file includes a small red 'x'.

See [“About file grooming in DLO”](#) on page 1664.

- 7 In the task pane, under Restore Tasks, click **Restore files** to open the Restore dialog.

- 8 Select the appropriate options, and then click **OK**.
See “[Restore options](#)” on page 1708.
- 9 Click **OK**.
If you customize NTFS permissions or directory attributes, such as compression or encryption for files or folders, you must reapply these settings after restoration. If you use a password for your PST file, you must reset the password after restoring your PST file.
- 10 In the Restore Summary dialog box, review the selected restore settings, and do one of the following:
 - Click **Print** to print a copy of the restore summary.
 - Click **Restore** to continue with the restore.
- 11 Click **OK** when the restore job completes.

Restore options

The administrator can restore files and folders to a desktop computer from the DLO Administration Console.

See “[Restoring files and folders from the DLO Administration Console](#)” on page 1707.

Table Q-55 **Restore options**

Item	Description
Restore to original computer	Restores the selected files or folders to the computer from which they were originally backed up. When files or folders are restored to the original desktop computer, the job is submitted to the Desktop Agent. The job runs when the Desktop Agent connects to the media server. The job may run immediately if the desktop computer is currently on the network. The job may be pending for some time if the desktop computer is not connected to the network.
Restore to original folder	Restores the file or folder to its original location.
Redirect the restore to an alternate folder	Restores the file or folder to a different location on the original desktop. You can browse to the folder where you want to restore the file. You can browse only if the Windows Firewall is turned off.

Table Q-55 Restore options (continued)

Item	Description
Restore to an alternate computer	Restores data to a computer other than the one from which they were originally backed up. When data is restored to a folder on an alternate computer, DLO processes the restore job immediately from the network user data folder. The job is not queued to the Desktop Agent.
Redirect the restore to a folder on an alternate computer	Restores the data to a selected folder on an alternate computer.
Stage this user data on an alternate computer for a new DLO installation	Migrates user data to a new computer. See “Migrating a desktop user to a new computer” on page 1702.
Preserve folder structure	Restores the data with its original directory structure intact. If you clear this option, all data (including the data in subdirectories) is restored to a single folder in the path you specify.
If file already exists:	Determines what to do when a file that was selected for restore already exists in the destination folder. The following options are available: <ul style="list-style-type: none"> ■ Do not overwrite Cancels the restoration of files that already exist in the destination folder. ■ Prompt Prompts the user before overwriting the file if it already exists in the destination folder. ■ Overwrite Overwrites the file without prompting if it already exists in the destination folder.
Restore deleted files	Restores files even though the original has been deleted.
Preserve security attributes on restored files	Preserves security information in restored files. You may need to uncheck this box to successfully restore a file if the source file security conflicts with the destination security. Unchecking this option causes the security information to be removed from the restored file.

Restore Summary options

The **Restore Summary** dialog box lists the files that you selected to restore.

Table Q-56 **Restore Summary** options

Item	Description
Settings	Lists the settings that you selected for this restore job.
Selections	Lists the files that you selected to restore.
Restore	Starts the restore job.
Cancel	Cancels the restore process. No files are restored.
Print	Prints the summary information on a printer that you select.

Searching for files and folders to restore with DLO

You can use the search feature to find the data that you want to restore.

To search for desktop files and folders to restore

- 1 On the DLO navigation bar, click **Restore**.
- 2 In the Computer pane, click the desktop on which you would like to search for files to restore.
- 3 In the task pane, under Restore Tasks, click **Search for files to restore**.
- 4 Select the appropriate options.
See [“Restore search options”](#) on page 1711.
- 5 Click **Search**.
- 6 In the results pane, check the items to be restored.
In some cases the Restore Search view may contain duplicate entries for the same file. You can select either file to restore and receive the same outcome.
- 7 Click **Restore**.
- 8 Select the appropriate options.
See [“Restoring files and folders from the DLO Administration Console”](#) on page 1707.
- 9 Click **OK**.

Restore search options

You can use the search feature to find the data that you want to restore.

See “[Searching for files and folders to restore with DLO](#)” on page 1710.

Table Q-57 Restore search options

Item	Description
Search for file names with this text in the file name	Indicates the name of the file or folder you want to find. Wildcard entries are accepted, for example *proj.doc.
Modified	Lets you search for the files that were modified during a specific time frame.
Today	Lets you search for the files that were modified on the current calendar day.
Within the past week	Lets you search for the files that were modified in the last calendar week.
Between	Lets you search for the files that were modified during a range of days.
Of the following type	Lets you search for a file type that is in the list.
Of the following size	Select this check box and then enter information as follows: <ul style="list-style-type: none"> ■ Select from equal to, at least, or at most in the first drop-down menu. ■ Type a file size. ■ Select KB, MB, or GB.

About DLO emergency restore and recovery passwords

DLO's Emergency Restore feature is used to recover Desktop Agent user data from the File Server in the event that the configuration database is lost. Emergency Restore can also simplify the task of restoring user data for the users that have been deleted from the DLO Administration Console. To use the Emergency Restore feature, a Recovery Password must have been established before the database was lost or the user was deleted. If user data is restored from another media, you must use the Recovery Password that was in effect when the user data was backed up to recover the data.

A Recovery Password is established when the DLO Administration Console is first launched. For older versions of DLO, a recovery password had to be manually established using the DLO command line interface. The recovery password is used to encrypt each user's encryption key so the key can safely be stored on the File Server. DLO encrypts user data using a user-specific, randomly generated encryption key. The encryption keys are stored in DLO's configuration database on the media server. The encryption-keys are also stored, in encrypted form, on the File Server.

The Emergency Restore feature prompts the administrator for the Recovery Password, which is used to decrypt the user's encryption key. The encryption key is then used to decrypt the user's data. If a recovery password has not been established the Emergency Restore feature cannot be used to restore encrypted user data.

About changing recovery passwords

If the Recovery Password must be changed, the administrator must be aware that the former Recovery Password will still be in effect for former backups of the File Server.

The Recovery Password should only be changed if mandated for security reasons, such as a compromised password. If possible the Recovery Password should never be changed. Changing or establishing a Recovery Password will never aide in restoring existing user data. In fact, it can make it more difficult: changing the Recovery Password can result in multiple Recovery Passwords being in use at the same time.

For example, consider the case where a recovery password "pwd1" is established when DLO is installed. Each user's encryption-key is encrypted with the Recovery Password stored on the File Server. When the File Server is backed up, the backup copies all use the Recovery Password "pwd1". If the recovery password is subsequently changed to "pwd2", the user encryption keys on the File Server will be changed to be encrypted with the new Recovery Password. Subsequent backups of the File Server will use the Recovery Password "pwd2". Now there are backups of the File Server using both "pwd1" and "pwd2" as the Recovery Password. When the Emergency Restore feature is used, the administrator must use the Recovery Password that was in effect when the File Server was backed up.

What happens when a user is deleted by the DLO Administration Console

When a user is deleted using the DLO Administration Console, all data that is associated with the user is deleted. This data includes the configuration data, which is stored on the media server, and the user data, which is stored on the File

Server. The method for restoring data for a deleted user depends upon whether a Recovery Password has been established or not.

Recovering data for a single user by using DLO Emergency Restore

The Emergency Restore feature can be used to restore data for a deleted user if the following conditions are met:

- The user data can be restored from a backup of the File Server.
- A Recovery Password was established before making the backup.

See [“About DLO emergency restore and recovery passwords”](#) on page 1711.

To recover data for a single user using DLO Emergency Restore

- 1 Restore the user-data to its original location on the File Server or to any other temporary location.

- 2 Use the DLO Command Line Interface to restore the data to DLO.

```
dlocommandu -emergencyrestore <usersharepath> -w <RecoveryPassword>  
-ap <destination-path>
```

Recovering data for a single user without using DLO Emergency Restore

If the Recovery Password was not established or has been lost, you must restore the media server and the file server to a single point in time before the user was deleted. Then, you can restore the data for the deleted user.

To recover data from a single user without using DLO Emergency Restore

- 1 Take both the File Server and media server offline

- 2 Back up both the File Server and media server.

Ensure that the backup includes the DLO configuration database and the all user data. This backup is used to restore DLO back to its current state after the data is recovered. If any DLO data is not backed up it may be impossible to return to the current state.

- 3 Restore the user data to the File Server.

If possible, restore just the data for the user being restored. If unsure, the entire volume on the File Server can be restored, provided that precaution was taken in step 2 to ensure that the entire volume was backed up.

- 4 Restore the configuration database to the media server.

The default database path is C:\Program Files\Symantec\Backup Exec\Data.

- 5 Restart the media server.

- 6 Use the DLO Administration Console to restore the user's data. Select "Restore to an alternate computer" and restore the data to a temporary location.
- 7 Restore both the File Server and media server back to the most recent state.

Recovering a media server or a file server if a non-system disk fails or is otherwise corrupted

You can recover a damaged or corrupt media server or file server. The media server stores the configuration database. The file server stores the user data.

Table Q-58 How to recover a media server or a file server

Step	Description
Step 1	Fix or replace the failed disk.
Step 2	Restore the entire disk from the backup copy.
Step 3	Restart the computer.

Recovering a media server if the hard drive fails or the computer needs to be replaced

You can recover a damaged or corrupt media server. The media server stores the configuration database.

To recover a media server if the hard drive fails or the media server computer needs to be replaced with a new computer

- 1 Set up the computer with the operating system software. Be sure to use the same computer name as the failed media server.
- 2 Install DLO on the new media server. Be sure to use the same version of DLO as was installed on the failed media server.
- 3 Restore the DLO database files, overwriting the database files created when DLO was installed. The default database path is C:\Program Files\Symantec\Backup Exec\Data.
- 4 Restart the computer

Recovering a file server if the hard drive fails or the computer needs to be replaced

You can recover a damaged or corrupt file server. The File Server stores the user data.

To recover a file server if the hard drive fails, or the file server computer needs to be replaced with a new computer

- 1 Set up the computer with the operating system software. Be sure to use the same computer name as the failed File Server.
- 2 If the File Server had the DLO Maintenance Server installed, then install the DLO Maintenance Server on the computer. Be sure to use the same version of DLO as was installed on the failed File Server.
- 3 Restore the DLO file data.

Computer History pane options and Job History pane options

Use the History view on the DLO Administration Console to view information about the status of Desktop Agent jobs. History logs are generated by each desktop that runs the Desktop Agent. The History view includes a computer history and a job history for each desktop.

You can view history logs in either the DLO Administration Console or the Desktop Agent Console. You can filter history logs so that old or less important logs do not appear, or so that only backup or restore job logs appear.

By default, the history logs are updated when a job runs and an hour has passed since the last update. However, if the job's status changes, the history log is updated immediately to reflect the new status.

The History view provides the following summary information:

Table Q-59 Computer History pane

Item	Description
User	The user name of the user who is logged on to the desktop that generated the message.
Computer	The name of the desktop that generated this message.
Last Backup Result	The outcome of a completed backup, for example, Success, Warnings, Failed, Canceled.
Profile	The name of the Profile to which the desktop user who is logged on to the desktop belongs. See “About DLO profiles” on page 1641.

Table Q-59 **Computer History pane** (*continued*)

Item	Description
Backup Mode	The backup mode that is specified in the profile. Backup modes include the following: <ul style="list-style-type: none"> ■ Continuous. The backup occurs whenever a file changes. ■ Scheduled. The backup occurs according to a schedule. ■ Manual. The backup occurs when initiated by the desktop user.
Desktop Data Folder Size	The current size of the desktop user data folder.
Network Data Folder Size	The current size of the network user data folder.
Network Data Folder Path	The location of the network user data folder.

The Job History pane displays the following information:

Table Q-60 **Job History pane**

Item	Description
Start Time	The time the job was started.
End Time	The time the job ended.
Operation	The operation that is performed in this job, such as backup or restore.
Status	The current status of the job, such as active, completed, completed with errors, completed with warnings, canceled, or failed.
Files Protected (Desktop)	The number of files that are copied to the desktop user data folder during the job.
Size Protected (Desktop)	The total bytes of data that are copied to the desktop user data folder during the job.
Files Protected (Network)	The number of files that are copied to the network user data folder during the job.
Size Protected (Network)	The total bytes of data that are copied to the network user data folder during the job.
Errors	The number of errors, if any, that were generated during the job.

Viewing history logs

History logs are listed for each job on a desktop computer.

To view a history log in the DLO Administration Console

- 1 On the DLO navigation bar, click **History**.
- 2 In the History pane, select the computer for which you want to view a history log.
- 3 In the Job History pane, click the log you want to view.
- 4 In the task pane, under General Tasks, click **View history log** file to display the log file viewer with all log messages for this job.
- 5 To filter the results, select the appropriate options.
See [“Log File Viewer options”](#) on page 1717.
- 6 Click **Search**.
- 7 Double click a log entry to view additional details.
- 8 Click **Close**.

Log File Viewer options

You can view the log file for each job that runs on a computer.

See [“Viewing history logs”](#) on page 1717.

Table Q-61 Log File Viewer options

Item	Description
All log files	Shows all log entries in the log file viewer.
Current log file	Searches only the log entries that are in the current log file.
With timestamp	Searches only those log entries within a specified time period. The following options are available: <ul style="list-style-type: none"> ■ Today - Show only log files that were created today. ■ Within the last week - Show all log files created in the last week. ■ Between dates - Show all log files created between the dates entered.

Table Q-61 Log File Viewer options (*continued*)

Item	Description
Of the following type	<p>Shows only logs of the indicated type.</p> <p>The available selections vary depending on the log file, but may include the following:</p> <ul style="list-style-type: none"> ■ Backup ■ Restore ■ Move User ■ Maintenance
With File names containing	<p>Searches for files by file name or file name type. Wildcard entries are supported.</p> <p>Example: *gold.doc</p> <p>When you use wildcards, you must use the '*' wildcard. For example, *.tmp returns all results with the .tmp extension while .tmp returns only files explicitly named .tmp.</p>
Limit search to	<p>Limits the log files displayed to one of the following types of log entries:</p> <ul style="list-style-type: none"> ■ Informational entries only ■ Error and warning entries only ■ Error entries only ■ Warning entries only ■ Local data folder entries only ■ Local data folder error entries only ■ Network data folder entries only ■ Network data folder error entries only

Setting filters for the job history view

The job history view can be filtered to show only the type of jobs you want to view. You can filter jobs by type, alerts that are received during the job, or by the time period in which the job was run.

To set filters for the job history view

- 1 On the DLO navigation bar, click **History**.
- 2 Click the desktop for which you want to view the history.

- 3 In the task pane, under Job History View Filters, click one of the following:

List all jobs	Lists the history logs for all jobs that have run on the selected desktop. These may include backup, synchronization, restore, or move user jobs.
List backup jobs only	Lists the history logs only for the backup jobs that have run on the selected desktop.
List restore jobs only	Lists the history logs only for the restore jobs that have run on the selected desktop.

- 4 Filter job history logs based on alerts that are received by selecting one or more of the following:

Show successful jobs	Lists the history logs for all successful jobs on the selected desktop.
Show jobs with warnings	Lists the history logs for all jobs that generated warnings on the selected desktop.
Show jobs with errors	Lists the history logs for all jobs that generated errors on the selected desktop.
Show canceled jobs	Lists the history logs for all jobs that were canceled on the selected desktop.

- 5 Select a time frame for filters to be displayed by selecting one of the following:

Show last 24 hours	Lists the history logs that have been generated in the last 24 hours, and that meet all other filtering criteria.
Show last 7 days	Lists the history logs that have been generated in the last 7 days, and that meet all other filtering criteria.
Show all	Lists all history logs that also meet all other filtering criteria.

Searching history logs

You can use the Log File Viewer to refine the list of jobs to only those of interest.

To search log files

- 1 On the DLO navigation bar, click **History**.
- 2 In the task pane, under General Tasks, click **Search log files** to display the log file viewer.
- 3 Set filtering options.
See "[Log File Viewer options](#)" on page 1717.
- 4 Click **Search**.
- 5 Double click on a log entry to view additional details.
- 6 Click **Close**.

About monitoring alerts on the DLO Administration Console

Alerts appear in DLO when the system needs administrator attention. Alerts help the DLO administrator understand the current condition of DLO jobs by displaying information on jobs.

Alerts can be generated to provide general information, or they can be in response to a problem. When an alert is generated due to a problem, the alert contains information about the problem. It may also include recommendations on how to fix the problem.

The DLO Administrator can choose to display all alerts, or to limit the type of alerts that appear.

Active alerts display the alerts that are active in the system and need a response from the operator. Alert history displays the alerts that have been responded to or the alerts that have been automatically cleared from the system.

In addition, the status bar at the bottom of the screen displays an alert icon. The icon that displays in the status bar is for the most severe type of alert in the Active alerts list. Therefore, if the current or most recent alert is not the most severe, the icon in the status bar does not match the icon for the most recent alert in the alert list.

Alerts are filtered by the Desktop Agent to minimize the load on DLO. By default, alerts are limited to one of each type in 24 hours. For example, you will see only one "Local Out of Disk Condition" alert in a 24-hour period from a desktop running the Desktop Agent.

Note: "Backup/Restore complete" alerts cannot be filtered. If you enable these alerts, they are generated each time a backup or restore job completes.

Active alerts that are older than a specified number of days are cleared and moved into the alert history. The alerts in the history are deleted if they have been cleared for more than a specified number of days. When alerts in the history have been cleared for a given number of days, which by default is seven days, they are deleted by a Backup Exec full backup job that backs up and delete files.

If an alert is manually cleared, it is moved into the alert history. Deleting an alert manually removes it permanently.

You can set up DLO to notify recipients when alerts occur.

Alert categories

The following tables lists DLO alert categories.

Table Q-62 Alert categories

Alert Type	Description
Informational	Notifies you that an expected action has occurred, such as the successful completion of a backup or restore job.
Warning	Notifies you of a potential issue. For example, an alert is generated when a backup has not been completed on a desktop within a given time frame, or if the disk quota limitations are being approached.
Error	Notifies you of an active or pending danger to the application or its data. An error would be generated, for example, if a backup failed to complete, or if a desktop has exceeded its disk quota limitations.

DLO informational alerts

The following table lists the types of informational alerts in DLO.

Table Q-63 Types of DLO informational alerts

Alert	Description
A backup job has completed	A backup job has completed successfully.

Table Q-63 Types of DLO informational alerts (*continued*)

Alert	Description
A restore job has been queued	A restore job was initiated from the media server.
A restore job has completed	A restore job has completed successfully.
PST file was skipped because it is not configured in Outlook	A PST file on the desktop computer was not backed up because it was not configured in Microsoft Outlook.
User was configured	A new user connected and was successfully configured.
A backup job has completed	A backup job has completed successfully.
A restore job has been queued	A restore job was initiated from the media server.
A restore job has completed	A restore job has completed successfully.
PST file was skipped because it is not configured in Outlook	A PST file on the desktop computer was not backed up because it was not configured in Microsoft Outlook.

DLO warnings

The following table lists the types of warnings in DLO.

Table Q-64 DLO warnings

Alert	Description
A backup job has completed with warnings	A backup job has completed, but warnings were generated.
A restore job has completed with warnings	A restore job has completed, but warnings were generated.
A restore job has not completed in 1 hour	A restore job was submitted, but an hour has passed and the restore job is not complete.
A restore job has not completed in 12 hours	A restore job was submitted, but 12 hours have passed and the restore job is not complete.

Table Q-64 DLO warnings (*continued*)

Alert	Description
A restore job has not completed in 24 hours	A restore job was submitted, but 24 hours have passed and the restore job is not complete.
Desktop user data folder approaching storage limit	The amount of stored backup data in a user's desktop user data folder is approaching the specified size limit.
Desktop user data folder disk space low	The volume containing the desktop user data folder is running low.
Evaluation period daily reminder	This reminder specifies the number of days remaining in the evaluation period for the Symantec Desktop and Laptop Option
Evaluation period has expired	The DLO evaluation period has expired. A license is required to continue to use DLO.
Network user data folder approaching storage limit	The amount of stored backup data in a user's network user data folder is approaching the specified size limit.
Network user data folder disk space low	A user's network user data folder is almost out of disk space.
No backups in 14 days	A user's data has not been backed up in 14 days. If a user works on multiple computers, this warning appears for each computer that the user accesses.
No backups in 28 days	A user's data has not been backed up in 28 days. If a user works on multiple computers, this warning appears for each computer that the user accesses.
No backups in 7 days	A user's data has not been backed up in 7 days. If a user works on multiple computers, this warning appears for each computer that the user accesses.
No matching automated user assignment	An automated user assignment does not exist with matching criteria for this user. Create a new automated user assignment, or edit an existing one and add criteria for this user.

DLO alerts

The following table lists the types of alerts in DLO.

Table Q-65 DLO alerts

Alert	Description
A backup job has completed with errors	A backup job has completed, but errors were generated.
A restore job has completed with errors	A restore job has completed, but errors were generated.
Desktop user data folder disk space full	The volume containing the desktop user data folder is full. There is insufficient free disk space to back up the current file. The file is copied directly to the network user data folder.
Desktop user data folder storage limit has been reached	The specified disk storage limit was reached when trying to add a new revision to the desktop user data folder.
Filename, directory name, or volume label syntax is incorrect.	Indicates either a storage system problem that requires attention, or a file name that was denied by SRM software. If the latter, these files should be added to DLO's global exclude list. See “About configuring global exclude filters in DLO” on page 1687.
Network user data folder disk space full	The volume containing the network user data folder is full. There is insufficient free disk space to back up the current file.
Network user data folder storage limit has been reached	The specified disk storage limit was reached when trying to add a new revision to the network user data folder.
Unable to configure the Desktop Agent	A new user has connected, but for an unknown reason, cannot be configured properly.

Configuring alerts

You can select the types of alerts that you want to receive. In addition, you can enable recipients for the alerts.

To configure alerts

- 1 On the DLO navigation bar, click **Alerts**.
- 2 In the task pane, under Alert Tasks, click **Configure alerts**.
- 3 Check the alerts you want to receive, and clear the check boxes for the alerts you do not want to receive.
- 4 To send notification to recipients when the selected alerts are generated, do the following:

- Select one or more alerts from the list. To select multiple alerts, click one item and press <Ctrl> or <Shift> while clicking the other items.
- Check the **Send notification of selected alert to recipients** check box.
- Select the recipients to receive notification of the alerts.

Alerts must be configured for notification before selecting recipients.

See [“About configuring recipients for notification in DLO”](#) on page 1728.

5 Click **OK**.

Configure Alerts options

You can select the types of alerts that you want to receive. In addition, you can enable recipients for the alerts.

See [“Configuring alerts”](#) on page 1724.

Table Q-66 **Configure Alerts options**

Item	Description
Alert categories	Lists all of the alerts that you can choose to receive.
Recipients	Lists all of the recipients that you can select to receive alerts..
Send notification of selected alert to recipients	Enables DLO to sent the selected alerts to the selected recipients.
New	Lets you set up a new recipient.
Remove	Lets you delete the selected recipient from the list.
Properties	Lets you view or change the properties for the selected recipient.

Managing DLO alerts

From the Alerts view in the DLO Administration Console, you can view a subset of alerts, clear alerts, and move alerts to a history log.

To view DLO alerts

- 1 On the DLO navigation bar, click **Alerts**.
- 2 Select **Active alerts** to view active alerts, or **Alert history** to view the alerts that have been cleared.

Alerts that are older than a specified number of days are cleared and moved into alert history. The number of days is specified in the Backup Exec Administration Console.

- 3 To filter alerts by type, select one or more options from Active Alerts View Filters or Alert History View Filters:

Show errors	Lists error alerts for the selected view.
Show warnings	Lists warning alerts for the selected view.
Show information	Lists informational alerts for the selected view.

- 4 To view the properties of an alert, right-click the alert in the Active Alerts or Alert History list and select properties.
- 5 If a log file is associated with the alert, a link is provided to the log file. Click this link to view the log file.
- 6 Click **Close** to close the Alert Information dialog.

Clearing DLO alerts

By default, alerts move to the alert history after a specified period of time. However, some alerts may appear frequently and fill the Active alerts pane. You may want to clear these alerts to the Alert history pane before they are moved automatically.

To clear DLO alerts

- 1 On the DLO navigation bar, click **Alerts**.
- 2 Filter the Alerts view.
See [“Managing DLO alerts”](#) on page 1725.
- 3 From the alert list, select one or more alerts that you want to clear.
- 4 In the task pane, under Alert Tasks, do one of the following:
 - Select **Respond** to clear only the selected alerts.
 - Select **Respond OK to all** to change the status of all alerts to cleared.

Alert Information options

You can view information about alerts and respond to them.

Table Q-67 Alert Information options

Item	Description
Operation	Lists the type of operation to which the alert applies.
Respond	Lets you clear the alert.
Computer	Lists the name of the computer to which the alert applies.
User	Lists the user name of the user who was logged on at the time of the alert.
Time	Lists the time when the alert occurred.

About configuring notification methods for DLO alerts

DLO has several methods to notify you of alerts.

You can choose any of the following methods:

- SMTP
You must have an SMTP-compliant email system, such as a POP3 mail server to receive alert notification messages using the SMTP notification method.
- MAPI
You must have a MAPI-compliant email system, such as Microsoft Exchange to receive alert notification messages using the MAPI notification method.
- Lotus Notes (VIM) email
You must have a VIM (Lotus Notes) compliant email system to receive alert notification messages using the VIM notification method.
- Pagers
You must have a modem set up on your system to use the pager notification method. You must be sure that the modem can communicate properly with your paging service in order for pager notification to work properly. Before you set up pager notification, contact your paging service for information about the recommended brand of modem to use with your paging service.
- Printers
- Net Send

To use notifications you must perform the following:

- Configure the methods you want to use to notify the recipient. Printer and Net Send notification methods do not require pre-configuration.
- Configure recipients. Recipients are individuals, computer consoles, printers, or groups. They can be configured to use one or more of the notification methods.
- Assign the recipients to alerts or jobs for notification.

See “[Configuring notification methods for DLO alerts](#)” on page 1728.

Configuring notification methods for DLO alerts

To configure a notification method for DLO alerts

- 1 On the Tools menu, click **Email and Pager Notification**.
- 2 Select the tab for the notification method that you want to configure, and then complete the options for that method.

The following methods are available:

- SMTP
- MAPI

If you install Outlook after you install DLO, you must stop and restart the DLO Administration Service.

- Lotus Notes (VIM) email
- Pagers
- Printers
- Net Send

- 3 Click **OK**.

About configuring recipients for notification in DLO

Recipients are individuals with a predefined notification method, computer consoles, printers, or groups. Recipient configuration consists of selecting a notification method and defining notification limits. After you create entries for the recipients, you can assign them to alerts or jobs.

The following types of recipients can be configured for notifications:

Table Q-68 Recipient types

Type	Description
Person	An individual that has a predefined method of notification such as SMTP, MAPI, or VIM email, or a pager. You must configure the notification method before you can enable it for the recipient.
SNMP Trap	SNMP Traps are sent to a computer that is configured to receive them.
Net Send	A computer that serves as a notification recipient.
Printer	A specific printer to which notifications can be sent.
Group	A group of one or more recipients, including person recipients, Net Send recipients, and other groups.

Enabling a person to receive DLO alert notifications by SMTP mail

You can configure a person to receive SMTP email notification messages if you have configured the SMTP notification method.

To enable a person to receive alert notifications by SMTP mail

- 1 On the Tools menu, click **Recipients**.
- 2 Click **New**.
- 3 Click **Person**.
- 4 Click **OK**.
- 5 In the Name field, type the name of the recipient that you want to configure.
- 6 On the SMTP Mail tab, select the appropriate options.
- 7 Click **OK**.

Enabling a person to receive DLO alert notifications by MAPI mail

You can configure a person recipient to receive MAPI email notification messages if you have configured the MAPI notification method.

To enable a person to receive alert notifications by MAPI mail

- 1 On the Tools menu, click **Recipients**.
- 2 Click **New**.
- 3 Click **Person**.
- 4 Click **OK**.

- 5 In the Name field, type the name of the recipient that you want to configure.
- 6 On the MAPI Mail tab, select the appropriate options.
- 7 Click **OK**.

Enabling a person to receive DLO alert notifications by VIM mail

You can configure a person recipient to receive VIM email notification messages if you have configured the VIM notification method.

To enable a person to receive alert notifications by VIM mail

- 1 On the Tools menu, click **Recipients**.
- 2 Click **New**.
- 3 Click **Person**.
- 4 Click **OK**.
- 5 In the Name field, type the name of the recipient that you want to configure.
- 6 On the VIM Mail tab, select the appropriate options.

Enabling a person to receive DLO alert notifications by pager

You can configure a person recipient to receive notification messages by pager if you have configured the pager notification method.

To enable a person to receive alert notifications by pager

- 1 On the Tools menu, click **Recipients**.
- 2 Click **New**, and then click **Person**.
- 3 Click **OK**.
- 4 In the Name field, type the name of the recipient that you want to configure.
- 5 On the Pager tab, select the appropriate options.
- 6 Click **Advanced** to configure advanced pager setup options and select the appropriate options.
- 7 Click **OK** to save the settings in the Advanced Pager Information dialog box, and then click **OK** to save the pager configuration settings.

Enabling SNMP Trap to receive DLO alert notifications

You can configure an SNMP trap to receive notification messages.

To enable SNMP Trap to receive alert notifications

- 1 On the Tools menu, click **Recipients**.
- 2 Click **New**.
- 3 Click **SNMP Trap**.
- 4 Click **OK**.
- 5 Select the appropriate options.
- 6 Click **OK**.
- 7 Click **Close**.

Enabling Net Send to receive DLO alert notifications

You can configure Net Send to send notification messages to a target computer or user.

Note: If the target computer has Internet pop-up advertisement blocking software installed, the Net Send notification message does not appear.

To enable Net Send to receive alert notifications

- 1 On the Tools menu, click **Recipients**.
- 2 Click **New**, and then click **Net Send**.
- 3 Click **OK**.
- 4 Select the appropriate options.
- 5 Click **OK**.

Enabling a printer to receive DLO alert notifications

You can select installed printers as a notification method for recipients; however, DLO does not support fax printer devices. Only the printers that were configured using the same user name and password as the DLO service account can be selected.

To enable a printer to receive alert notifications

- 1 On the Tools menu, click **Recipients**.
- 2 Click **New**, and then click **Printer**.
- 3 Click **OK**.
- 4 Select the appropriate options.

Enabling a group to receive DLO alert notifications

Groups are configured by adding recipients as group members. A group contains one or more recipients and each recipient receives the notification message. Members of the group can be a combination of individual persons, computers, printers, or other groups.

To enable a group to receive alert notifications

- 1 On the Tools menu, click **Recipients**.
- 2 Click **New**, and then click **Group**.
- 3 Click **OK**.
- 4 In the Group Name field, type the name of the group for whom you want to configure the notification.
- 5 Do one of the following:

To add members to the group Select recipients from the All Recipients list, and then click **Add** to move them to the Group Members list.

To remove members from the group Select recipients from the Group Members list, and then click **Remove** to move them to the All Recipients list.

- 6 Click **OK**.

Scheduling notification for recipients in DLO

You can select the times of the day and the days of the week the recipient is available to receive the notification messages. You can modify the schedule after the recipient is configured by editing recipient notification properties.

See [“About configuring recipients for notification in DLO”](#) on page 1728.

To schedule notification for recipients

- 1 On the Recipient Properties dialog box, under the Limit when notifications can be sent group box, click **Enable** to activate the option.
You can access the Recipient Properties dialog box from the Tools > Recipients menu.
- 2 Click **Schedule**.

3 Do any of the following:

Include work days	Clear the Include work days check box to exclude Monday through Friday from 8 A.M. to 6 P.M.
Include weeknights	Clear the Include weeknights check box to exclude Monday through Friday from 6 P.M. to 8 A.M.
Include weekends	Clear the Include weekends check box to exclude Saturday and Sunday, 24 hours a day.

You can select any combination of Include work days, Include weeknights, or Include weekends, or click any single hour of the chart to select or clear that hour.

4 Click **OK**.

Changing information about a recipient in DLO

You can edit the recipient notification properties at any time and change the recipient information, such as an email address, telephone number, or schedule.

You can edit any of the properties except for the recipient name in the Name field. To modify the recipient name, you must create a new recipient, and then delete the old one.

To change information about a recipient

- 1 On the Tools menu, click **Recipients**.
- 2 Select the recipient you want to edit.
- 3 Click **Properties**.
- 4 Edit the properties for the selected recipient.
- 5 Click **OK**.

Changing the notification method for a recipient in DLO

You can configure new notification methods or edit existing notification methods after you configure recipients.

To change the notification method for a recipient

- 1 On the Tools menu, click **Recipients**.
- 2 Select the recipient to be edited and click **Properties**.
- 3 Edit notification properties for the following types of notification methods:
 - SMTP Configuration.

See “[Enabling SNMP Trap to receive DLO alert notifications](#)” on page 1730.

- MAPI Configuration.
See “[Enabling a person to receive DLO alert notifications by MAPI mail](#)” on page 1729.
- VIM Configuration.
See “[Enabling a person to receive DLO alert notifications by VIM mail](#)” on page 1730.
- Pager Configuration. Click **Enable** to activate or clear the notification method, and then select a modem from the Configured Modems list.

4 Click **OK**.

Removing recipients for DLO alerts

You can delete the recipients that do not want to receive notification messages; however, the recipient is permanently removed upon deletion. If you want to keep the recipient, but do not want the recipient to receive notifications, clear the Enable check box in the recipient properties.

To remove a recipient

- 1 On the Tools menu, click **Recipients**.
- 2 Select the recipient you want to delete, and then click **Remove**.
- 3 Click **OK**.
- 4 You can start the job after configuring the new recipients or edit recipient properties or select other options from the Properties pane.

About DLO reports

DLO provides a variety of reports that show detailed information about your DLO operations. When you generate a report, you can specify filter parameters or a time range for the data that you want to include in the report. If Adobe Acrobat is detected, reports are displayed in Adobe Portable Document Format (PDF). If Adobe Acrobat is not detected, the reports are displayed using HTML. Both PDF and HTML reports can be saved and printed.

The following reports are available on the Reports view:

Table Q-69 DLO reports

Report Name	Description
Active Alerts	A list of all currently active alerts. The alerts are arranged chronologically.
Active Alerts by Computer	A list of all currently active alerts. The alerts are sorted by computer name.
Active Alerts by User	A list of all currently active alerts from all computers. The alerts are sorted alphabetically by Desktop Agent user name.
Alert History	A chronological list of alerts that have been sent by all computers in the past.
Alert History by Computer	A list of alerts that have been sent by all computers in the past. The alerts are sorted by computer name.
Alert History by User	A list of alerts that have been sent by all computers in the past. The alerts are sorted by Desktop Agent user name.
Failed Backups	A chronological list of computers that have a failed status for the last backup.
Failed Backup by Computer	A list of computers that have a failed status for the last backup. The alerts are sorted by computer name. Only the last backup result is stored in the DLO database. Therefore, it is possible to report only the last backup result for each desktop computer and not a complete history of failed jobs.
Failed Backup by User	A list of computers that have a failed status for the last backup. The alerts are sorted by Desktop Agent user name. Only the last backup result is stored in the DLO database. Therefore, it is possible to report only the last backup result for each desktop computer and not a complete history of failed jobs.
Last Backup Status	A chronological list of the last backup status for all Desktop Agent computers.
Last Backup Status by Computer	A list of the last backup status for all Desktop Agent computers. The alerts are sorted by computer name.
Last Backup Status by User	A list of the last backup status for all Desktop Agent computers. The alerts are sorted by Desktop Agent user name.

See [“Viewing DLO report properties”](#) on page 1736.

See [“Running a DLO report”](#) on page 1736.

Running a DLO report

When you run a report, you can specify filtering criteria to determine which items are included in the report. After the report is generated, only the items that match the entered criteria appear in the report. If no criteria are entered, all available entries are included in the report.

To run a report

- 1 On the navigation bar, click **Reports**.
- 2 In the Reports pane, select the report you want to run.
- 3 In the task pane, under Reports Tasks, click **Run report now**.
- 4 Select the appropriate parameters for the data you want to include in the report.
 See [“Run Report Now options”](#) on page 1736.
- 5 Click **OK** to run the report. The report can be printed or saved before it is closed.
- 6 Click **OK** to close the report.

Run Report Now options

When you run a report, you can specify filtering criteria to determine which items are included in the report.

See [“Running a DLO report”](#) on page 1736.

Table Q-70 Run Report Now options

Item	Description
Computer	Creates a report for a specific computer. You must enter a desktop computer name.
User	Creates a report for a specific desktop user. You must enter the user’s name.
Days	Creates a report for a specific number of days. You must enter the number of days.

Viewing DLO report properties

Report properties provide a summary of information about each report. The properties can be viewed, but not edited.

See [“Report options”](#) on page 1737.

To view report properties

- 1 On the navigation bar, click **Reports**.
- 2 In the Reports pane, select the report for which you want to view properties.
- 3 In the task pane, under General tasks, click **Properties**.
- 4 After you review the properties, click **OK**.

Report options

Report properties provide a summary of information about each report. The properties can be viewed, but not edited.

See [“Viewing DLO report properties”](#) on page 1736.

Table Q-71 Report options

Item	Description
Title	Shows the name of the report.
Description	Shows the type of data that is included in the report.
Category	Shows the classification for the report. The following report categories are available: <ul style="list-style-type: none"> ■ Alerts ■ Last Backup Status ■ Failed Jobs
Author	Shows the creator of the report.
Subject	Shows the version of the product for which the report was created.
Keywords	Shows the primary information that is used to categorize the report.
File name	Shows the file name of the report template.
File size	Shows the size of the report template.
Creation Date	Shows the date the report was installed on the system.

About maintaining the DLO database

The Desktop and Laptop Option installs its own Microsoft SQL Express 2005 or SQL Server database in the same location as the Backup Exec database. These databases operate independently of one another. If you move the Backup Exec database at a later date, the DLO database remains in its original location.

You can maintain both the Backup Exec database and the DLO database using the Backup Exec database maintenance options. Use BEUtility to perform database operations on BKUPEXCDLO. The DLO database is backed up and restored automatically each time the Backup Exec database is backed up or restored.

If you use BEUtility to repair or recover the DLO database, all DLO Administration Consoles must be closed. Otherwise, the operation fails.

About clustering the Desktop and Laptop Option

To cluster DLO in a Backup Exec cluster configuration, you must install DLO on each cluster node.

You cannot add DLO to an existing Backup Exec cluster. Either add DLO when you configure a cluster, or uncluster an existing cluster, add DLO, and then reconfigure the cluster.

If the Desktop Agent was installed from a cluster node that is now inactive, it does not reconnect to the cluster following the unclustering process.

See [“About installing the Backup Exec Desktop and Laptop Option”](#) on page 1610.

See [“Using Backup Exec with Veritas Cluster Server”](#) on page 842.

See [“Uninstalling Backup Exec from a Microsoft cluster”](#) on page 820.

See [“Installing Backup Exec on a Microsoft Cluster Server ”](#) on page 818.

Installing Backup Exec and the Desktop and Laptop Option to an existing cluster

Follow these steps to install Backup Exec and the Desktop and Laptop Option to an existing cluster.

To install Backup Exec and the Desktop and Laptop Option to an existing cluster

- 1 Install Backup Exec on the active cluster node. Be sure to include the Desktop and Laptop Option during the Backup Exec installation.
- 2 Install Backup Exec on each failover node. Be sure to include the Desktop and Laptop Option during the Backup Exec installation.

Install Backup Exec when the node has access to the shared cluster disk.

- 3 From the server you used in step 1, complete the Cluster Configuration Wizard from the Backup Exec console to configure the cluster-aware Backup Exec media server. Be sure to include all cluster nodes where DLO was installed.

Upgrading an existing Backup Exec 9.x or 10.x cluster that includes DLO

Follow these steps to upgrade an existing Backup Exec 9.x or 10.x cluster that includes DLO.

To upgrade an existing Backup Exec 9.x or 10.x cluster that includes DLO

- 1 Install Backup Exec on the active Backup Exec cluster node. Be sure to include the Desktop and Laptop Option during the Backup Exec installation.
- 2 Install Backup Exec on each failover node. Be sure to include the Desktop and Laptop Option during the Backup Exec installation.
Install Backup Exec when the node has access to the shared Backup Exec cluster disk.
- 3 From the server you used in step 1, use the Cluster Configuration Wizard to reconfigure cluster-aware Backup Exec with the same virtual server name. Add all cluster nodes that were upgraded.

Upgrading an existing Backup Exec 9.x or 10.x cluster and adding DLO to the cluster

Follow these steps to upgrade an existing Backup Exec 9.x or 10.x cluster and add DLO to the cluster.

To upgrade an existing Backup Exec 9.x or 10.x cluster and add DLO to the cluster

- 1 Install Backup Exec on the active Backup Exec cluster node. Do not select the DLO option.
- 2 Install Backup Exec on each failover node. Do not select the DLO option.
Install Backup Exec when the node has access to the shared Backup Exec cluster disk.
- 3 Use the Cluster Configuration Wizard on the active Backup Exec cluster node to uncluster all nodes. Click No when you are prompted to remove data on the shared drive. Click Yes when you are prompted to make the data available to the local node.
- 4 Install DLO on all cluster nodes. All nodes must have DLO before you can recluster them.
- 5 From the server you used in step 3, use the Cluster Configuration Wizard to reconfigure cluster-aware Backup Exec with the same virtual server name.
- 6 Add all cluster nodes that were upgraded.
The DLO services are added to the Cluster Administrator.

Reconnecting a Desktop Agent to a cluster node after you uncluster DLO

Follow these steps to reconnect a Desktop Agent to a cluster node after you uncluster DLO.

To reconnect a Desktop Agent to a cluster node after you uncluster DLO

- 1 In the `.dlo\notify` directory on the desktop user's Storage Location, create a text file named `NewMediaServerDesktopMachineName`. For example, in the file named `NewMediaServerAdmin123`, the desktop computer name is `Admin123`.
- 2 In the text file, type the name of the new DLO server on the first line and save the file.
- 3 Repeat step 1 and step 2 for all desktops that will use a new DLO server.

If the desktop is running, it should connect to the new server. If the desktop is not running, it should connect to the new server the next time it runs.

Moving a Storage Location in a DLO cluster environment before taking DLO out of the cluster

Follow these steps to move a Storage Location in a DLO cluster environment before you take DLO out of the cluster.

To move a Storage Location in a DLO cluster environment before you take DLO out of the cluster

- 1 If the Storage Location is on a shared drive or virtual server, you must move the user data for all of the Storage Location users. Move the user data to a Storage Location on the local node.
- 2 Modify all automated user assignments that are configured to use the Storage Location on the shared drive or virtual server so that they use another Storage Location on the local node.
- 3 Verify that all user data was moved off the Storage Location on the shared drive or virtual server and then delete it from the shared drive or virtual server.

About the DLO command syntax

DLO Command Line Interface commands are run from the installation directory and are executed with the `DLOCommandu` command.

The default installation directory for Backup Exec DLO is:

C:\Program Files\Symantec\Backup Exec\DLO.

If Backup Exec DLO is upgraded from a previous version that was installed in a different location, the installation is moved to this new location.

DLOCommandu is executed as follows:

```
DLOCommandu [remote-server-options] command
[command-options-and-arguments] [log-file-option]
```

About remote server options for the command line

Remote server options allow you to specify the name of the remote server on which you want to run a command. You can also enter your user name and password if required.

Remote server options are as follows:

Table Q-72 Remote server options

Option	Description
-C <computer>	Remote computer name, default to local computer
-N <user>	Fully qualified user name, e.g. Enterprise\GFord. The default is the current user
-W <password>	User password if -n is specified

DLO commands in detail

The following commands are available:

Table Q-73 Types of commands

Command	For more information
-AssignSL	See “About the -AssignSL Command” on page 1742.
-EnableUser	See “About the -EnableUser Command” on page 1743.
-ChangeServer	See “About the -ChangeServer Command” on page 1744.
-KeyTest	See “About the -KeyTest Command” on page 1745.
-ListProfile	See “About the -ListProfile Command” on page 1746.
-ListSL	See “About the -ListSL Command” on page 1747.
-ListUser	See “About the -ListUser Command” on page 1748.

Table Q-73 Types of commands (*continued*)

Command	For more information
-LogFile	See “About the -LogFile Command” on page 1748.
-Update	See “About the -Update Command” on page 1749.
-EmergencyRestore	See “About the -EmergencyRestore Command” on page 1752.
-SetRecoveryPwd	See “About the -SetRecoveryPwd Command” on page 1752.
-NotifyClients	See “About the -NotifyClients Command” on page 1752.
-InactiveAccounts	See “About the -InactiveAccounts Command” on page 1753.
-RenameDomain	See “About the -RenameDomain Command” on page 1753.
-RenameMS	See “About the -RenameMS Command” on page 1753.
-LimitAdminTo	See “About the -LimitAdminTo Command” on page 1754.
-IOProfile	See “About the -IOProfile Command” on page 1754.

About the -AssignSL Command

The -AssignSL command is used to assign a new Storage Location to existing users when the existing Storage Location is no longer available. The new Storage Location must be managed by the same media server

Caution: If the existing Storage Location is accessible, use the Move User command to move users to a new Storage Locations.

See [“Moving Desktop Agent users to a new network user data folder”](#) on page 1701.

Desktop Agent users can be assigned to new Storage Locations based on User account name, profile name, profile ID, Storage Location, Storage Location ID, and File server.

The Desktop Agent that is being moved is disabled until the media server is notified that the move is complete.

Use the following syntax:

```
DLOCommandu -assignsl -NI [-A | -F | -P | -PI | -S | -SI | -U ]
```

Note: Wildcard matches (*) are permitted in profile, Storage Location, and user names. Quotation marks are required around names if the name contains a space or colon.

Use the following command options:

Table Q-74 -AssignSL options

Option	Description
-NI <newSLID>	The -NI option is used to specify the name of the new storage location.
-A	Assigns a new storage location to all users.
-F <file server>	Assigns a new storage location to users with storage locations on the named file server.
-P <profile name>	Assigns a new storage location to users with named profile.
-PI <profile id>	Assigns a new storage location to users with given profile id.
-S <SL name>	Assigns a new storage location to users with named storage location.
-SI <SL id>	Assigns a new storage location to users with the given storage location ID.
-U <user>	Assigns a new storage location to named user account only.

The following examples show you how to use the command options:

```
DLOCommandu -assignsl -NI DLO_SL02 -A
```

```
DLOCommandu -assignsl -NI DLO_SL03 -U mmouse
```

About the -EnableUser Command

The -EnableUser command is used to enable or disable a user. Users can be enabled or disabled by All, file server (all storage locations), profile name, profile ID, storage location name, storage location ID, or user name.

Use this command if you want to force the desktop computer to refresh from the media server.

Use the following syntax:

```
DLOCommandu -enableuser [ -E | -D ] [ -A | -F | -P | -PI | -S | -SI | -U ]
```

Note: Wildcard matches (*) are permitted in profile, Storage Location, and user names. Quotation marks are required around names if the name contains a space or colon.

Use the following command options:

Table Q-75 -EnableUser options

Option	Description
-A	Enables or disables all users on the media server.
-E -D	Enables or disables a user account. The default value is to enable a user (i.e. -E).
-F <file server>	Enables or disables users with storage locations on the named file server.
-P <profile name>	Enables or disables users with the specified profile name.
-PI <profile id>	Enables or disables the users that are assigned to the specified profile.
-S <SL name>	Enables or disables the users that are assigned to the specified storage location.

The following examples show you how to use the command options:

DLOCommandu -enableuser -E -A

DLOCommandu -enableuser -D -U mmouse

About the -ChangeServer Command

The -ChangeServer command is used to reassign users to another media server.

Each desktop user must back up to a network user data folder that is managed by the same media server to which the user is assigned. If a matching automated user assignment is available on the new media server, the user is automatically assigned a profile and storage location. If a matching automated user assignment is not available, the user can be manually configured.

When a Desktop Agent user is reassigned from one media server to another, the user's current profile settings and existing backup files are not moved. They remain on the original file server.

Use the following syntax:

DLOCommandu -ChangeServer -M <media server> [-A | -F <file server> | -P <profile name> | -PI <profile id> | -S <SL name> | -SI <SL id> | -SP <SL path> | -U <user>]

Note: Wildcard matches (*) are permitted in profile, Storage Location, and user names. Quotation marks are required around names if the name contains a space or colon.

Use the following command options:

Table Q-76 -ChangeServer options

Option	Description
-A	Switches all users (default).
-F <file server>	Switches users with storage locations on the named file server.
-M <media server>	The new media server name.
-P <profile name>	Switches users based on profile name.
-PI <profile id>	Switches users based on profile id.
-S <SL name>	Switches users based on storage location name.
-SI <SL id>	Switches users based on storage location id.
-SP <SL path>	Switches users based on storage location path.
-U <user>	Switches users based on user name.

The following examples show you how to use the command options:

```
DLOCommandu -ChangeServer -M sunshine -P Desktop*
```

```
DLOCommandu -ChangeServer -M sunshine -SP \\moonlight\EngDept
```

```
DLOCommandu -ChangeServer -M sunshine -SP \\moonlight\EngDept\Enterprise-MNoel
```

About the -KeyTest Command

The -KeyTest command scans network user data to identify encrypted data that cannot be restored with the current encryption key.

Use the following syntax:

```
DLOCommandu -KeyTest
```

Use the following command options independently or in combination:

Table Q-77 -KeyTest options

Option	Description
-f	The -f option forces a full scan for all users even if the data has already been validated.
-quar	The -quar option quarantines any unrestoreable data encountered. Data that cannot be restored with the current encryption key is quarantined in the .dloquarantine folder in the user's network user data folder. If this option is not specified, the data is scanned and reported but is not quarantined.
-purge	The -purge option deletes any previously quarantined data.

The following examples show you how to use the command options:

Table Q-78 Command option examples

Item	Command
Check for unrestoreable data that has not previously been validated, or that was backed up by an old version of the Desktop Agent:	DLOCommandu -keytest
Scan all data, even if it has been previously validated, to identify unrestoreable data. Quarantine unrestoreable data.	DLOCommandu -keytest -f -quar

About the -ListProfile Command

The -ListProfile command is used to list profiles of Desktop Agent users.

Use the following syntax:

DLOCommandu -listprofile [-A | -P]

Note: Wildcard matches (*) are permitted in profile, Storage Location, and user names. Quotation marks are required around names if the name contains a space or colon.

Use the following command options:

Table Q-79 -ListProfile options

Options	Description
-A	Lists the settings for all profiles (default).
-P <profile name>	Lists the settings for only the specified profile.

The following examples show you how to use the command options:

DLOCommandu -listprofile -A

DLOCommandu -listprofile -P yourprofile

About the -ListSL Command

The -ListSL command is used to list the DLO storage locations.

Use the following syntax:

DLOCommandu -listsl [-A | -F | -S]

Note: Wildcard matches (*) are permitted in profile, Storage Location, and user names. Quotation marks are required around names if the name contains a space or colon.

Use the following command options:

Table Q-80 -ListSL options

Option	Description
-A	Lists all storage locations (default)
-F <file server>	Lists the storage locations for the named server
-S <SL name>	Lists only the named storage location

The following examples show you how to use the command options:

DLOCommandu -listsl -A

DLOCommandu -listsl -F yourserver

DLOCommandu -listsl -S yourSL

About the -ListUser Command

The -ListUser command is used to list by All, file server, profile name, profile ID, storage location name, storage location ID, or user name.

Use the following syntax:

```
DLOCommandu -listuser [ -A | -F | -P | -PI | -S | -SI | -U ]
```

Note: Wildcard matches (*) are permitted in profile, Storage Location, and user names. Quotation marks are required around names if the name contains a space or colon.

Use the following command options:

Table Q-81 -ListUser options

Option	Description
-A	Lists settings for all users (default)
-F <file server>	Lists settings for users with storage locations on the named file server
-P <profile name>	Lists settings for users by profile name
-PI <profile id>	Lists settings for users by profile id
-S <SL name>	Lists settings for users by storage location name
-SI <SL id>	Lists settings for users by storage location id
-U <user>	Lists settings for users by user name

The following examples show you how to use the command options:

```
DLOCommandu -listuser -A
```

```
DLOCommandu -listuser -P yourprofile
```

```
DLOCommandu -listuser -U mmouse
```

```
DLOCommandu -listuser -U m*
```

About the -LogFile Command

The LogFile option allows administrators to change the path or name of the LogFile. And, since every command overwrites the LogFile, to track all events (logs), you must change the path\name of the next LogFile to retain older versions.

The default path is the "\Logs" folder under the installed path:

C:\Program Files\Symantec\Backup Exec\DLO\Logs

If DLO was upgraded from a previous version, the original directory structure is used.

The default path for the "\Logs" folder in previous releases was:

C:\Program Files\VERITAS\Backup Exec\DLO\Logs

Use the following syntax:

-LogFile <path\file>

Note: Wildcard matches (*) are permitted in profile, Storage Location, and user names. Quotation marks are required around names if the name contains a space or colon.

Use the following command options:

Table Q-82 -LogFile options

Option	Description
<path>	Specifies the path to the new LogFile
<file>	Specifies the filename for the new LogFile

The following examples show you how to use the command options:

DLOCommandu -logfile test.log

DLOCommandu -logfile "c:\test.log"

About the -Update Command

The -Update command is used to list, add, remove, and publish Desktop Agent updates.

See [“About updating DLO”](#) on page 1632.

Use the following syntax:

DLOCommandu -update [-list | -add | -remove | -publish]

The following subcommands allow you to list, add, remove, or publish updates:

Table Q-83 -Update subcommands

Subcommand	Description
-List [-A -UI <update id>]	Lists settings for previously used updates.
-Add -F <file name>	Adds an "update definition file" to the updates list and assigns it a unique update ID number. The update ID number is used when the update is published with the -publish command.
-Remove [-UI <update id> -A]	Removes a file or files from the update list.
-Publish [-R] -UI <update id> [-P <profile name> -PI <profile id> -U <user>]	Makes the specified updates available to users. Users can be identified by using the following options: -P Profile name -PI Profile RecordID. To obtain the profile RecordID, run the -listprofile command. -U User name

Use the following command options:

Table Q-84 -Update command and subcommand options

Option	Description
-A	Updates all
-F <file name>	Specifies a text file that contains update records
-U <user name>	Specifies a fully qualified user name, such as Enterprise\JFord
-P <profile name>	Specifies a profile name
-PI <profile id>	Specifies a profile record id
-R	Designates to unpublish
-UI <update id>	Specifies an update record id

Note: Wildcard matches (*) are permitted in profile, Storage Location, and user names. Quotation marks are required around names if the name contains a space or colon.

The following examples show you how to use the command options:

Table Q-85 Command option examples

To do this	Description	Command
List published updates	Lists settings for all published updates	DLOCommandu -update -list -A
List details of a specific update		DLOCommandu -update -list -UI <updateID>
Add a file to the update list and assign it an ID number	Prepares an update file to be published and assigns it a unique Record ID number. The Record ID number is returned when the following command is executed	DLOCommandu -update -add -f cntlfile.txt
Publish an update to make it available to Desktop Agents	Makes updates available to users. You can specify whether to make this update available to all users, specific users, or users in a profile. You can also use wildcards to specify profile and user names	<p>To publish an update for a profile:</p> <p>DLOCommandu -update -publish -UI <updateID> -P <profile name></p> <p>DLOCommandu -update -publish -UI 63 -P yourprofile</p> <p>To publish an update for a specific user:</p> <p>DLOCommandu -update -list -UI <updateID> -U <username></p> <p>To publish an update for all users:</p> <p>DLOCommandu -update -list -UI <updateID> -U *</p>
Remove a file from the update list	Removes a file from the update list. If the file was previously published, it must be unpublished before removing it.	<p>To unpublish:</p> <p>DLOCommandu -update -publish -R -UI 33</p> <p>To remove:</p> <p>DLOCommandu -update -remove -UI 3</p>

About the -EmergencyRestore Command

The -Emergency Restore command uses the DLO administrator's recovery password to restore user data that would otherwise be unavailable if the DLO database is damaged or corrupted. The recovery password must be known to execute this command. The data is restored to the specified location in the original data structure, but it is no longer encrypted.

See [“About setting a recovery password”](#) on page 1616.

Use the following syntax:

```
DLOCommandu -EmergencyRestore <usersharepath> -W <recovery password>  
-AP <destination path>
```

You can use the following command options:

Table Q-86 -EmergencyRestore options

Option	Description
<usersharepath>	Specifies the full path to the user share directory
-W <recovery password>	Specifies the recovery password
-AP <destination path>	Specifies the path to which data is restored

About the -SetRecoveryPwd Command

The -SetRecoveryPwd command is used to change the recovery password, which enables you to retrieve encrypted data that would otherwise be lost if the DLO database is damaged or corrupted. The -SetRecoveryPwd command now updates the password for existing users as well as new users.

This recovery password can be changed only by using the DLO command line interface tools.

See [“About setting a recovery password”](#) on page 1616.

Use the following syntax:

```
DLOCommandu -SetRecoveryPwd <password>
```

About the -NotifyClients Command

The -NotifyClients command forces the Desktop Agents to refresh the profile settings immediately, or the next time the Desktop Agent connects if it is offline.

Use the following syntax:

```
DLOCommandu -notifyclients
```

About the -InactiveAccounts Command

The -InactiveAccounts command is used to list and delete the accounts that have not been used in a specified number of days.

Use the following command to list inactive accounts:

```
dlocommandu -inactiveaccounts -list -days <#days>
```

This command returns a list of inactive accounts.

The list includes the following information, which is used to delete specific accounts:

- computer name
- computer ID
- domain\user name
- userID

Use the following command to delete specific inactive accounts:

```
dlocommandu -inactiveaccounts -delete -U <domain\user name> -M <computer name> -days <#days>
```

```
dlocommandu -inactiveaccounts -delete -UI <userID> -MI <computer ID> -days <#days>
```

Where -U and -M are used to delete the user and computer by name and -UI and -MI are used to delete the user and computer by ID.

Use the following command to delete ALL accounts inactive for a specified number of days:

```
dlocommandu -inactiveaccounts -delete -a <#days>
```

About the -RenameDomain Command

The -RenameDomain command is used after a Windows domain has been renamed. Running the RenameDomain command changes each Desktop Agent user's record to reflect the new domain name and changes the path for the network user data folder. It also notifies each Desktop Agent of the change.

Use the following syntax:

```
DLOCommandu -RenameDomain <OldDomainName> <NewDomainName>
```

About the -RenameMS Command

The -RenameMS command is used when a media server has been renamed. Running the RenameMS command updates the installation share, storage location paths,

and network user data folder paths. It also notifies each Desktop Agent of the change.

Before you can use the `-RenameMS` command, you must do the following steps in the order listed:

- Use the Windows control panel to rename the media server.
See the Microsoft Windows documentation.
- Use Backup Exec Utility to update the configuration for the new media server name.
See the Backup Exec Utility online help.

After you rename the media server and use Backup Exec Utility to update the configuration, you can use the `-RenameMS` command.

Use the following syntax:

```
DLOCommandu -RenameMS <OldServerName> <NewServerName>
```

About the `-LimitAdminTo` Command

The `-LimitAdminTo` command limits administration of DLO to the specified group or user.

Use the following syntax:

```
DLOCommandu -LimitAdminTo -NAU <domain\NewAdminName>
```

```
DLOCommandu -LimitAdminTo -NAU <domain\NewAdminGroup>
```

You can use the following command options:

Table Q-87 -LimitAdminTo options

Option	Description
-NAU	The <code>-NAU</code> option is used to add a new DLO administrator or to add a group that can be used of DLO administrators.
-DAU	The <code>-DAU</code> option is used to delete a DLO administrator or a DLO administration group.
-L	The <code>-L</code> option lists all of the current DLO administrators and groups.

About the `-IOProfile` Command

The `-IOProfile` command enables a profile to be exported from one media server, and then imported to another media server. An option is also provided to import global settings.

Note: When a profile is imported, it does not initially have any users assigned to it, so there is no immediate impact. When global settings are imported, they immediately apply to all Desktop Agent users who are assigned to the server.

Use the following command to export a profile:

```
DLOCommandu -C <master server name> -IOProfile -DBF <export file name> -E <profile name>
```

This exports the requested named profile (-E) from the specified server (-C) into the named file (-DBF). It is not necessary to specify the master server name with the -C option if the profile is on the same server where the command is run.

Use the following command to import a profile:

```
DLOCommandu -C < server name> -IOProfile -DBF <export file name>
```

This imports the profile in the given file (-DBF) into the named server (-C.)

Use the following command to import the console settings for DLO administrator account management in addition to the profile:

```
DLOCommandu -C < server name> -IOProfile -DBF <export file name> -IPRGCS
```

Use the following command to import the global settings in addition to the profile:

```
DLOCommandu -C < server name> -IOProfile -DBF <export file name> -IPRGS
```

About the Desktop Agent

The Desktop Agent is the component of the Backup Exec Desktop and Laptop Option that protects files on desktop and laptop computers (collectively referred to as desktops). It backs up data to the desktop's local drive and to a storage location on the network.

The DLO administrator initially configures the Desktop Agent. Your profile determines the level of interaction between you and the Desktop Agent. The administrator may also configure the Desktop Agent to run without a user interface, with a fully functional user interface, or somewhere in between.

If the DLO administrator has set your profile so that you can view the complete Desktop Agent and modify settings, then you can do the following:

- Restore files
- Synchronize files between multiple desktops

- Configure backup selections
- Set schedules
- View history

See [“Desktop Agent terminology”](#) on page 1756.

Desktop Agent terminology

The following terms are used in the DLO documentation:

Table Q-88 DLO terminology

Option	Description
Desktop	The desktop or laptop computer on which the Desktop Agent runs.
Desktop Agent	The DLO program that runs on desktop and laptop computers, and its user interface.
desktop user data folder	A folder on the desktop in which backup files are stored for offline availability.
network user data folder	A folder on a network file server where your backup data is stored.
Profile	Specifies detailed configuration settings for Desktop Agent operation. Profiles are assigned to groups of similar desktop users.
Synchronization	The process of maintaining the current revision of a given file on more than one desktop.
Automated User Assignment	Assigns a profile and a Storage Location to the desktop user when the Desktop Agent is first installed on a desktop.

A full glossary of DLO terms is available.

Features and benefits of the Desktop Agent

The Desktop Agent provides the following features:

Table Q-89 Features of the Desktop Agent

Item	Description
Data protection	Selected files are copied automatically to user data folders on the desktop's local drive and on the network. The Desktop Agent can be configured so that no user interaction is required. Files are protected automatically when the desktop is online or offline. Backup Exec protects the data by backing up the network user data folders on the DLO file server.
Data availability	You can access data from multiple desktops in multiple locations by using the same login credentials on each desktop. You can also restore previous file revisions if you save at least one file revision in the desktop user data folder.
Synchronization	<p>A user that accesses multiple computers with the same login credentials can configure folders to be synchronized on each of the computers.</p> <p>When a synchronized file is changed on one computer, the updated file is copied to the following locations all other computers that are configured for synchronization:</p> <ul style="list-style-type: none"> ■ The network user data folder ■ The desktop user data folder

System requirements for the Desktop Agent

The following are the minimum system requirements for running this version of the Desktop Agent.

Table Q-90 Minimum system requirements for the Desktop Agent

Item	Description
Operating system	<p>The following operating systems are supported:</p> <ul style="list-style-type: none"> ■ Microsoft Windows 2000 ■ Microsoft Windows XP Service Pack 2 or later ■ Microsoft Windows XP Professional x64 Edition ■ Microsoft Windows Vista ■ Microsoft Windows 7 <p>The Desktop Agent is not supported on any server operating system, including Windows Server 2003/Storage Server 2003/2000 Server.</p>
Processor	Pentium system

Table Q-90 Minimum system requirements for the Desktop Agent (*continued*)

Item	Description
Memory	Required: 256 MB RAM Recommended: 512 MB (or more for better performance).
Internet browser	Internet Explorer 5.01 or later; however, version 5.5 is recommended.
Disk space	25 MB hard disk space is required after Microsoft Windows is installed (typical installation). Additional space may be required if the desktop user data folder is enabled.
Other hardware	Network interface card or a virtual network adapter card.

Installing the Desktop Agent

The DLO administrator determines who installs the Desktop Agent. It can be either the administrator or the desktop user. Administrator rights are required to install the Desktop Agent. If you need to restart the desktop during installation, you must use the same administrator logon account again to ensure that the installation completes successfully.

After the Desktop Agent is installed on a desktop, anyone who logs on to that desktop can use the Desktop Agent. The logged on user has access to only the DLO backup files that are associated with the logged-on account.

All computers that run the DLO Administration Console or the Desktop Agent should be set to a common time. This can be accomplished by configuring the Windows Time Synchronization service on the network. See the Microsoft Web site for additional information.

To install the Desktop Agent

- 1 From the desktop on which you want to install the Desktop Agent, browse to the network server where the installation files for the Desktop Agent are stored.

The default location is \\<Backup Exec media server name>\DLOAgent. If you are unsure of the location, contact the administrator.
- 2 Double-click setup.exe.
- 3 On the Welcome screen, click **Next**.
- 4 Read the license agreement, and then click **I accept the terms in the license agreement**.
- 5 Click **Next**.

6 Do one of the following:

To install the Desktop Agent in the default location Continue with step 7.

The default installation location is C:\Program Files\Symantec\Backup Exec\DLO.

To install the Desktop Agent in a location of your choice Do the following in the order listed:

- Click **Change**.
- Enter the path for the location where you want to install the Desktop Agent.
- Click **OK**.

7 Click **Next**.

8 Click **Install**.

9 Click **Finish** to install the Desktop Agent.

How to configure the Desktop Agent

You can configure the Desktop Agent in the following ways:

- Connect to the media server.
See [“About connecting from the Desktop Agent to the media server”](#) on page 1759.
- Use local accounts on desktops.
See [“About using local accounts on desktop computers”](#) on page 1761.
- Use alternate credentials.
See [“Alternate Credentials options”](#) on page 1760.
- Reset dialog boxes and account information.
See [“Resetting dialog boxes and account information in DLO”](#) on page 1762.
- Change your connection status.
See [“Changing your connection status”](#) on page 1762.
- Enable or disable the Desktop Agent.
See [“Enabling the Desktop Agent”](#) on page 1763.
See [“Disabling the Desktop Agent”](#) on page 1763.

About connecting from the Desktop Agent to the media server

The Desktop Agent communicates with the DLO database and services on the media server during normal operation. When you use the Desktop Agent, you must connect to the media server by using a domain account.

Note: If you connect to the media server with one set of credentials, and then try to connect to the server with a different set of credentials, authentication may fail. Restart the computer to reconnect.

When new information is available for the Desktop Agent, the Desktop Agent receives a notification of this new information and retrieves it. For example, when settings or synchronized files change or if a software update is available. The Desktop Agent and the media server do not contact each other directly.

Caution: If you try to connect to a server using characters in the share name that do not exist on the code page for the local system, the connection fails. Code pages map character codes to individual characters, and are typically specific to a language or group of languages.

Alternate Credentials options

The Desktop Agent uses the logon account by default. However, if an alternate account was specified, it might be used, such as to connect across domains.

If you are logged on with credentials that the Desktop Agent does not recognize, you can specify alternate credentials for Desktop Agent operation and save the account information for future sessions. If you prefer, you can disable an account for Desktop Agent operations so that the Desktop Agent does not run when you are logged on with the account currently being used. You can save this account info for future connections.

Note: If you have a previously established network connection to the media server and it does not match the account the Desktop Agent uses, the Desktop Agent tries to reconnect as the Desktop Agent user. If this fails, the following error displays: "Multiple connections to a server or shared resource by the same user, using more than one user name, are not allowed. Disconnect all previous connections to the server or shared resource and try again." The account used by the Desktop Agent is the logon account by default, but could be an alternate account if one has been specified; for example, to connect across domains.

In a cross-domain configuration where there is no trust relationship, if multiple users are running the same Desktop Agent, each user must provide a unique user name and password in the media server domain. If different users use the same credentials, DLO displays an error message stating that the user is already connected to the media server.

Table Q-91 **Alternate Credentials** options

Item	Description
Use this account	Enables the Desktop Agent to run when you use the account under which you are currently logged on.
Username	Indicates the user name for an account that is authorized for Desktop Agent operation.
Password	Indicates the password for the account to be used for Desktop Agent operation.
Domain	Indicates the domain for the account to be used for Desktop Agent operation.
Save my password	<p>Enables DLO to save and use this password in the future. You can then automatically authenticate to the media server or storage location in the event of an authentication failure.</p> <p>This option appears only if the DLO administrator has enabled this option. On newly-deployed Desktop Agents, this option does not appear until the second time the Desktop Agent connects to the media server.</p>
Disable this account	Prevents the Desktop Agent from running when you use the account under which you are currently logged on.

About using local accounts on desktop computers

If you log on to your desktop with a local account, the Desktop Agent prompts you for your user name and password for your domain account.

You should consider the following information when you use local accounts on the desktops that run the Desktop Agent:

- You can use a set of domain credentials with one local account. If you use more than one local account on a desktop or laptop computer, you should either disable DLO for other accounts or have unique domain credentials for each account.

See [“Alternate Credentials options”](#) on page 1760.

For example, if you usually log on to the desktop computer as "myusername", you should have a domain account to use for DLO with this account. If you also occasionally log on as 'administrator', DLO can be disabled when you are logged on to this account. Alternately, you can provide a unique set of domain credentials to use for DLO when you are logged on as "administrator".

- Multiple users of the same desktop computer can all use DLO. However, they must provide unique credentials for the desktop computer and unique domain credentials for connection with the Desktop Agent.
- DLO does not support the Fast User Switching feature of Windows XP.

Resetting dialog boxes and account information in DLO

You can prevent certain dialog boxes from appearing by checking the **Don't show me this message again** check box. However, you can reset disabled dialog boxes. If passwords and account information are cleared, the Desktop Agent prompts for this information if it is required to access a resource.

To reset dialog boxes and account information

- 1 On the Tools menu, click **Options**.
- 2 On the **Preference** tab, do one of the following:

To reset the dialog boxes that were disabled

- Do the following in the order listed:
- Click **Reset dialogs**.
 - Click **Yes** at the prompt.

To clear passwords and account information

- Do the following in the order listed:
- Click **Reset accounts**.
 - Click **Yes** at the prompt.

- 3 Click **OK**.

Changing your connection status

When you use the Desktop Agent, your connection status is displayed in the lower right corner of the Desktop Agent Console.

When the Desktop Agent is in offline mode, the following are true until you choose to work online again:

- Files are not transferred to the network user data folder. Pending files remain in the pending files list with a status of "Pending network."
- Job logs are not copied up to the network user data folder.
- Alerts are not posted to the media server.

The DLO Administrator sets a maximum time after which the Desktop Agent is automatically returned to the online mode, if a network connection is available.

To change your connection status

- 1 Click the connection status on the lower right corner of the Desktop Agent.
- 2 Do one of the following:
 - Click **Work Offline** to place the Desktop Agent in offline mode.
 - Click **Work online** to place the Desktop Agent in online mode.

Enabling the Desktop Agent

If the Desktop Agent has been disabled, and your Profile allows it, you can re-enable the Desktop Agent.

The **Enable** option is not available if you do not have permission to take this action.

To enable the Desktop Agent

- 1 On the Windows system tray, right-click the Desktop Agent icon.
- 2 Click **Enable**.

Disabling the Desktop Agent

If your Profile allows it, you can disable the Desktop Agent.

The **Disable** option is not available if you do not have permission to take this action.

To disable the Desktop Agent

- 1 On the Windows system tray, right-click the Desktop Agent icon.
- 2 Click **Disable**.

About the Desktop Agent Console

The Desktop Agent Console is the user interface for the Desktop Agent. The DLO administrator controls access to the Desktop Agent Console.

Note: To ensure that you have the latest status and settings at any time while using the Desktop Agent, use the Refresh feature.

The DLO administrator may choose from the following:

Table Q-92 User interface options for the Desktop Agent

Item	Description
Display the complete interface	Enables desktop users to access all Desktop Agent options
Display only the status	Enables desktop users to view the status of backup jobs, but they cannot change Desktop Agent settings or access options other than status. Desktop users can right-click the system tray icon to open the status view or exit the program.
Display only the system tray icon	The desktop user sees only the Desktop Agent icon in the system tray in the lower right corner of the screen. Desktop users can right-click the system tray icon to exit the program.
Do not display anything	The Desktop Agent runs in the background. The desktop user cannot view the Desktop Agent.

The Desktop Agent Console has the following components:

Table Q-93 Desktop Agent Console features

Item	Description
Menu bar	The menu bar appears across the top of the screen. To display a menu, click the menu name. Some menu items are not available until an item is selected from the console screen.
Tasks bar	The Tasks bar appears on the left side of the Desktop Agent Console. To hide the Tasks bar, from the View menu, select Tasks Bar . Actions are initiated from the tasks bar, and these actions vary with the selected view.
Views menu	<p>The Views menu appears in the Tasks bar and enables you to navigate to the following views:</p> <ul style="list-style-type: none"> ■ Status See “About the status of the Desktop Agent” on page 1782. ■ Backup selections See “About using the Desktop Agent to back up your data” on page 1765. ■ Synchronized selections See “About synchronizing desktop user data” on page 1778. ■ Restore See “Restoring files by using the Desktop Agent” on page 1786. ■ History

Table Q-93 Desktop Agent Console features (*continued*)

Item	Description
Tasks menu	Actions are initiated from the tasks menu. These actions vary with the selected view.
Tools menu	Includes Options, which enables you to do the following: <ul style="list-style-type: none">■ Reset dialog boxes that were suppressed by the Don't show me this message again check box.■ Clear passwords and account information that the Desktop Agent has stored. See "Resetting dialog boxes and account information in DLO" on page 1762.

About using the Desktop Agent to back up your data

When data is backed up by the Desktop Agent, it is transferred to the user data folder on the desktop's local drive. Then, the data is transferred to a network user data folder, which is assigned by the DLO Administrator. Network user data folders are typically also backed up by Backup Exec, which provides an additional level of protection.

Caution: If you try to connect to a server using characters in the share name that do not exist on the code page for the local system, the connection fails. Code pages map character codes to individual characters, and are typically specific to a language or group of languages.

See ["About using DLO to back up Outlook PST files incrementally"](#) on page 1769.

See ["About restoring Microsoft Outlook Personal Folder files"](#) on page 1789.

Select files that you want to protect from the Backup Selections view. The DLO administrator assigns the initial backup selections. However, if the DLO administrator has set your profile to view the complete Desktop Agent and modify settings, then you can choose your backup selections.

You can change Desktop Agent settings and backup selections when you work offline. The settings are stored until you are working online, at which time they are automatically transferred. If the administrator has also made changes that conflict with the changes made on the Desktop Agent, the administrator's changes are used.

You can view and modify backup selections using two views: standard and advanced. The standard view lists the contents of your local drives, allowing you

to check off files and folders to be backed up. It also uses default backup selection settings to add new selections. The advanced view provides more configuration options for selections.

A backup selection consists of the following items:

- A folder or list of folders
- Criteria for the files to be included or excluded from the backup
- Limits on the number of file revisions to retain
- Settings for compression, backup file deletion, and encryption

About revisions

Revisions are versions of a file at a specific point in time. When a file is changed and backed up, DLO stores a new revision. DLO stores and maintains a specific number of revisions for all files in a backup selection. Because each backup selection is configured separately, the number of revisions that are retained can vary for different backup selections.

When the number of revisions is exceeded, DLO removes the oldest revision. It maintains only the specified number of revisions in the desktop and network user data folders.

You can limit the number of revisions DLO retains in a given period of time. If you back up a document frequently, all of your revisions could potentially be a few minutes apart. By specifying that you want to retain only 2 revisions every 24 hours, at least 120 minutes apart, you can retain older revisions for a longer period of time. While some intermediate versions are not retained, it does support situations in which returning to an older revision is needed.

Another consideration in determining the number of revisions to retain is the amount of storage space that is required to store the data. The amount of space that is required for backups can be estimated by multiplying the number of revisions that are retained by the amount of data protected.

For example, if you retain three revisions of each file and have 10 MB to back up, approximately 30 MB of disk space is required.

Although compression can improve the space utilization, it varies significantly with file type and other factors.

DLO protects all of the alternate streams for a file, including security streams. If a new version of a file contains only alternate stream data modifications, the new version replaces the old version without impacting the revision count.

See [“Modifying backup selections in the Desktop Agent's standard view”](#) on page 1767.

See [“Adding backup selections in the Desktop Agent's advanced view”](#) on page 1768.

See [“About using DLO to back up Outlook PST files incrementally”](#) on page 1769.

See [“About restoring Microsoft Outlook Personal Folder files”](#) on page 1789.

See [“About restoring files with alternate stream data”](#) on page 1790.

Modifying backup selections in the Desktop Agent's standard view

The backup selections standard view provides a list of drives, folders, and files that you can select for backup.

When you create new backup selections in the standard view, the default backup selection settings are used. When you add new subfolders and files to the backup selection using the standard view, the new backup selections have the same settings as the original selections.

In the standard view, files and folders are represented in a tree view where users can select or deselect files and folders for backup. When the check box next to a file or folder is grayed out, the selection was defined by the administrator and can only be changed in the advanced view if the administrator has granted this right in the profile definition.

See [“Modifying backup selections in the Desktop Agent's advanced view”](#) on page 1768.

When a red X appears in the check box next to a file or folder, this item has been globally excluded from all backups by the administrator and cannot be selected.

After clicking Save, previously backed-up selections that were unchecked are treated like deleted backup selections and will no longer be backed up. The backup files for this selection will be deleted after the number of days specified in the backup selection settings. The source files for the deleted backup selection will not be deleted by the Desktop Agent.

Checked folders that were not previously checked are added to the backup selections for this desktop.

To modify backup selections in the standard view

- 1 In the Tasks bar, under Views, click **Backup Selections**.
- 2 Click **Standard view**.

- 3 Check the folders and files you want to back up, and uncheck the files and folders that you no longer want to back up.

Expand selections by clicking the plus sign (+) and collapse selections by clicking the minus sign (-).
- 4 Click **Save changes** to save the new settings or **Undo changes** to return to the last saved settings.

Adding backup selections in the Desktop Agent's advanced view

The advanced view provides more configuration options than the standard view.

To add a backup selection in the Backup Selections advanced view

- 1 Under Views in the Desktop Agent Tasks Bar, click **Backup Selections**.
- 2 Click **Advanced view**.
- 3 Click **Add**.
- 4 Do any of the following to customize the backup selection properties:
 - On the General tab, set general backup selection properties including the name, description, and folder to be backed up.
See [“General options for DLO backup selections”](#) on page 1661.
 - On the Include/Exclude tab, include specific files in or exclude specific files from this backup selection.
See [“Include/Exclude options for DLO backup selections”](#) on page 1663.
 - On the Revision Control tab, set revision control for this backup selection.
See [“Revision Control options for DLO backup selections”](#) on page 1664.
 - On the Options tab, set Delta File Transfer, encryption, and compression options for this backup selection.
See [“Options for a DLO backup selection”](#) on page 1666.
- 5 Click **OK** to save your changes.

Modifying backup selections in the Desktop Agent's advanced view

From the advanced view, backup selections created on the Desktop Agent and those created by the DLO administrator in the profile can be modified if the profile grants sufficient rights to the Desktop Agent user.

To modify backup selections in the advanced view

- 1 Under Views in the Desktop Agent Tasks Bar, click **Backup Selections**.
- 2 Click **Advanced view**.

- 3 Select the backup selection you want to change, and then click **Modify**.

Profile backup selections are those set by the DLO administrator. If the backup selection is a profile backup selection, and if the user has been granted sufficient rights, it can be modified by selecting Use custom selection in the drop-down menu. Once this option is selected, your backup selection will no longer be updated when the administrator updates the profile backup selection.

You can return to the profile backup selection settings at any time by selecting Use Profile selection in the drop-down menu. Once you make this selection, your profile will be updated if the DLO administrator modifies the profile backup selection.

- 4 Modify the backup selection properties as needed.
- 5 Click **OK**.

Deleting backup selections in the Desktop Agent's advanced view

When you delete a backup selection, the backup files are deleted by the Backup Exec grooming process after the number of days specified in the backup selection.

To delete a backup selection

- 1 Under Views in the Desktop Agent Tasks Bar, click **Backup Selections**.
- 2 Click **Advanced view**.
- 3 Select the backup selection you want to delete.
You cannot delete profile backup selections.
- 4 Click **Remove**.
- 5 Click **Yes** to verify that you want to delete this backup selection, or click **No** to cancel.

About using DLO to back up Outlook PST files incrementally

DLO is configured to back up PST files incrementally by default. Incremental backup of PST files is controlled by the administrator in the Profile, or by the desktop user in Options dialog if the desktop user has been granted sufficient rights.

Note: Outlook must be the default mail application to perform incremental backups of Outlook PST files.

The following should be considered when backing up PST files incrementally:

- When Outlook PST files are backed up incrementally, only one revision is maintained regardless of the number of revisions set in the backup selection.
- When you restore Microsoft Outlook PST files, the restored PST file will differ from the original PST file.
See [“About restoring Microsoft Outlook Personal Folder files”](#) on page 1789.
- Synchronized files cannot be backed up incrementally.
- When a DLO profile is configured to limit the bandwidth usage during data transfer to the network user data folder, bandwidth is not limited during the incremental transfer of PST files.

If you do not intend to use Outlook as your default mail application, you can disable the warning message about incremental backups.

See [“Setting customized options on the Desktop Agent”](#) on page 1775.

When an Outlook PST file is included in a DLO backup selection, it will appear in the Desktop Agent pending queue whenever the PST file is closed. Because PSTs are a shared resource, opening and closing of PST files is controlled by a process called MAPI. Both DLO and Outlook access PSTs via the MAPI process. MAPI opens a PST upon request from the application.

Depending on the version of MAPI that is in use, MAPI may or may not close a PST in response to the following:

- An application such as DLO or Microsoft Outlook detaches from the PST, such as when Outlook is closed
- DLO startup
- After 30 minutes of inactivity in the PST

When the PST is closed DLO does one of the following. If the PST is being handled incrementally via MAPI (see section on incremental PSTs) DLO determines if the PST has been backed up in its entirety. If it has already been backed up then the entry is simply removed from the Desktop Agent pending queue because DLO knows the PST is in sync. If the PST is not being handled incrementally, the PST will be backed up in its entirety at this time.

See [“About restoring Microsoft Outlook Personal Folder files”](#) on page 1789.

About backing up Lotus Notes NSF files incrementally

The following types of Lotus Notes NSF Files can be backed up incrementally:

Table Q-94 NSF files that can be backed up incrementally

File Name	Location	Description
BOOKMARK.NSF	Notes\Data directory	Contains saved bookmarks and Welcome Page information.
NAMES.NSF	Notes\Data directory	This file contains contacts, connections, locations and Personal Address Book information.
A_<name>.NSF		This is an email archive file. Email must be archived to be incrementally backed up by DLO. See Lotus Notes documentation for additional information on archiving email.

When a file is backed up incrementally, there is no progress indicator in the Desktop Agent Status view, and only one revision is retained.

Note: When a DLO profile is configured to limit the bandwidth usage during data transfer to the network user data folder, bandwidth is not limited during the incremental transfer of Lotus Notes NSF files.

Lotus Notes must already be installed before the Desktop Agent is installed. If Lotus Notes is installed after the Desktop Agent, you must run the Desktop Agent installer again to repair the installation. Additionally, if Lotus Notes is open during the Desktop Agent installation, Lotus Notes must be restarted.

Lotus Notes email files can only be backed up incrementally with DLO if the emails have been archived. Once emails are archived, the resulting archive file can be backed up incrementally. See the Lotus Notes documentation for information on archiving emails.

Deleted Lotus Notes email files are not backed up.

See [“Configuring the Desktop Agent for incremental backup of Lotus Notes files”](#) on page 1771.

Configuring the Desktop Agent for incremental backup of Lotus Notes files

If Lotus Notes is installed, you can back up email files incrementally.

See [“About backing up Lotus Notes NSF files incrementally”](#) on page 1770.

To configure the Desktop Agent for incremental backup of Lotus Notes files

- 1 Verify that Lotus Notes was installed before the Desktop Agent was installed, or that the Desktop Agent installer was run again after Lotus Notes was installed to repair the installation.
- 2 Verify that emails to be backed up have been archived in Lotus Notes.
- 3 Verify that the Lotus Notes NSF files to be backed up have been selected in the appropriate backup selection.
See [“About using the Desktop Agent to back up your data”](#) on page 1765.
- 4 In the Tasks bar, under Tools, click **Options**.
- 5 On the Options tab, check **Enable message level incremental backups of Lotus Notes email files**.
- 6 Click **OK**.

About using the Desktop Agent when Lotus Notes is not configured for the current user

When a user logs in to a computer that has both DLO and Lotus Notes installed, but that user is not yet configured in Lotus Notes, a debugging DOS-window may appear which contains the following errors:

```
<time_date_stamp> Created new log files as C:\Documents and  
Settings\\Local Settings\Application Data\Lotus\Notes\Data\log.nsf.
```

```
<time_date_stamp> A previous process with the process ID <####> failed to  
terminate properly.
```

The DOS-window cannot be closed without manually exiting the DLO process, but can be remedied by configuring the current user for Lotus Notes. Once the user is configured, the errors will no longer be generated at login for that user.

About modifying Desktop Agent settings

If the DLO administrator has set your profile so that you can view the complete Desktop Agent and modify settings, you can use the Settings view to modify the following:

- Backup job schedule options
- Desktop user data folder location
- Desktop user data folder disk space limits
- Log file disk space limits

- Logging level
- Bandwidth usage

The Desktop Agent continues to use the settings that are specified in the profile until you specifically elect to use customized schedules or options.

See [“Changing schedule options for a DLO backup job”](#) on page 1773.

See [“Setting customized options on the Desktop Agent”](#) on page 1775.

You can change Desktop Agent settings and backup selections when you work offline. The settings are stored until you are online, at which time they are automatically transferred. If the administrator has also made changes that conflict with the changes made on the Desktop Agent, the administrator’s changes are used.

Note: Changing settings on one Desktop Agent causes settings to be loaded on other Desktop Agents that use the same authentication. Running jobs are canceled and then restarted.

Changing schedule options for a DLO backup job

You can change backup job schedule options if the DLO administrator has set your profile so that you can view the complete Desktop Agent and modify settings.

To change schedule options for a backup job

- 1 In the Tasks bar, under Tools, click **Options**.
- 2 On the Schedule tab, select the appropriate options, and then click **OK**.

See [“Schedule options”](#) on page 1773.

Schedule options

You can use the schedule that is associated with the profile or you can set up a customized schedule, if your profile allows it.

See [“Changing schedule options for a DLO backup job”](#) on page 1773.

Table Q-95 **Schedule options**

Item	Description
Use Profile schedule	Select Use Profile schedule from the drop-down menu to use the scheduling options that are specified in the profile. If this option is selected, additional settings on the Schedule tab cannot be modified.
Use custom schedule	Lets you specify a customized schedule that differs from the profile schedule.
Whenever a file changes	Backs up files automatically whenever they change. This feature is available only for NTFS file systems. For FAT file systems, type a number of minutes or hours between backups in the Back up changed files every field.
According to a schedule	Backs up files according to a schedule. The default is to run a backup at 11:00 P.M. every Monday, Tuesday, Wednesday, Thursday, and Friday. Click Modify to change this default.
Manually	Runs a backup only when you initiate it.

Table Q-95 Schedule options (continued)

Item	Description
Logout/Restart/Shutdown	<p>Lets you choose from the following options:</p> <ul style="list-style-type: none"> <li data-bbox="720 369 1245 545"> <p>■ Do nothing Proceeds with a logout, restart, or shutdown even when files require a backup. If a job is already running, a prompt asks if the user wants to log out, restart or shutdown when the job is complete.</p> <li data-bbox="720 552 1245 730"> <p>■ Prompt user to run job Prompts the user to run a backup job before proceeding with the logout, restart, or shutdown. If a job is already running, a prompt asks if the job should be canceled to continue with the logout, restart, or shutdown.</p> <li data-bbox="720 737 1245 916"> <p>■ Run job immediately Backs up waiting files without prompting before proceeding with a logout, restart, or shutdown. If a job is already running, a prompt asks if the job should be canceled to continue with the logout, restart, or shutdown.</p> <li data-bbox="720 923 1245 1102"> <p>■ Run job as scheduled Proceeds with a logout, restart, or shutdown and back up files according to the schedule. If a job is already running, a prompt asks if the job should be canceled to continue with the logout, restart, or shutdown.</p> <li data-bbox="720 1109 1245 1194"> <p>■ Run job at next login Proceeds with a logout, restart, or shutdown without prompting, and run a job at the next logon.</p>

Setting customized options on the Desktop Agent

You can change additional Desktop Agent settings if your profile allows it.

To set customized options

- 1 In the Tasks bar, under Tools, click **Options**.
- 2 On the Options tab, select **Use custom options** from the drop-down menu.
- 3 Select the appropriate options, and then click **OK**.

See [“Options for the Desktop Agent”](#) on page 1776.

Options for the Desktop Agent

If your profile lets you modify settings, you can change the logging options and disk space usage for your computer.

See [“Setting customized options on the Desktop Agent”](#) on page 1775.

Table Q-96 Options for the Desktop Agent

Item	Description
Use Profile options	<p>Indicates that the Desktop Agent uses the scheduling options that are specified in the profile.</p> <p>If this option is selected, additional settings on the Schedule tab cannot be modified.</p>
Use customized options	<p>Lets you set up a customized schedule that differs from the profile schedule.</p> <p>This option must be selected to enable access to additional settings on the Options tab.</p>
Limit disk space usage on my computer to:	<p>Limits the amount of space that is used on the computer to store backup files.</p> <p>Select % to enter a percentage of the hard disk space that can be used to store backup files.</p> <p>Select MB to enter the maximum number of megabytes of disk space that can be used to store backup files.</p>
Keep log files for a minimum of (days)	<p>Indicates the minimum number of days to keep log files. Log files are not deleted until they are at least as old as specified.</p> <p>Log grooming occurs each time a log is created. Log files are not deleted until the minimum age has been reached and, when the combined size of all log files, is also reached.</p>
After minimum number of days, delete oldest log files when combined size exceeds	<p>Indicates the maximum combined size of all log files to be retained before the oldest log files are deleted.</p> <p>You may have more than the specified number of MB of log files stored if none of the log files are as old as specified in the Keep log files for a minimum of (days) setting.</p>
Log groom messages	Creates logs for grooming operations.
Log information messages for backup	Creates logs for all backup operations.
Log warning messages	Creates logs for all operations that generate warnings.

Table Q-96 Options for the Desktop Agent (*continued*)

Item	Description
Enable message level incremental backups of Outlook PST files	<p>Enables incremental backups of Microsoft Outlook Personal Folder (PST) files. Incremental backups must be enabled to allow PST files to be backed up while they are open.</p> <p>If this option is not checked, PST files that are configured in Outlook are fully backed up each time the PST file is saved. In general, PST files are saved when Outlook is closed.</p> <p>See “About using DLO to back up Outlook PST files incrementally” on page 1769.</p>
Enable message level incremental backups of Lotus Notes email files	<p>Enables the configuration of DLO for incremental backup of certain Lotus Notes NSF files. Additional steps may be necessary to insure backup of these files.</p> <p>See “Configuring the Desktop Agent for incremental backup of Lotus Notes files” on page 1771.</p> <p>Unchecking this box prevents the incremental backup of Lotus Notes files.</p>

Moving the desktop user data folder

You can change the location of the desktop user data folder if your profile allows it.

To move the desktop user data folder

- 1 In the Tasks bar, under Tools, click **Settings**.
- 2 On the Backup Folders tab, click **Move**.
- 3 In the Browse for folder dialog box, choose a new location for the desktop user data folder.
- 4 Click **OK**.
- 5 When prompted to continue, click **Yes**.
- 6 Click **OK**.

Customizing connection policies

The Desktop Agent can be configured to disable or limit backups for certain connection types. For example, if the DLO administrator has granted you sufficient rights, you can choose to disable backups when you are connected by a dialup

connection. Then, you can continue backing up when you are connected to a higher speed connection.

When backups are limited by a connection policy, files are backed up to the desktop user data folder. Files are transferred to the network user data folder when connection policies are no longer limiting backups. If the desktop user data folder is disabled, no offline protection is provided.

When connection policies are created using Active Directory settings, and two or more policies match a specific user or computer, the most restrictive policy is used.

Example:

One connection policy that matches a specific user or computer disables backups to the network user data folder of all files over 500 KB. A second connection policy that also matches the computer or user disables all backups to the network user data folder. The second policy is used because it is more restrictive to limit all backups than just backups of large files.

To customize connection policies

- 1 In the Tasks bar, under Tools, click **Settings**.
- 2 On the Connection Policies tab, select the appropriate options, and then click **OK**.
See [“Add/Edit Connection Policy options”](#) on page 1656.
- 3 If you selected Active Directory in step 2, configure the Active Directory settings, and then click **OK**.
See [“Active Directory Object options”](#) on page 1685.
- 4 Click **OK** twice.

About synchronizing desktop user data

Your backed up data is stored in the desktop user data folder on the local drive of each desktop, and in the network user data folder. If you have multiple desktops, your network user data folder contains copies of backed up files from each desktop. When a folder is synchronized, only one copy of the folder and its contents is included in the network user data folder. When the file is changed on one desktop, it is stored in the desktop user data folder on that computer, and then uploaded to the network user data folder the next time a DLO job is run. It is then available for download to another synchronized desktop computer the next time that computer runs a job.

After a folder is synchronized, the Desktop Agent checks the network user data folder each time the desktop is connected to the network and a job is run. If new file versions are available in any of the synchronized folders, the Desktop Agent downloads the new version to the user data folder on the desktop. If you change a file on your current desktop and change the same file on one of your other backed up computers without synchronizing the files, a conflict occurs and you are prompted to select which file revision to use.

By synchronizing backed-up data, you can work on a file on any of your desktops with the assurance that you are working on the most recent version.

The Synchronized Selections view displays folders that were backed up on your other desktops and are available for synchronization. Select any of these folders that you want to synchronize with the current desktop computer.

Note: If you customize NTFS permissions or folder attributes for compression or encryption, you must reapply these settings after restoration or synchronization.

See [“How synchronization works”](#) on page 1779.

How synchronization works

When a DLO job runs, DLO does the following to back up and synchronize files:

- Backs up files that changed on the desktop.
- Makes synchronized files available to the other computers with which the desktop is synchronized.
- Downloads synchronized files that were changed on another computer and uploaded since the last DLO job ran.
- Retains all conflicting versions of files. You can then choose which version to use.

When you back up files, you can set various filters, such as which types of files to include, exclude, compress, or encrypt. When you synchronize files between computers, the filters are combined. For example, if one of the synchronized files is compressed and encrypted, all synchronized files are compressed and encrypted automatically. If the original backup selection backed up only .jpg files, the synchronized file set includes only .jpg files.

If the settings for a synchronized folder are changed after the folder is synchronized, and the folder is later unsynchronized, the folder reverts to the original backup selection settings. For example, the original backup selection backed up only .jpg files and the folder is later synchronized and set to back up

all files. If the folder is then unsynchronized, it will once again back up only .jpg files.

If the number of files that are backed up on different computers varies, DLO synchronizes the largest number of files. For example, if you back up three files on computer A and back up five files on computer B, DLO synchronizes five files.

Synchronized selections are subject to limitation by global excludes in the same manner as backup selections.

See [“About configuring global exclude filters in DLO”](#) on page 1687.

You can manage synchronization using the following options:

- Standard view, which enables you to create new synchronization sets.
- Advanced view, which enables you to modify settings for each synchronization set.

To use the synchronization feature, all synchronized computers must run the same version of the Desktop Agent and the clocks on all computers must be synchronized. In addition, the computers that run the Desktop Agent must use the same version of the Windows operating system. For example, you can synchronize data between two computers that both run Microsoft Windows Vista. You cannot synchronize data between a computer that runs Windows Vista and a computer that runs Windows XP.

Synchronizing a folder across multiple desktops

By synchronizing backed-up data, you can work on a file on any of your desktops with the assurance that you are working on the most recent version.

See [“How synchronization works”](#) on page 1779.

To synchronize a folder across multiple desktops

- 1 In the Tasks bar, under Views, click **Synchronized Selections**.
- 2 Click **Standard view**.

Desktops available for synchronization appear in the Remote Computers pane.

A desktop must have the same owner and must be backed up with the Desktop Agent to appear in the Synchronized Selections view. Only backed up folders are available for synchronization.

- 3 Select the folders that you want to synchronize.
- 4 When the Choose Local Folder dialog box appears, type or browse to the location where the synchronized files are to be stored.

- 5 Click **OK**.
- 6 Click **Save changes**.

Changing or viewing a synchronized folder

You can change or view the settings for a synchronized folder.

See [“How synchronization works”](#) on page 1779.

To change or view a synchronized folder

- 1 In the Tasks bar, under Views, click **Synchronized Selections**.
- 2 Click **Advanced view**.
- 3 Select the folder that you want to change or view.
- 4 Click **Modify**.
- 5 Configure the synchronization folder settings.
- 6 Click **OK**.

Removing a synchronized folder

When a synchronized selection is deleted, the backup files are deleted in the same manner as when source files are deleted. They are groomed away after the number of days specified in the backup selection.

See [“How synchronization works”](#) on page 1779.

To remove a synchronized folder

- 1 In the Tasks bar, under Views, click **Synchronized Selections**.
- 2 Click **Advanced View**.
- 3 Click the synchronization selection that you want to remove.
- 4 Click **Remove**.
- 5 Click **Yes** at the prompt.

Resolving conflicts with synchronized files

If a synchronized file is modified on more than one computer without updating the file with the Desktop Agent, a conflict occurs and you are prompted to determine which file version to keep. For example, a conflict occurs if the same file is modified on both your desktop computer and your laptop and your laptop is disconnected from the network. When your laptop is subsequently connected to the network, the conflict is detected.

See [“How synchronization works”](#) on page 1779.

To resolve a conflict with a synchronized file

- 1 In the Tasks bar, under Views, click **Status**.

If a conflict is identified, a resolve conflicts option appears in the Status view.

- 2 Click the **Conflicts have been found** link to open the Resolve Conflicts wizard.
- 3 Review the information on synchronization conflicts, and then click **Next**.
- 4 Select the file that contains the conflict.
- 5 Click **Open Folder**.
- 6 Manage the revisions as required.

For example, to keep an older revision, you can delete the newer revision and rename the conflicting revision back to its original name.

- 7 Click **Finish**.

About the status of the Desktop Agent

The Desktop Agent Status view provides a summary of Desktop Agent operations, including:

Table Q-97 Desktop Agent operations

Item	Description
Status	Displays the current state of Desktop Agent jobs, displays when backups will run, and summarizes the results of the last backup.
Details	This link is located just below the status summary if a backup selection has been made for a FAT drive. It provides scheduling details that are based on current Desktop Agent settings.
Show/Hide Pending Files	Hides or displays pending files. This selection toggles between Hide pending files and Show pending files when you click the link.
Network Usage	Displays the total amount of data that is stored in the network user data folder for this computer.
Local Usage	Displays the total amount of data that is stored in the desktop user data folder on this computer.
Details	This link is located just below the status summary and provides detailed information on folder usage for user data. See “Viewing usage details” on page 1783.

Starting a pending job from the Status view

From the Status view, you can run any type of pending job, such as backup, synchronization, or restore.

To start a pending job from the Status view

- 1 In the Tasks bar, under Views, click **Status**.
- 2 In the Tasks bar, under Tasks, click **Run job**.

About suspending or canceling a job

You can suspend or cancel a job if your profile allows it.

Note: The DLO administrator sets the maximum time after which a suspended job will resume.

The available options depend on the type of job being suspended:

Table Q-98 Options for suspending jobs

Type of job running	Options
Continuous	Suspend the job and resume after a specified number of minutes.
Manual	The following options are available: <ul style="list-style-type: none">■ Suspend the job and resume after a specified number of minutes.■ Cancel the job until it is started again manually.
Scheduled	The following options are available: <ul style="list-style-type: none">■ Suspend the job and resume after a specified number of minutes.■ Cancel the job until it is scheduled to run again.

Viewing usage details

The Desktop Agent Status view provides a summary of information on both local and network disk space that is used to store your data.

The following additional usage details and a grooming function are available in the Usage Details dialog box:

- Total disk space currently used on the network and desktop computer to store your backup data.
- Quotas, or maximum allowed storage space that can be used to store your data on the network and desktop computers.
- The disk space available on the network and desktop computer for storing your data.
- An option to immediately delete old revisions and deleted files.
- Links to additional information and help.

Note: The link to usage details is only available when the Desktop Agent is idle. It is not shown when a job is running.

To view usage details and groom files

- 1 In the Tasks bar, under Views, click **Status**.
- 2 In the Status pane, under Usage Summary, click **Details**.
- 3 Review the usage information, and take the appropriate actions as required:
See “[Usage Details](#)” on page 1784.

Usage Details

The **Usage Details** dialog box provides the following information.

See “[Viewing usage details](#)” on page 1783.

Table Q-99 Usage Details

Item	Description
Local	<p>Summarizes the disk space usage on the desktop computer for storing your data. The following information is provided:</p> <p>Using - The total disk space on the desktop computer currently being used to store your backup data.</p> <p>Quota - The maximum amount of disk space you can use to store your backup data on the desktop computer. The administrator sets the quota limit in the profile. However, you can modify it if your profile allows it.</p> <p>See “About modifying Desktop Agent settings” on page 1772.</p> <p>Available - The amount of free disk space available on the desktop computer for storing your data without exceeding a quota. If there is no quota, the Desktop Agent reserves a small amount of disk space so the drive does not fill completely with backup data.</p>
Network	<p>Summarizes the disk space usage on the network for storing your data. The following information is provided:</p> <p>Using - The total disk space on the network currently being used to store your backup data.</p> <p>Quota - The maximum amount of disk space you can use to store your backup data on the network.</p> <p>Available - The amount of free disk space available on the network for storing backup data for the current user without exceeding a quota.</p>
Synchronized Files	<p>Summarizes the disk space usage for storing synchronized data. The following information is provided:</p> <p>Using - The total disk space on the network currently being used to store your synchronized data.</p>

Table Q-99 **Usage Details** (*continued*)

Item	Description
<p>Remove deleted files</p>	<p>Deletes all files that are marked as deleted in your Network and desktop user data folders. The periodic maintenance cycle deletes these files after the amount of time that is specified in your assigned profile.</p> <p>Choose from the following options on the Remove Deleted Files dialog box:</p> <ul style="list-style-type: none"> ■ Remove only the deleted files that currently meet the backup selection deleted files criteria. ■ Remove all deleted files. <p>Check the Remove files from the network user data folder check box to additionally groom deleted files from the network user data folder.</p>
<p>Click here to view last job log</p>	<p>Opens the Log File Viewer.</p> <p>See “About monitoring job history in the Desktop Agent” on page 1790.</p>

Restoring files by using the Desktop Agent

You can use the Desktop Agent to restore files to the original or an alternate directory if your profile allows it. If a Desktop Agent user has more than one desktop computer running DLO, files can be selected from all available backups on each of the user’s desktops. However, those files can only be restored to the current desktop computer.

See [“About using DLO to back up Outlook PST files incrementally”](#) on page 1769.

See [“About restoring Microsoft Outlook Personal Folder files”](#) on page 1789.

If you customize NTFS permissions or directory attributes, such as compression or encryption for files or folders, you must reapply these settings after restoration.

If you disconnect from the network while the Desktop Agent is running, you may encounter a slow response when browsing the Restore view. On the Tasks menu, select Refresh to fix this problem.

DLO does not restore a file to its original location if the file is in use by another application.

If DLO encounters a file that is in use, you can do one of the following tasks to restore the file:

- Schedule a time to restore the file. The file is restored after the computer restarts. You are not notified when the file is restored.
- Log on to the desktop computer with an administrative account. Then, run a restore job to overwrite the locked file and restore it.
- Close the file in the other application.
- Restore the file to an alternate location.

To restore files

- 1 In the Tasks bar, under Views, click **Restore**.
- 2 In Show, select one of the following revision display options:

All revisions	All file revisions are displayed and available as restore selections.
Latest revision	Only the latest file revision is displayed and available as a restore selection.
Revisions modified on or after	If selected, enter a date and time after which revisions will be displayed and available as restore selections, then click OK .

- 3 Select the items you want to restore.

In some cases the Restore Search view may contain duplicate entries for the same file. If this occurs, you can select either file to restore and receive the same outcome.

When you delete a file, the backup files are retained until the file grooming process deletes them. If an original file has been deleted, but backup files are still available, the icon for the file in the restore view includes a red X to indicate the deletion of the original file.

See [“About file grooming in DLO”](#) on page 1664.

- 4 Click **Restore**.
- 5 Select the appropriate options, and then click **OK**.
See [“Restore options”](#) on page 1787.

Restore options

The **Restore** dialog box lets you determine how to handle restored files.

See [“Restoring files by using the Desktop Agent”](#) on page 1786.

Table Q-100 Restore options

Item	Description
Restore to the original folders on this computer	Restores files and folders to their original location.
Redirect the restore to an alternate folder on this computer	Restores files and folders to an alternate folder on the same computer.
Preserve folder structure	Restores the data with its original directory structure. If you clear this option, all data (including the data in subdirectories) is restored to the path you specify.
If file already exists	Determines what to do if a file that you want to restore already exists. You can choose one of the following: <ul style="list-style-type: none">■ Do not overwrite■ Prompt■ Overwrite
Restore deleted files	Lets you restore files even though the source file has been deleted.
Preserve security attributes on restored files	Preserves security information in restored files. You may need to uncheck this box to successfully restore a file if the source file security conflicts with the destination security. Unchecking this option causes the security information to be removed from the restored file.

Searching for desktop files and folders to restore

Use the Search feature to locate the files and folders that you want to restore.

To search for desktop files and folders to restore

- 1 In the Tasks bar, under Views, click **Restore**.
- 2 In the Tasks bar, under Tasks, click **Search for files to restore**.
- 3 Select the appropriate options, and then click **OK**.

See “[Search options](#)” on page 1788.

Search options

Use the **Search** dialog box to find files to restore.

See [“Searching for desktop files and folders to restore”](#) on page 1788.

Table Q-101 Search options

Item	Description
Search for file names with this text in the file name	Lets you search by file name or folder.
Modified	Lets you search for files that were modified during a specific time frame. Then specify the time frame.
Today	Lets you search for files that were modified on the current calendar day.
Within the past week	Lets you search for files that were modified in the last calendar week.
Between	Lets you search between calendar dates.
Of the following type	Lets you search for a specific type of file. Select a file type from the list provided.
Of the following size	Lets you search for a file that is equal to a specific size, at least a specific size, or at most a specific size..

About restoring Microsoft Outlook Personal Folder files

When you restore Microsoft Outlook Personal Folder (PST) files, the following differences will exist between the restored PST and the original PST:

- The file size will be different.
- Any rule that points to a folder inside a PST file will no longer work. You must edit the rule to point to the correct folder.
- Restored PST files will have Inbox, Outbox, and Sent Items folders, even if the original files did not have them.
- If you use a password for your PST file, you must reset the password after restoring your PST file.

See [“About using DLO to back up Outlook PST files incrementally”](#) on page 1769.

About restoring deleted email messages

The default behavior when deleting a message from a mail archive may differ depending on the mail application. With Lotus Notes, there is a "soft delete" feature that allows a message to be maintained in a special folder, the "Trash," for a

measured interval (default is 48-hours). After that, the message is permanently deleted. Outlook behaves in much the same manner. Deleted messages are moved to the "Deleted Items" folder but there is no time limit associated with this action. Outlook permanently deletes a message when the user empties the Deleted Items folder.

In either case, the Desktop Agent replicates the delete during the next backup operation. If a user accidentally deletes a message from a mail archive, the message must be recovered. Because there are no versions maintained for email archives, permanently deleted messages are unavailable after the time limit has expired or the user has manually emptied the folder.

About restoring files with alternate stream data

DLO now protects all of the alternate streams for a file, including security streams. If a new version of a file contains only changes to alternate stream data, the file replaces the previous version and does not affect the revision count. Only revisions with actual data changes are treated as new revisions.

FAT partitions do not use alternate data streams. If a file is restored from an NTFS partition to a FAT partition, the alternate stream data is not included in the restored file.

When a file is restored, one of the options is to preserve the security attributes on restored files. If this option is not checked, the security attributes are removed from the restored file. This option is set in the restore dialog box.

About using Backup Exec Retrieve to restore files

When DLO is configured for use with Backup Exec Retrieve, you can use a Web browser to search for DLO files and restore them to your computer. The search results include all backed up versions of the DLO file in the network user data folder that match the search criteria. You can also search based on recent activity. A unique icon distinguishes DLO files. For detailed information on restoring files using Backup Exec Retrieve, see the Symantec Backup Exec Continuous Protection Server Administrator's Guide

About monitoring job history in the Desktop Agent

When a backup, restore, or synchronization operation takes place, details of that operation are stored in log files. Log files can be viewed, searched, and saved as text files. The History View summarizes the following information and provides access to the full logs.

You can choose to view the backup history or restore history by selecting the appropriate tab at the bottom of the History window.

Table Q-102 Job History View information

Item	Description
Started	The date and time the operation started
Ended	The date and time the operation ended
Status	The status of the job, such as Active, Completed, Canceled, or Failed.
Files Transferred (Local)	The total number of files that were transferred to the desktop user data folder during the listed job.
Size Transferred (Local)	The total number of bytes of data that were transferred to the desktop user data folder during the listed job.
Files Transferred (Network)	The total number of files that were transferred to the network user data folder during the listed job. This information is only available for the backup history, not the restore history.
Size Transferred (Network)	The total number of bytes of data that were transferred to the network user data folder during the listed job. This information is only available for the backup history, not the restore history.
Errors	The number of files that failed to copy and produced errors.

See [“Viewing log files”](#) on page 1791.

See [“Searching for log files”](#) on page 1793.

Viewing log files

Log files contain information about the jobs that have run on a computer.

See [“About monitoring job history in the Desktop Agent”](#) on page 1790.

To view log files

- 1 In the Tasks bar, under Views, click **History**.
- 2 Do one of the following:
 - To view backup logs, click **Backup**.
 - To view restore logs, click **Restore**.

3 In Show, select one of the following items:

- | | |
|-----------------------|--|
| All logs | All history logs are displayed |
| All logs with errors | History logs for all jobs that generated errors are displayed. |
| Logs filtered by date | All logs that are generated after a specified date and time are displayed. Enter the date and time after which logs are to be displayed in the Filter by date dialog box and click OK . |

- 4** Click the job history entry for which you want to view the history log.
- 5** Click **View log** to open the log file viewer.
 See “[Log File Viewer options](#)” on page 1792.
- 6** If required, click **Save As** to save the log file as a text file.
- 7** Click **Close**.

Log File Viewer options

You can view details about log entries and search for specific log entries that you want to view.

See “[Viewing log files](#)” on page 1791.

Table Q-103 Log File Viewer options

Item	Description
All log files	Lets you search for log entries in all log files.
Current log file	Lets you search for log entries in the log file that you selected.
With timestamp of	Lets you search for log entries with a specific timestamp.
Today	Lets you search for log entries that were generated on the current day.
Within the last week	Lets you search for log entries that were generated in the last calendar week.
Between <date> and <date>	Lets you search for log entries that were generated between two specific dates.

Table Q-103 Log File Viewer options (continued)

Item	Description
Of the following type	Lets you search for a specific type of log entry, such as backup or restore.
With filenames containing	Lets you search for log entries that contain specific file names.
Limit search to	Lets you search only for specific types of log entries, such as informational entries or error entries.
Search	Lets you search for the log entries that match the criteria you selected.
Save As	Lets you save the log file as a text file.
Open Log File	Lets you open a log file that you saved previously.

Searching for log files

The Log File Viewer has a powerful search mechanism to help you locate the log files you want to view.

See [“About monitoring job history in the Desktop Agent”](#) on page 1790.

To search for log files

- 1 In the Tasks bar, under Views, click **History**.
- 2 In the History pane, click the **Search** link.
- 3 Enter filtering parameters:

All log files	Select this option to show all log entries in the log file viewer.
Current log file	Select this option to search only those log entries in the current log file.

With timestamp of	<p>Check the With Timestamp of box to search only those log entries within a specified time period. The options include:</p> <p>Today - Show only log files that were created today.</p> <p>Within the last week - Show all log files created in the last week.</p> <p>Between dates - Show all log files created between the dates entered.</p>
Of the following type	<p>Check the Of the following type check box to show only logs of the indicated type. You may select one of the following types:</p> <ul style="list-style-type: none">■ Backup■ Restore■ Move User■ Maintenance■ Error■ Warning
With Filenames containing	<p>Check the With Filename like check box and enter a filename, or file type. Wildcard entries are supported.</p> <p>Example: *gold.doc</p> <p>When you use wildcards you must use the '*' wildcard. For example, *.tmp returns all results with the .tmp extension while .tmp returns only files explicitly named .tmp.</p>
Informational entries only	Select Informational entries only to display only informational entries.
Error and warning entries only	Select Error and warning entries only to display both error and warning entries.
Error entries only	Select Error entries only to display only error entries.
Warning entries only	Select Warning entries only to display only entries for warnings.

4 Click **Search**.

- 5 If required, click **Save As** to save the log file as a text file.
- 6 Click **Close**.

About grooming log files

Log grooming occurs each time a log is created. Log files are not deleted until they have reached both the minimum age and maximum combined size of all log files settings. If the administrator has granted you sufficient rights in your profile, you can modify these settings in the Desktop Agent settings Options tab.

See [“Setting customized options on the Desktop Agent”](#) on page 1775.

About using DLO with other products

The following are known compatibility issues:

Table Q-104 Compatibility issues

Product	Description
Symantec Storage Exec	<p>Symantec Storage Exec is a policy-based storage resource manager for controlling file and application disk usage in Microsoft Windows environments. DLO and Storage Exec are compatible, but care must be taken to avoid conflicts between DLO backup selections and Storage Exec policies. If DLO is configured to back up a specific file type and Storage Exec is set to prevent this file type from being copied to the server, a conflict will result. DLO will try to back up the file, but the operation will fail. The DLO history log will indicate that the file failed to copy to the network user data folder.</p> <p>To prevent this conflict, DLO backup selections and Storage Exec policies must be reviewed to identify any potential conflicts. If a conflict is found, the policies must be manually revised to eliminate the conflict.</p>
WinCVS	<p>When DLO runs concurrently with WinCVS, permission denied errors are sometimes generated when checking out source. You can avoid this by excluding any directories that are named cvs when you use global excludes or backup selection excludes.</p>
Windows XP Service Pack 2	<p>If you use Windows XP with Service Pack 2, you must enable file sharing to use the browse button in the DLO Administration Console Restore view.</p>
PGP Desktop 8.1	<p>When running DLO with PGP Personal Desktop 8.1, you cannot create a mounted drive or unmount a drive which is in a DLO Backup Selection unless DLO is turned off.</p>

Troubleshooting the DLO Administration Console

If you have questions about the DLO Administration Console, review the following information to find answers.

Table Q-105 DLO Administration Console questions and answers

Question	Answer
<p>I modified an Automated User Assignment, but the change isn't reflected for existing Desktop Agent users.</p>	<p>Automated User Assignments are only used once to assign a profile and Storage Location to a new Desktop Agent user. An Automated User Assignment can be modified to change the profile and Storage Location settings, but these changes apply only to new users. Users that have already been configured are not affected by subsequent changes in the Automated User Assignment.</p> <p>This also applies to existing users who install the Desktop Agent on another desktop. The new installation uses the existing user settings and stores data in the user's existing user data folder. Automated User Assignment changes do not affect an existing user, even if the Desktop Agent installation is on a new computer.</p> <p>Settings for an existing desktop user can be changed by modifying the profile to which the user is assigned, or by reassigning that user to a new profile or Storage Location.</p> <p>See "Changing the profile for a Desktop Agent user" on page 1699.</p> <p>See "About managing Desktop Agent users" on page 1696.</p> <p>See "About Automated User Assignments" on page 1683.</p> <p>See "About DLO profiles" on page 1641.</p> <p>See "Moving Desktop Agent users to a new network user data folder" on page 1701.</p>

Table Q-105 DLO Administration Console questions and answers (*continued*)

Question	Answer
<p>A desktop user ran the Desktop Agent and received an error indicating "Unable to configure the Desktop Agent. No settings found for the current user and no automatic user assignments match." What does this mean?</p>	<p>This message means that DLO could not find the user or an Automated User Assignment that matched the user's domain and group.</p> <p>Users are added to DLO in the following ways:</p> <ul style="list-style-type: none">■ An Automated User Assignment that matches the user's domain and group assigns a profile and Storage Location to the Desktop Agent and adds the user to DLO. Check that you have created Automated User Assignments that match the Domain and Group to which the user running the Desktop Agent belongs. You can also create an Automated User Assignment that covers all domains and all groups to catch any users who might not match a more specific Automated User Assignment. Such a "catchall" Automated User Assignment would typically be set to the lowest priority.■ Users are manually added to DLO. This process requires that you assign a profile and either a Storage Location or user data folder to the new user. <p>Be sure that the user has a matching Automated User Assignment, or is added manually before the user runs the Desktop Agent.</p> <p>See "About Automated User Assignments" on page 1683.</p>

Table Q-105 DLO Administration Console questions and answers (*continued*)

Question	Answer
<p>When do I need a network user data folder, and when do I need a Storage Location?</p>	<p>Every Desktop Agent user must have a network user data folder, which is used to store backup data. Storage Locations are locations on the network where network user data folders are automatically created and maintained. They are not required if existing network shares are used to store user data.</p> <p>If you want DLO to automatically create network user data folders, use a Storage Location. When new users are added to a Storage Location, network user data folders are automatically created for them within the Storage Location.</p> <p>Alternately, if you would like to use existing network shares as network user data folders, or if you want to create network user data folders manually, then do not use Storage Locations.</p> <p>See “How to configure DLO” on page 1640.</p>
<p>I'm trying to create a Storage Location on a remote file server, and I am receiving an error indicating the MSDE Database Instance for the Desktop and Laptop Option needs to have access to the remote file server. What do I need to do?</p>	<p>To create Storage Locations on a remote file server, you must use an account that has administrative rights on the remote file server.</p> <p>You can change the account credentials that were used to create the Storage Location.</p> <p>See “Changing DLO service credentials” on page 1618.</p>
<p>I manually added a new user and assigned the user to an existing Storage Location. I don't see a new user data folder for the new user in this Storage Location. Isn't it supposed to create one?</p>	<p>User data folders are created only after the Desktop Agent is both installed on the desktop and run by the new user.</p>
<p>How do I prevent a user from backing up data?</p>	<p>You can disable the user.</p> <p>See “Enabling or disabling DLO access for a desktop user” on page 1700.</p>

Table Q-105 DLO Administration Console questions and answers (*continued*)

Question	Answer
In a backup selection, I selected to encrypt or compress my user's data. However, data that has already been backed up is not encrypted or compressed. Why is this?	DLO does not retroactively apply changes to encryption and compression settings to user data that is already backed up. Any data that is backed up after these settings have changed will use the new settings.
I would like to prevent files of specific types from being backed up. How can I set up DLO to always exclude files like *.mp3 or *.gho?	On the Tools menu, select Global Excludes . In this dialog box, you can add specific file types that will be excluded in all backup selections for all profiles.
Backups do not seem to be running for all users, or specific files are not being backed up.	If backup jobs are not running for a group of users, check the profile for these users to verify that backups are scheduled. If specific files are not being backed up, review the backup selections in the profile to verify that the files are selected for backup.
I just tried to restore a file, but it doesn't appear to have been restored.	When restoring existing files to their original location, verify that you have selected Prompt or Overwrite in the Restore dialog box to replace the file. If you select Do not overwrite, the file is not restored.
In a profile, I configured backup selections to encrypt files. Now I need to recover files for a user. Do I need an encryption key to restore this data?	As an Administrator running the DLO Administration Console, you can redirect a restore of encrypted user data to an alternate computer or location, and it will be decrypted during the restore.
I would like to restore data to a user's computer, but that user is out of the office. Do I have to wait until that user returns to the office before I can start the restore?	DLO can queue restore jobs to desktops. If the user is offline now, you can queue a restore job through the Restore view in the DLO Administration console. Another option is to restore the data to an alternate location, such as the administration computer or a network drive.

Table Q-105 DLO Administration Console questions and answers (*continued*)

Question	Answer
How can I protect open files?	<p>DLO does not protect open files. It tries to back up files when they are closed or saved. If a file cannot be backed up because it is open (for example, a Word document you are editing) it remains in the Desktop Agent's pending list. The Desktop Agent tries to back up the file at the next backup time. This also means that certain files that are opened by the operating system are not backed up. They never close when the operating system is running.</p> <p>The exception to this is protection of open PST files. The Desktop Agent is designed to protect open PST files if they are part of the profile or user's backup selections.</p> <p>Incremental backups must be enabled for open file backups of PST files.</p> <p>See "About excluding files that are always open" on page 1695.</p>
The History view in the DLO Administration Console doesn't show the most recent backup for all users.	The DLO Administration Console is automatically updated when a job runs, but not more than once per hour.

Troubleshooting the Desktop Agent

If you have questions about the Desktop Agent, review the following information to find answers.

Table Q-106 Desktop Agent questions and answers

Question	Answer
Do I have to install Backup Exec on every desktop I want to protect?	No. You must install the Desktop Agent on every desktop you want to protect. It is not necessary to install Backup Exec.

Table Q-106 Desktop Agent questions and answers (*continued*)

Question	Answer
I installed the Desktop and Laptop option, but I do not know how to install the Desktop Agent on users' computers.	<p>The Desktop Agent can be installed by running the installation program from the share where DLO is installed. Or, it can be push installed using the Backup Exec installer.</p> <p>The Desktop Agent installation program is located in a share where you installed DLO. This share will have a name in the following format:</p> <p>\\<Server>\DLOAgent.</p> <p>Using Windows Explorer, browse to this share from the desktop that you want to protect with the Desktop Agent. Run Setup.exe from this share. You must be an administrator on the desktop to install the Desktop Agent software.</p> <p>Symantec recommends that DLO administrators run the Configuration Wizard to familiarize themselves with the application.</p>
Can I install the Desktop Agent on Windows Servers or media servers?	Because the Desktop Agent is designed to protect user data rather than critical server data, it cannot be installed on Windows Servers or media servers.
I am receiving the following error while authenticating through the Desktop Agent to the media server: "Failed to Initialize database. 0x800A0E7D"	You tried to connect to the media server with an account that is not in the same domain as the media server, or in a trusted domain. For DLO to function properly, the media server must be in a Windows Domain.
I have a desktop and a laptop computer protected by the Desktop Agent. Why can't I move my laptop to a new Storage Location?	When a user has multiple computers running the Desktop Agent, all backup data is stored in the same network user data folder. If you want to move your data to a new Storage Location, you must move the entire network user data folder for all of your computers to that new location.

Table Q-106 Desktop Agent questions and answers (*continued*)

Question	Answer
<p>I am trying to synchronize files between my desktop and laptop computers, but I cannot see my other computer in the Synchronization View in the Desktop Agent.</p>	<p>To synchronize data between two computers, the same user account must be used when running the Desktop Agent on each computer. For example, the user Domain\MyUser must have backed up data on Computer A and Computer B in order for synchronization to take place between these two computers.</p> <p>If you are sure you have backed up data while running the Desktop Agent under the same user account on both of your computers, select Refresh in the Desktop Agent's Synchronization View to make the synchronization selections available. If this is not successful, Exit from the File menu and restart the Desktop Agent application.</p>
<p>What files or folders can I synchronize between my computers?</p>	<p>Any data that is backed up by a backup selection is eligible for synchronization. These backup selections may be defined by the DLO Administrator in the profile or in a backup selection that was created with the Desktop Agent.</p>
<p>I would like to share my synchronized data with my co-workers. How can I do this?</p>	<p>The Desktop and Laptop Option does not provide functionality for sharing files between users. Synchronization is designed to share files between a single user's computers.</p>

Accessibility and DLO

The following table lists keyboard navigation within tabbed dialog boxes:

Table Q-107 Keyboard Navigation within Tabbed Dialog Boxes

Keyboard input	Result
<p>CTRL+PAGE DOWN or CTRL+TAB</p>	<p>Switches to the next tab and displays the page.</p>
<p>CTRL+ PAGE UP</p>	<p>Switches to the previous tab and displays the page.</p>
<p>RIGHT ARROW or LEFT ARROW</p>	<p>When the focus is on a tab selector, chooses the next or previous tab in the current row and displays the page.</p>

The following table lists keyboard shortcuts for the Administration Console:

Table Q-108 Keyboard Shortcuts Unique to Backup Exec Desktop and Laptop Option Administration Console

Accelerator	Mnemonic	Result
ALT	F	The File menu expands. From the File menu, you can create new profiles and Storage Locations, and add users.
ALT	E	The Edit menu expands. From the Edit menu, you can restore files, search for files to restore, manage alerts, and delete items.
ALT	V	The View menu expands. From the View menu, you can change the information that displays on the screen.
ALT	N	The Network menu expands. Use the Network menu to work with administrator accounts, connect to the DLO Administration Servers on the network, or to reconnect to a local DLO Administration Server.
ALT	T	The Tools menu expands. Use the Tools menu to set global excludes, access all DLO wizards, and manage service credentials.
ALT	W	The Window menu expands. Use the Window menu to move to a new window or view.
ALT	H	The Help menu expands. Use the Help menu to access documentation and various Symantec Web sites.

The following table lists keyboard shortcuts for the Desktop Agent:

Table Q-109 Keyboard Shortcuts Unique to Desktop and Laptop Option Desktop Agent

Accelerator	Mnemonic	Result
ALT	F	The File menu expands. From the File menu, you can minimize or exit the Desktop Agent.
ALT	V	The View menu expands. From the View menu, you can change the information that displays on the screen.
ALT	K	The Tasks menu expands. Use the Tasks menu to run a job or refresh the view.

Table Q-109 Keyboard Shortcuts Unique to Desktop and Laptop Option Desktop Agent (*continued*)

Accelerator	Mnemonic	Result
ALT	O	The Tools menu expands. Use the Tools menu to reset dialog boxes and accounts.
ALT	H	The Help menu expands. Use the Help menu to access the online help for the Desktop Agent.

Symantec Backup Exec Intelligent Disaster Recovery Option

This appendix includes the following topics:

- [About the Intelligent Disaster Recovery Option](#)
- [Requirements for using IDR](#)
- [About installing the IDR Option](#)
- [About preparing computers for IDR](#)
- [About the the Intelligent Disaster Recovery Configuration Wizard](#)
- [About creating and updating recovery media](#)
- [Copying the disaster recovery files](#)
- [Preparing IDR media by using other media servers](#)
- [About preparing to recover from a disaster by using IDR](#)
- [About the Intelligent Disaster Recovery Wizard](#)
- [About using IDR with the Central Admin Server Option](#)
- [About using IDR with Veritas Storage Foundation for Windows](#)
- [Best Practices for IDR](#)

About the Intelligent Disaster Recovery Option

The Symantec Backup Exec 2010 Intelligent Disaster Recovery Option (IDR) enables you to quickly and efficiently recover Windows computers after a hard drive failure. The IDR wizards guide you in preparing for disaster recovery and in recovering a local or remote computer to its pre-disaster state.

Before you can recover computers, you must prepare for a disaster by performing the following steps in the order listed:

- On the media server, use the Intelligent Disaster Recovery Configuration Wizard to specify a location where a copy of the computer-specific disaster recovery file (*.dr file, where the asterisk represents the name of the computer being protected) will be stored.

The Intelligent Disaster Recovery Configuration Wizard guides you through setting an alternate data path for the *.dr file. The default data path for the *.dr file is on the media server's hard drive, but Symantec recommends that you specify an alternate data path to store another copy of the *.dr file in case the media server's hard drive is damaged.

- Run full backups of the hard drives on the computers to be protected. Include System State for Windows 2000 and Windows XP computers, and Shadow Copy Components for Windows Server 2003/Windows Vista/Windows Server 2008/Windows Server 2008 R2/Windows 7 computers. Do not exclude any files from the full backups; otherwise the *.dr file will not be created.

Backup Exec creates the *.dr file during a full backup and stores it in the default and alternate storage locations. Catalog entries from subsequent backups are added to the *.dr file as these backups are completed.

- Run the Intelligent Disaster Recovery Preparation Wizard to create bootable media for each computer.

The Intelligent Disaster Recovery Preparation Wizard guides you through the preparation of the bootable media that you use to recover protected computers. The Intelligent Disaster Recovery Preparation Wizard also lets you copy the *.dr file to any location.

After you have performed these steps for each computer you want to protect, you are prepared to recover those computers using any of the following recovery methods:

- Restore a media server (Backup Exec server) using a locally attached storage device.
- Restore a media server (Backup Exec server) using a remote backup-to-disk folder.

- Restore a Windows computer by moving the media and the storage device to the computer being restored, and then restoring the computer through the locally attached storage device.
- Restore a remote Windows computer using a network connection to the media server

See [“About the the Intelligent Disaster Recovery Configuration Wizard”](#) on page 1810.

See [“About creating and updating recovery media”](#) on page 1813.

Requirements for using IDR

The following items are required before you use IDR:

- Symantec Backup Exec 2010.
- The Symantec Backup Exec Remote Agent for Windows or Backup Exec must be installed on any remote computers to be protected with IDR.
- Sufficient hard drive space to hold an entire Windows installation (600 MB to 2 GB).

Note: Media servers can be recovered using remote backup-to-disk folders. Mixed media loaders are not supported for local IDR.

- Encryption key files for all hard drives that you encrypt with Windows BitLocker Drive Encryption (Windows Vista/Windows Server 2008/Windows Server 2008 R2/Windows 7 only).
- A third party ISO 9660-compliant CD burning application to burn the IDR-created bootable CD image to a CD.
- A writable or rewritable CD device.

See [“About requirements for running the Intelligent Disaster Recovery Preparation Wizard”](#) on page 1814.

See [“About using a trial version of the IDR Option”](#) on page 1808.

About installing the IDR Option

You can install IDR as an option during the initial installation of Backup Exec 2010, or it can be installed later.

The Remote Agent must be purchased separately from the IDR Option, and must also be installed on any remote computer you want to protect with IDR. The

Remote Agent is a system service that runs on remote servers and enhances backup and restore performance. It is required for IDR functionality.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

See [“Push-installing the Remote Agent and Advanced Open File Option to remote computers”](#) on page 134.

See [“About installing the Remote Agent for Windows Systems”](#) on page 139.

See [“About the the Intelligent Disaster Recovery Configuration Wizard”](#) on page 1810.

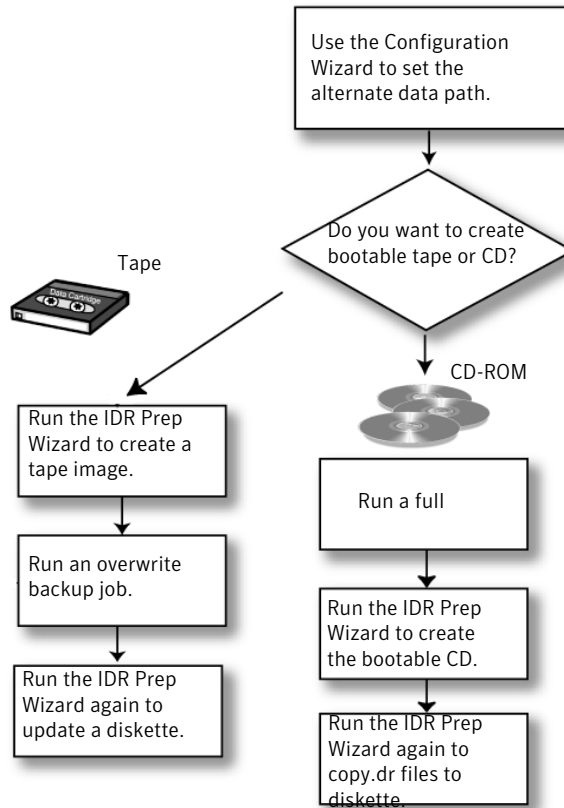
About using a trial version of the IDR Option

The IDR Option can be installed and evaluated for up to 60 days or until Backup Exec is licensed. However, you must run a full backup job of the critical volumes, along with system state and shadow copy components. You also must recreate the IDR recovery media after the IDR Option has been installed.

About preparing computers for IDR

The key to successfully recovering computers after a disaster is to carefully and properly prepare those computers for a disaster.

Figure R-1 How to prepare computers for IDR



Preparing computers for IDR involves the following:

- Using the Intelligent Disaster Recovery Configuration Wizard to determine the alternate location where a copy of the *.dr file will be stored.
- Performing full backup jobs on the computers to be protected.
- Creating the bootable recovery media using the Intelligent Disaster Recovery Preparation Wizard.

You can create the following types of bootable media with the Intelligent Disaster Recovery Preparation Wizard:

- CD-R (CD-Recordable) or CD-RW (CD-Rewritable).
- Bootable tape (the tape device must support bootable specifications)

Consider what type of Windows computer is being protected, the available hardware, and the system BIOS when selecting the type of bootable media to create. You can also combine media to make updating the *.dr files easier. If you are using bootable CD-R or CD-RW, or tape, you can still back up the *.dr files to any location using the Intelligent Disaster Recovery Preparation Wizard so that you can easily update them when required.

Use the following table to decide which type of media to use.

Table R-1 Bootable media comparison chart

Type of Media	Advantages	Disadvantages
CD-R, CD-RW	<ul style="list-style-type: none">■ Can also be used to protect remote Windows computers on the network.■ Can create bootable CD images for remote computers.	<ul style="list-style-type: none">■ Requires a BIOS that supports booting from a CD.■ Requires a CD burner.
Bootable tape	<ul style="list-style-type: none">■ Does not require a CD burner.	<ul style="list-style-type: none">■ Requires a BIOS that supports booting from a SCSI CD and a bootable tape device that emulates a SCSI CD drive.■ Cannot create bootable tape images for remote computers.

About the the Intelligent Disaster Recovery Configuration Wizard

The Intelligent Disaster Recovery Configuration Wizard appears on the Backup Exec Getting Started section of the Home view. This wizard prompts you to set an alternate data path for the computer-specific disaster recovery file, called a *.dr file.

The asterisk (*) represents the name of the computer for which the file was created. A *.dr file contains specific information for the computer being protected, including the following:

- Hardware-specific information for each computer, such as hard disk partition information, mass storage controller information, and Network Interface Card information.

- A list of catalog entries that identify the backup media that is used to recover the computer.
- For Microsoft Vista/Windows Server 2008/Windows Server 2008 R2/Windows 7, Windows Automated System Recovery (ASR) configuration information file (asr.xml). The ASR file is necessary to recreate partitions on Windows Vista/Server 2008/Windows Server 2008 R2 computers during the recovery process.
- For Windows XP and Windows Server 2003 computers, Windows Automated System Recovery (ASR) configuration information files (asr.sif and asrpn.sif). The ASR files are necessary to recreate partitions on Windows XP and Windows Server 2003 computers during the recovery process.

See [“Running the Intelligent Disaster Recovery Preparation Wizard”](#) on page 1816.

See [“Recovering a computer by using the Intelligent Disaster Recovery Wizard”](#) on page 1831.

See [“About creating and updating recovery media”](#) on page 1813.

See [“Creating a bootable tape image”](#) on page 1819.

See [“How to back up data”](#) on page 329.

About manually editing the default data paths for the *.dr files

If you did not use the Intelligent Disaster Recovery Configuration Wizard to set an alternate data path for the *.dr files, you can set it manually.

See [“Manually editing the default data paths for the *.dr files”](#) on page 1812.

Copies of the *.dr files, which contain the computer-specific information for the computer being protected, are necessary to automate the recovery of an IDR-protected computer.

Backup Exec automatically creates the *.dr file during a backup and stores it in the Disaster Recovery Data Path default location on the media server’s hard drive, which is

```
C:\Program Files\Symantec\Backup Exec\IDR\Data\
```

Symantec recommends that you do not change the default.

You can also specify an alternate location where a second copy of the *.dr file is stored so that the *.dr file is available even if the media server has been damaged. It is recommended that the alternate location be on another computer or on a different physical drive than the default location, and that it be a mapped network drive.

See [“Intelligent Disaster Recovery data paths”](#) on page 1812.

Manually editing the default data paths for the *.dr files

Use the following steps to edit the default data paths for the *.dr files.

See [“About manually editing the default data paths for the *.dr files”](#) on page 1811.

To edit the default data paths for the *.dr files

- 1 From the **Tools** menu, click **Options**.
- 2 Under **Settings**, click **Intelligent Disaster Recovery**.
- 3 Enter the paths where you want to store the .dr file.

Intelligent Disaster Recovery data paths

You can edit the default data path and the alternative data path where you want to store the Intelligent Disaster Recovery *.dr file.

Symantec recommends that you do not change the default data path.

See [“About manually editing the default data paths for the *.dr files”](#) on page 1811.

Table R-2 Data path storage locations for the *.dr file

Item	Description
Data path	Enter a directory path where a copy of the *.dr file for the protected computer will be stored. Backup Exec automatically creates the *.dr file during a backup and stores it in the default location on the media server's hard drive, which is C:\Program Files\Symantec\Backup Exec\IDR\Data\<computer name>.dr.

Table R-2 Data path storage locations for the *.dr file (*continued*)

Item	Description
Alternate path	<p>Enter an alternate directory path where a copy of the *.dr files for the protected computers will be stored. Backup Exec automatically creates or updates the *.dr file during a backup and stores it in the specified location during a backup.</p> <p>It is recommended that you specify an alternate data path that is not on the media server, or is on a different physical drive than the default location. During a recovery, you can copy the *.dr file from the alternate path to any location to recover the target computer if the media server's hard drive is unavailable.</p> <p>To use a remote computer's hard drive as the alternate data path, establish a valid connection to the remote computer. Specify a UNC path as the alternate path, and then check the directory to make sure the *.dr files were copied.</p> <p>When using Backup Exec's Remote Administrator, do not specify a floppy drive (A:, B:) as the alternative data path.</p>

About creating and updating recovery media

Before running the Intelligent Disaster Recovery Preparation Wizard to create or update the recovery media, run a full backup of the hard drive (unless you are creating bootable tape media).

See [“Creating a bootable tape image”](#) on page 1819.

The *.dr file is created when a full backup of the entire hard drive is run.

Note: If you exclude files from backups, the *.dr file is not created.

After the *.dr file is created, Backup Exec automatically updates it with data from all subsequent backups (except copy backups) in its default location on the computer and in the alternate location you specified. You can view the default locations from Tools > Options > Intelligent Disaster Recovery.

For each backup set that is backed up, an alert appears reminding you to use the Intelligent Disaster Recovery Preparation Wizard to back up the *.dr files to any location. If you use a diskette, you should label it and then store it with the rest of the disaster recovery media.

If you do not run a full backup before running the Intelligent Disaster Recovery Preparation Wizard, you can still create all the media, but the computer-specific *.dr file will not contain the catalog entries for the backup sets, and during the recovery phase, you will have to manually search for and restore the backup sets necessary to recover the computer.

Note: For the local media server, update the bootable media after every successful full backup, or whenever you patch or upgrade your operating system software. Symantec also recommends that you update the bootable media when you reconfigure or update your storage drivers or network drivers. For remote computers, you do not need to create or update bootable media until a disaster occurs, as long as a *.dr file for the remote computer is available on the media server.

The bootable media contains the system files necessary to make a failed Windows computer operational after a disaster. Create a new bootable image whenever hardware, SCSI drivers, or storage device drivers change on the computer that is protected.

Prepare and test bootable media before a disaster to make sure that the media was prepared correctly.

See [“Recovering a computer by using the Intelligent Disaster Recovery Wizard”](#) on page 1831.

The bootable media also contains a text file called <computer name>-diskconf.txt, which contains information about the computer’s hard disk layout.

See [“Creating a bootable CD image”](#) on page 1817.

See [“Creating a bootable tape image”](#) on page 1819.

See [“Creating the Intelligent Disaster Recovery nonbootable CD image only”](#) on page 1822.

About requirements for running the Intelligent Disaster Recovery Preparation Wizard

Before you run the Intelligent Disaster Recovery Preparation Wizard, run a full backup of the hard drive before creating the boot and recovery media (unless you are creating a bootable tape image).

When running full backups for IDR preparation, do the following:

- Make sure that volumes (C, D, etc.) have been backed up. The *.dr files are not created or updated if only individual directories are backed up.
- For Windows 2000/XP, back up System State.
- For Windows Server 2003/Vista/Server 2008, back up Shadow Copy components.
- If utility partitions are present on the computer, select them for backup. See [“About the Computer Name node in the backup selections list”](#) on page 282.
- Do not include or exclude files from the backup using the Advanced File Selection feature.
- Ensure that if the computer is a remote computer, a compatible version of the Remote Agent has been installed on it. To determine if the Remote Agent is installed on a remote computer, from Windows Explorer, right-click the remote server and then from the shortcut menu, click Properties. The status of the Remote Agent, if installed, appears.
- If you install Backup Exec into an existing SQL instance, Symantec recommends that you periodically back up the SQL system databases using the optional SQL Agent.

About running the Intelligent Disaster Recovery Preparation Wizard

The Intelligent Disaster Recovery Preparation Wizard guides you through the process of creating the bootable media that is used to recover protected computers. You can also use the Intelligent Disaster Recovery Preparation Wizard to copy disaster recovery *.dr files to any location. For example, you can use local drives, network drives, USB thumb drives, and so on. It can also guide you through process of creating a non-bootable disaster recovery CD image. A non-bootable disaster recovery CD can be used to run the Intelligent Disaster Recovery Wizard if the manufacturer of the computer that you protect requires you to start the computer using the manufacturer’s bootable CD.

For example, if you are running a RAID system on a Dell, HP or other type of computer, you may be required to start the computer with the manufacturer’s bootable CD in order to install the required RAID drivers.

When running the Intelligent Disaster Recovery Preparation Wizard, the local computer on which the IDR Option is installed is used by default to create or update the disaster recovery media. However, if the computer does not have the IDR Option installed locally, select Choose a media server that has the IDR Option installed to select another media server on which the IDR Option is installed in order to create or update the media.

See [“Running the Intelligent Disaster Recovery Preparation Wizard”](#) on page 1816.

See [“About creating and updating recovery media”](#) on page 1813.

See [“About requirements for running the Intelligent Disaster Recovery Preparation Wizard”](#) on page 1814.

See [“Preparing IDR media by using other media servers”](#) on page 1825.

See [“Creating a bootable CD image”](#) on page 1817.

See [“Creating a bootable tape image”](#) on page 1819.

See [“Creating the Intelligent Disaster Recovery nonbootable CD image only”](#) on page 1822.

See [“Copying the disaster recovery files”](#) on page 1824.

Running the Intelligent Disaster Recovery Preparation Wizard

Use the following steps to run the Intelligent Disaster Recovery Preparation Wizard.

To run the Intelligent Disaster Recovery Preparation Wizard

- ◆ On the **Tools** menu, click **Wizards > Intelligent Disaster Recovery Preparation Wizard**.

About creating recovery media after a disaster

If a disaster occurs on a computer before you create the recovery media for it, you can still create recovery media if you made a full backup of the computer before the disaster.

Note: For remote computers, this feature is available only if the Remote Agent version 10.0 or higher is installed on the remote computer.

When you create a full backup of a computer, IDR creates a *.dr file that contains system and catalog information. IDR uses the *.dr file to create the recovery media needed to recover the computer.

If a disaster occurs on the local media server, you can create recovery media for it after a disaster if you have another media server and a copy of the *.dr file from the local media server in an alternate location. Also, you can use the Remote Administrator to recover the local media server.

See [“Creating a bootable CD image”](#) on page 1817.

See [“Creating a bootable tape image”](#) on page 1819.

See [“Creating the Intelligent Disaster Recovery nonbootable CD image only”](#) on page 1822.

About creating a bootable CD image

The Intelligent Disaster Recovery Preparation Wizard also lets you copy the *.dr file to any location on a regular basis and to recreate the bootable CD image whenever hardware, SCSI drivers, or tape drivers change on the computer.

In addition to the requirements for running the Intelligent Disaster Recovery Preparation Wizard, note the following:

- Backup Exec does not include support for burning the disaster recovery CD image to supported CD-R and CD-RW drives. To write the CD image to a CD, use a third party ISO 9660-compliant application. You should verify the image created by third-party CD burning software before you need it for disaster recovery.
- CD-R is the recommended media for creating a bootable CD image. If CD-RW media is used, the CD drive must have MultiRead ability; otherwise, inconsistent behavior may occur when running IDR. Test the media with the CD drive before relying on it for disaster recovery.
- Prior to a disaster, test the bootable CD to ensure that the computer can start from it.

See [“Creating a bootable CD image”](#) on page 1817.

Creating a bootable CD image

Use the Intelligent Disaster Recovery Preparation Wizard to create a bootable CD image.

See [“About creating a bootable CD image”](#) on page 1817.

To create a bootable CD image

- 1 Verify that the computer to be protected has been backed up using the full backup method.
- 2 On the **Tools** menu, click **Wizards > Intelligent Disaster Recovery Preparation Wizard**.

By default, the Intelligent Disaster Recovery Preparation Wizard uses this computer to create a bootable CD image. If this computer does not have the IDR Option installed locally, select another media server on which the IDR Option is installed to create a bootable CD image.

3 Do one of the following:

To use this computer to create the bootable CD image On the Welcome screen, click **Next**.

To use another computer to create the bootable CD image Click **Choose a media server that has the IDR Option installed**. See [“Preparing IDR media by using other media servers”](#) on page 1825.

- 4** On the **Create IDR Boot Media** screen, under **Create**, select **Bootable CD Image** for use with CD Writers (ISO 9660) and then click **Next**.
- 5** On the **Starting CD Image Creation** screen, click **Next**.
- 6** In the **Available Computers** pane, select the computers for which you want to create bootable media, and then click the right arrow to move the computer to the **Selected Computers** pane.
- 7** If a computer you want to protect does not appear in the **Available Computers** pane, click **Browse** to search for the computer. You can also type the name of the computer in the field next to the **Add** button and then click **Add**.
- 8** Click **Next**.
- 9** On the **Select Location for CD Image** screen, type a path where you want to store the bootable CD image until you burn a CD, or click **Browse** to navigate to a storage location.
- 10** Click **Next**.

- 11** On the **Select Path to Windows Operating System Installation Files** screen, type a path to where copies of the operating system setup files are located. You can also click **Browse** to navigate to the location.

The Windows operating system you specify in the installation path must match the Windows version and language of the computer being protected.

You can enter one of the following:

If the files are on a CD Type the CD drive letter.
CD

If the files are stored on a network or the local computer's hard drive Type the path to files.

If an .ISO image for operating system CD is available Specify the path to that image.

- 12** Click **Next**.

The wizard begins creating the bootable image.

- 13** When the bootable CD image is complete, click **Next**.

- 14** Click **Finish**.

- 15** Use a third-party CD burning software tool to burn the bootable CD image to a CD as a CD disk image. Do not burn the CD image as a file.

About updating a bootable CD image

If you initially created a bootable image on CD and then you change the media server's hardware, you must create a new bootable CD image. Run another full backup of the protected computer. After backing up the media server, run the Intelligent Disaster Recovery Preparation Wizard again to create a new bootable CD image.

See [“Creating a bootable CD image”](#) on page 1817.

Creating a bootable tape image

This option is available only for the local media server with a compliant bootable tape device.

Use the Intelligent Disaster Recovery Preparation Wizard to create a bootable tape image, and then run an overwrite backup job so that the image is written to the tape.

The Intelligent Disaster Recovery Preparation Wizard also enables you to update the *.dr file on a regular basis, and to recreate the bootable tape image whenever hardware, SCSI drivers, or tape drivers change on the computer. Prior to a disaster, test the bootable tape to ensure that the computer can start from it. Follow the tape drive manufacturer's documentation for testing tape drive booting capability.

Before starting this procedure, review the requirements for running the Intelligent Disaster Recovery Preparation Wizard.

See [“About requirements for running the Intelligent Disaster Recovery Preparation Wizard”](#) on page 1814.

See [“Recovering a computer by using the Intelligent Disaster Recovery Wizard”](#) on page 1831.

See [“Updating a bootable tape image”](#) on page 1821.

To create a bootable tape image

A bootable tape drive and its driver must be detected by the Intelligent Disaster Recovery Preparation Wizard before the option to create a bootable tape image is displayed.

- 1 On the **Tools** menu, click **Wizards > Intelligent Disaster Recovery Preparation Wizard**.

By default, the Intelligent Disaster Recovery Preparation Wizard uses this computer to prepare the bootable tape image. If this computer does not have the IDR Option installed locally, select another media server on which the IDR Option is installed to create the boot image.

- 2 Do one of the following:

To use this computer to create the bootable tape image	On the Welcome screen, click Next .
--	--

To use another computer to create the bootable tape image	Click Choose a media server that has the IDR Option installed . See “Preparing IDR media by using other media servers” on page 1825.
---	--

- 3 Under **Create**, select **Bootable Tape Image for use with bootable tape devices**, and then click **Next**.

- 4 Read the **Starting Tape Image Creation** screen, and then click **Next**.
- 5 Type a path to store the bootable image until you perform a full backup, or click **Browse** to navigate to a storage location.
- 6 Click **Next**.
- 7 Type a path to the location of the Windows operating system files, or click **Browse** to navigate to the location.
- 8 Click **Next**.
- 9 When the bootable tape image has completed, click **Next**.
- 10 To view the computer's hard disk configuration, click **View Disk Configuration**.
- 11 Click **Finish**.
- 12 Run an overwrite backup job so that the bootable image is written to the tape.

Updating a bootable tape image

To update a bootable tape image, use the following procedure.

Note: A bootable tape drive and its driver must be detected by the Intelligent Disaster Recovery Preparation Wizard before the option to create a bootable tape image is displayed.

To update the bootable tape image

- 1 On the **Tools** menu, click **Wizards > Intelligent Disaster Recovery Preparation Wizard**.

By default, the Intelligent Disaster Recovery Preparation Wizard uses this computer to update the bootable tape image. If this computer does not have the IDR Option installed locally, select another media server on which the IDR Option is installed to update the boot image.

- 2 Do one of the following:

To use this computer to update the bootable tape image

On the Welcome screen, click **Next**.

To use another computer to update the bootable tape image

Click **Choose a media server that has the IDR Option installed**.

See [“Preparing IDR media by using other media servers”](#) on page 1825.

The **Create IDR Boot Media** screen appears.

- 3 Under **Create**, click **Bootable Tape Image for use with bootable tape devices**, and then click **Next**.

The **Starting Tape Image Creation** screen appears.

If you have previously prepared a bootable image for tape, the **Disaster Recovery Image Found** screen appears.

- 4 Click **Delete the existing image** to write the new bootable image to the bootable tape when the first overwrite backup job runs.
- 5 Continue to follow the prompts until the wizard is complete.
- 6 When the bootable image is complete, run an overwrite backup job so that the image is written to the tape.

See [“How to back up data”](#) on page 329.

Creating the Intelligent Disaster Recovery nonbootable CD image only

You can create a non-bootable disaster recovery CD image to complete the disaster recovery media set if the computer being protected has a bootable tape or CD image already created, or if the boot image has just been updated. The nonbootable

CD image, named `nonbootable_idrcd.iso`, includes the necessary drivers, the Intelligent Disaster Recovery Wizard, and the computer-specific `*.dr` file.

Note: You must burn the nonbootable CD image to a CD as a disc image using third-party software. Do not burn the CD image as a file.

See [“How to back up data”](#) on page 329.

To create an Intelligent Disaster Recovery nonbootable CD image only

1 On the **Tools** menu, click **Wizards > Intelligent Disaster Recovery Preparation Wizard**.

By default, the Intelligent Disaster Recovery Preparation Wizard uses this computer to create the nonbootable CD image. If this computer does not have the IDR Option installed locally, select another media server on which the IDR Option is installed to create the nonbootable CD image.

2 Do one of the following:

To use this computer to create the nonbootable CD image

On the **Welcome** screen, click **Next**.

To use another computer to create the Intelligent Disaster Recovery nonbootable CD image

Click **Choose a media server that has the IDR Option installed**. See [“Preparing IDR media by using other media servers”](#) on page 1825.

3 Under **Create**, click **Nonbootable disaster recovery CD Image** and then click **Next**.

4 On the **Starting nonbootable CD Image Creation** screen, click **Next**.

5 Type a path to a storage location for the nonbootable CD image, or click **Browse** to navigate to a storage location.

6 Click **Next**.

The wizard begins creating the nonbootable image.

7 When the nonbootable CD image is complete, click **Next**.

8 Click **Finish**.

- 9 Use a third-party CD burning software tool to burn the nonbootable CD image to a CD as a CD disk image. Do not burn the CD image as a file.
- 10 Label the nonbootable disaster recovery CD image appropriately and then store it with your bootable IDR CD.

Copying the disaster recovery files

Symantec recommends that you copy the disaster recovery information *.dr files you created during the backup process to an alternate safe location.

The *.dr files reside in the \Program Files\Symantec\Backup Exec\IDR\Data directory, on the media server where IDR is installed.

To copy the disaster recovery files

- 1 Run a full backup of the target computer.

When running full backups for IDR preparation:

- Ensure that full backups of each hard disk volume (C:, D:, etc.) were made. The *.dr files are not created or updated if you back up only individual directories.

In addition, do the following:

- For Windows 2000/XP, back up System State.
- For Windows Server 2003/Vista/Windows Server 2008/Windows Server 2008 R2/Windows 7, back up the Shadow Copy and System State components.
- If utility partitions are present on the computer, select them for backup. See [“About the Computer Name node in the backup selections list”](#) on page 282.
- Do not include or exclude files from the backup using the Advanced File Selection feature.

- 2 On the **Tools** menu, click **Wizards > Intelligent Disaster Recovery Preparation Wizard**.

By default, the Intelligent Disaster Recovery Preparation Wizard uses this computer to copy the disaster recovery information to an alternate location. If this computer does not have the IDR Option installed on it, select another media server on which the IDR Option is installed to copy the disaster recovery files.

3 Do one of the following:

To use this computer to copy the disaster recovery information files On the **Welcome** screen, click **Next**.

To use another computer to copy the disaster recovery information files On the **Welcome** screen, click **Choose a media server that has the IDR Option installed**.
See [“Preparing IDR media by using other media servers”](#) on page 1825.

4 Under **Copy**, click **Disaster recovery information (.dr) files**, and then click **Next**.

5 Select the computer or computers for which you want to copy the disaster recovery information files.

6 Enter a destination folder name in the **Copy To** field or click **Browse** to navigate to the destination folder, and then click **Next**.

A destination folder can be on local drives, network drives, and USB thumb drives.

7 On the **Copy Disaster Recovery Information Files** screen, click **Next**.

8 When the **Finish** screen appears, the disaster recovery information files are copied.

See [“Preparing IDR media by using other media servers”](#) on page 1825.

Preparing IDR media by using other media servers

When running the Intelligent Disaster Recovery Preparation Wizard, the local computer where the IDR Option is installed is used by default to create or update the disaster recovery media. However, if the computer does not have the IDR Option installed locally, you can select another media server on which the IDR Option is installed in order to create or update the media.

To perform disaster recovery preparations on a different media server

- 1 On the **Tools** menu, click **Wizards> Intelligent Disaster Recovery Preparation Wizard**.
- 2 On the Intelligent Disaster Recovery Preparation Wizard **Welcome** screen, click **Choose a media server that has the IDR option installed**, and then click **Next**.
- 3 Click **Browse** to browse the network and select a media server that has the IDR Option installed.
- 4 Enter the credentials required to access the media server.
See “[Media server logon credential options](#)” on page 1826.
- 5 Click **Next** to continue preparing disaster recover media.

The media server that you select is the computer that actually creates the media.

See “[Recovering a computer by using the Intelligent Disaster Recovery Wizard](#)” on page 1831.

See “[Performing an automated restore by using the Intelligent Disaster Recovery Wizard](#)” on page 1832.

See “[Performing a manual restore by using the Intelligent Disaster Recovery Wizard](#)” on page 1841.

Media server logon credential options

Enter the credentials required to access the media server.

See “[Preparing IDR media by using other media servers](#)” on page 1825.

Table R-3 Media server logon credential options

Item	Description
Media Server name	Indicates the name of the remote media server that is selected to run the restore job.
User name	Indicates the user name that has administrator rights to the remote media server.
Password	Indicates the password required for access.

Table R-3 Media server logon credential options (*continued*)

Item	Description
Domain	Indicates the domain in which the remote media server is a member. If the media server is in a workgroup, leave this field blank.

About preparing to recover from a disaster by using IDR

When a disaster occurs, you can use IDR to return the computer to its pre-disaster state. Recovering a computer is a multi-step process that involves both manual and automatic processes. To recover a computer, you must follow these steps in order:

Caution: Disconnect any storage area network (SAN) or cluster that is attached to the computer being recovered; otherwise, the hard drives on those computers may also be repartitioned and reformatted.

Table R-4 Preparing to recover from a disaster by using IDR

Step	Description
Step 1	Plan any hardware changes to the computer to be recovered. See “About changing hardware in the computer to be recovered” on page 1829.
Step 2	Review additional requirements for IBM computers if the computer to be recovered is an IBM computer. See “About using IDR to recover IBM computers” on page 1830.
Step 3	Start the computer using the bootable media created with the Intelligent Disaster Recovery Preparation Wizard to start the recovery process.

Table R-4 Preparing to recover from a disaster by using IDR (*continued*)

Step	Description
Step 4	<p>Use the Intelligent Disaster Recovery Wizard to restore the computer to its pre-disaster state and restore the data files.</p> <p>See “Recovering a computer by using the Intelligent Disaster Recovery Wizard” on page 1831.</p>

Note: Boot managers, such as System Commander or the OS/2 Boot Manager, cannot be restored with IDR. Boot managers are usually installed at a very low level that Backup Exec cannot protect. For example, the OS/2 Boot Manager resides in its own hard drive partition that Backup Exec cannot access. Because of the many different boot managers available, you may not be able to restart the computer after an IDR recovery, even though the operating system was restored. If this happens, re-installing the boot manager should fix the problem.

Before recovering the computer, note the following:

- There must be enough disks to restore all of the critical system disks. A disk is considered critical if it is required for the computer to start successfully.
- The storage capacity of each critical disk must be greater than or equal to the corresponding original disk. Disk geometries, which may also be called disk parameters, must be compatible.
- Floppy and CD devices cannot be external PC-card drives. Because external PC-card devices are not supported during the GUI-mode Windows Setup phase, they cannot be used to access data, and recovery cannot be completed.
- If a *.dr file is unavailable for the computer being restored, you can still use IDR to recover the computer, but you must first manually restore the non-critical partition information, including utility partitions.
- IDR does not recover software mirrored volumes or any kind of software RAID with the auto-partitioning feature. You must manually apply the mirror with the Disk Manager. In addition, hardware RAID components must be set up before you perform the disaster recovery.

See [“About changing hardware in the computer to be recovered”](#) on page 1829.

See [“Recovering a computer by using the Intelligent Disaster Recovery Wizard”](#) on page 1831.

About changing hardware in the computer to be recovered

During the creation of the *.dr files, IDR creates a device driver database on the media server running IDR. This database contains the drivers required for the various hard drives and network interface cards installed in each of the computers. If you experience a hard drive or network interface card failure in a particular computer and you replace either of the failed components with the same type found in your other computers, IDR automatically installs the correct device driver during the recovery.

You can also use IDR to recovery a computer that is no longer functioning. For example, if the computer's main system board fails, you can restore the computer's data after you replace the system board, even if the new board is a different model or contains multiple processors.

If you plan to change the hardware in the computer being recovered, note the following:

- **Hard drives.** Any hard drives you replace should be the same size or larger than the original drives, and the number of hard drives you replace must equal or exceed the number of hard drives that were in the original computer configuration; otherwise repartitioning problems may occur.
- **System boards.** After you replace a faulty system board and after you use IDR to recover the computer, you must use the system board manufacturer's driver CD to re-install additional functionality such as onboard sound and video.
- **Network interface cards.** If you change the network interface card in the computer you are recovering, you must install the necessary network drivers. Without the network drivers, you cannot access the network if you want to use a remote media server or remote backup-to-disk folders to recover the computer. After you complete the recovery, you must install new network interface card drivers that match the network card presently in the computer. The backup sets that you use to recover the computer contains the original network interface card drivers for the faulty network interface card that you replaced.

See [“Recovering a computer by using the Intelligent Disaster Recovery Wizard”](#) on page 1831.

See [“About encrypted backup sets and the Intelligent Disaster Recovery Wizard”](#) on page 1831.

See [“Performing a manual restore by using the Intelligent Disaster Recovery Wizard”](#) on page 1841.

About using IDR to recover IBM computers

To recover an IBM computer equipped with an IBM ServeRAID card, perform the following additional procedures before starting the IDR process:

- Install and configure the IBM ServeRAID controller card and ServeRAID software so that a boot partition will be visible to the Windows operating system.
- Start the server using the IBM server's ServeRAID Configuration and Management CD in the CD-ROM drive prior to using the IDR bootable media. This will start IBM ServeRAID utilities configuration and installation process to view and update the current BIOS and firmware levels.

Refer to the IBM ServeRAID documentation for complete installation instructions for installing Windows on an IBM Server with the ServeRAID controller. Create and initialize the ServeRAID disks in order for partitions to be visible to the Windows operating system.

See [“Recovering a computer by using the Intelligent Disaster Recovery Wizard”](#) on page 1831.

About the Intelligent Disaster Recovery Wizard

When you use the Intelligent Disaster Recovery Wizard to perform a recovery, the Intelligent Disaster Recovery Wizard lets you access the media device that is required for restore from three sources. You can:

- Use locally attached media devices at the computer that you want to recover.
- Use remote backup-to-disk folders that reside on remote computers.
- Run restore jobs from remote media servers.

To restore data by using the Intelligent Disaster Recovery Wizard, the following items are required:

- The media set that contains the full backup of the target computer being restored.
- If you want to recover a local computer, a storage device must be connected to the computer that you want to recover.
- If you are using a bootable CD, a media server that can restore the backup sets to the target computer must be connected on the network.

See [“Performing an automated restore by using the Intelligent Disaster Recovery Wizard”](#) on page 1832.

See [“Performing a manual restore by using the Intelligent Disaster Recovery Wizard”](#) on page 1841.

About encrypted backup sets and the Intelligent Disaster Recovery Wizard

The Intelligent Disaster Recovery Wizard supports the recovery of computers with previously encrypted backup sets.

When you use the Intelligent Disaster Recovery Wizard’s automated recovery option to recover a local media server, the wizard prompts you for the pass phrase of each encrypted backup set that is required to complete the restore job.

When you use a remote media server to recover a computer with encrypted backup sets, one of the following occurs:

Table R-5 Encryption keys, passphrases, and the Intelligent Disaster Recovery Wizard

Item	Description
If the backup set was created on the remote media server	The Intelligent Disaster Recovery Wizard automatically retrieves the encryption keys.
If the backup set was not created on the remote media server	The Intelligent Disaster Recovery Wizard prompts you to enter the pass phrase.

When you use the Intelligent Disaster Recovery Wizard’s manual recovery option, the wizard prompts you for the pass phrase of each encrypted backup set that is required to complete the recovery.

See [“Encryption keys”](#) on page 416.

Recovering a computer by using the Intelligent Disaster Recovery Wizard

To recover a computer with the Intelligent Disaster Recovery Wizard, the following process must be followed.

Note: To fully automate the recovery, you must have the current *.dr file for the computer that you want to restore. If a *.dr file is unavailable or if the *.dr file is not current, you can still use IDR to manually recover the computer.

Table R-6 Process for recovering a computer by using the Intelligent Disaster Recovery Wizard

Step	Action
Step 1	Start the computer by using the bootable tape or CD that was created with the Intelligent Disaster Recovery Preparation Wizard.
Step 2	Use Windows Setup to prepare the computer for recovery.
Step 3	Use the Intelligent Disaster Recovery Wizard to restore the computer to an operational state and to restore the computer's data from the last backup set.

Performing an automated restore by using the Intelligent Disaster Recovery Wizard

Use the following steps to perform an automated restore by using the Intelligent Disaster Recovery Wizard.

See [“Restoring from a locally attached media device”](#) on page 1835.

See [“Restoring from remote backup-to-disk folders”](#) on page 1837.

See [“Restoring from a remote media server”](#) on page 1839.

To perform an automated restore by using the Intelligent Disaster Recovery Wizard

- 1 Place the bootable IDR CD in the CD drive of the computer to be recovered and then start the computer.
- 2 After reading the **IDR boot** screen, click **Enter**.

- 3 In the initial **Symantec Intelligent Disaster Recovery** panel, click **Automated Recovery**, and then click **Next**.

If the Intelligent Disaster Recovery Wizard fails to run and returns you to the initial IDR Recovery screen, click **Start > View Log File**. Use this log file when you contact Symantec technical support.

If SCSI or RAID controller drivers are required, the drivers are automatically installed if the Intelligent Disaster Recovery Wizard finds them in its driver database. If the SCSI or RAID drivers are not found, click **Have Disk** to install the required drivers, and then click **OK**.

Note: Each listed controller that appears in the **SCSI, RAID and USB Controllers Discovery** panel may have updated drivers available. To determine if updated drivers are available, manually select each listed controller. If updated drivers are available, then the **Update Driver** option becomes active. Symantec recommends that you update the drivers.

- 4 Select the *.dr file for the computer that you want to recover, and then click **Next**.

Each *.dr file is labeled using the name of the computer from which it was created. It also shows the date and time it was created. Make sure you select the correct *.dr file.

- 5 If the *.dr file does not appear, click **Browse** to navigate to the destination folder where you stored the backup copy of the *.dr file.
- 6 If the *.dr file resides on a network drive, click **Install Network** to enable networking.
- 7 After installing the network drivers, click **Browse** to find the *.dr file.
See "[Installing network drivers](#)" on page 1840.
- 8 In the **Hard Disks Layouts** panel, do one of the following:

This step pertains to Windows 2000/Vista/Server 2008 only.

To use the current hard disk layout Click **Keep current layout**, and then click **Next**.

To restore the original hard disk layout Click **Restore original layout**, and then click **Next**.

- 9 Do one of the following:

If BitLocker Drive Encryption is enabled on any of the existing hard drives to which you are recovering

Do the following in the order listed:

- Click **Unlock**.
- In the **BitLocker Drive Recovery** panel, select the file that contains the encryption key, or enter the recovery password.
- Click **Next**.

If BitLocker Drive Encryption is not enabled

Go to step 11.

10 In the **Restore Hard Disk Layout** panel, click **Next**.

11 In the **Modify Hard Disk Layout** panel, do one of the following:

To use the original configuration from the *.dr file

Click **Next**.

To make additional changes to the partition information

Do the following in the order listed:

- Click **Run Disk Management**.
- Modify the disk layout.
- Click **Next**.

For more information on the Windows Disk Management program and fault tolerant configurations, see your Windows documentation.

See [“About altering hard drive partition sizes”](#) on page 1841.

If a *.dr file does not exist for the computer being recovered

Do the following in the order listed:

- Click **Run Disk Management**.
- Modify the partition layout.
- Click **Next**.

12 Select one of the following methods to access the storage device.

Use locally attached media device

Select this option if you have locally-connected backup media such as tape drives, autoloaders, or backup-to-disk folders. If you use bootable tape, you must use this option.

See [“Restoring from a locally attached media device”](#) on page 1835.

Install networking, and then restore from remote backup-to-disk folders

Select this option if your backup-to-disk folders are located on remote computers.

See [“Restoring from remote backup-to-disk folders”](#) on page 1837.

Install networking and then restore from a remote media server

Select this option if you want to submit the restore jobs from remote media servers.

See [“Restoring from a remote media server”](#) on page 1839.

Restoring from a locally attached media device

Use the following steps to restore from a locally-attached media device.

See [“About changing hardware in the computer to be recovered”](#) on page 1829.

To restore from a locally-attached media device

1 From the **Select Restore Method** screen, select **Use locally attached media device** and then click **Next**.

2 After media devices are detected, click **Next**.

Depending on the level of assistance selected, all backup sets may be automatically restored, or you can select individual backup sets to restore.

After the restore is complete, you can provide additional media to restore.

If a *.dr file does not exist, or if there are no catalog entries in the *.dr file, then perform a manual restore and select **I will provide my own media**, and then perform a manual restore by using the Intelligent Disaster Recovery Wizard.

See [“Performing a manual restore by using the Intelligent Disaster Recovery Wizard”](#) on page 1841.

3 Do one of the following:

- If you are restoring from a stand-alone tape drive, insert the correct media and then click **Next**.
- If you are restoring from a backup-to-disk folder, click **Next**.

If you use a robotic library to recover your computer, the first drive in the robotic library is used for the restore.

4 Click **Next**.

On the **Restore Data** screen, all of the backup sets that you need to fully restore the computer are checked by default.

5 Click **Next**.

6 On the **Insert Media into Restore Device** screen, select the backup-to-disk folder or drive that contains the required media that you want to restore.

7 Click **Next**.

8 When the automated restore process finishes, click the appropriate response.

Yes	The Intelligent Disaster Recovery Wizard prompts you to select another media set from which to continue the restore process.
No	The Intelligent Disaster Recovery Wizard updates the hard disk drivers and then finishes the recovery process.

- 9 On the **You have completed the Intelligent Disaster Recovery Wizard** screen, you can run the following:

Run CMD.exe	The Intelligent Disaster Recovery Wizard lets you open a command window that you can use to make further modifications to your computer.
Edit Boot.ini	The Intelligent Disaster Recovery Wizard lets you edit the boot.ini file by using the Windows Notepad application to modify the boot.ini file on the root of the system partition. (Windows 2000/XP/2003 only)
Messages	Click the Messages button to review messages that were generated by the Intelligent Disaster Recovery Wizard during the recovery process.

- 10 Click **Finish**. Remove the CD after the computer restarts but before the start process finds the startup CD drive.

As the computer restarts, a Symantec screen momentarily appears while the Intelligent Disaster Recovery Wizard makes final recovery modifications. After the modifications are finished, the computer restarts again and the recovery process is complete.

Restoring from remote backup-to-disk folders

Use the following steps to restore from remote backup-to-disk folders.

To restore from remote backup-to-disk folders

- 1 From the **Select Restore Method** screen, click **Install networking and then restore from remote backup-to-disk folders**.
- 2 Click **Next**.
The Intelligent Disaster Recovery Wizard attempts to detect and install the correct network drivers.
- 3 After networking is installed, click **Next**.
The Intelligent Disaster Recovery Wizard attempts to detect the remote backup-to-disk folders. After it finds them, a **Connect to <computer_name>** screen appears, prompting you for access credentials to the remote computer where the backup-to-disk folders reside.
- 4 Enter the credentials that are required to access the remote computer.

5 Click **Next**.

The **Detecting Media Devices** screen appears, showing the backup-to-disk folders.

6 Click **Next**.

On the **Restore Data** screen, all of the backup sets that you need to fully restore the computer are checked by default.

7 Click **Next**.

8 On the **Insert Media into Restore Device** screen, select the backup-to-disk folder that contains the required media that you want to restore.

9 Click **Next**.

10 When the automated restore process finishes, click the appropriate response.

Yes	The Intelligent Disaster Recovery Wizard prompts you to select another media set from which to continue the restore process.
-----	--

No	The Intelligent Disaster Recovery Wizard updates the hard disk drivers and then finishes the recovery process.
----	--

11 On the **You have completed the Intelligent Disaster Recovery Wizard** screen, you can run the following:

Run CMD.exe	The Intelligent Disaster Recovery Wizard lets you open a command window that you can use to make further modifications to your computer.
-------------	--

Edit Boot.ini	The Intelligent Disaster Recovery Wizard lets you edit the boot.ini file using the Windows Notepad application to modify the boot.ini file on the root of the system partition.
---------------	---

Messages	Click the Messages button to review messages generated by the Intelligent Disaster Recovery Wizard during the recovery process.
----------	--

12 Click **Finish**. Remove the CD after the computer restarts but before the start process finds the startup CD drive.

As the computer restarts, a Symantec screen momentarily appears while the Intelligent Disaster Recovery Wizard makes final recovery modifications. After the modifications are finished, the computer restarts again and the recovery process is complete.

Restoring from a remote media server

Use the following steps to restore from a remote media server.

To restore from a remote media server

1 From the **Select Restore Method** screen, select **Install networking and then restore from a remote media server**.

2 Click **Next**.

The Intelligent Disaster Recovery Wizard attempts to detect and install the correct network drivers.

3 After networking is installed, click **Next**.

4 On the **Connect to Media Server** screen, enter the credentials that are required to access the media server.

Server Name The name of the remote media server that you selected to run the restore job.

Domain The domain in which the remote media server is a member.

User Name The user name that has administrator rights to the remote media server.

Password The password that is required for access.

5 Click **Next**.

On the **Restore Data** screen, all backup sets that are required to fully restore the computer are checked by default.

6 Click **Next**.

7 When the automated restore process finishes, click the appropriate response.

Yes The Intelligent Disaster Recovery Wizard prompts you to select another media set from which to continue the restore process.

No The Intelligent Disaster Recovery Wizard updates the hard disk drivers and then finishes the recovery process.

- 8 On the **You have completed the Intelligent Disaster Recovery Wizard** screen, you can run the following:

Run CMD.exe	The Intelligent Disaster Recovery Wizard lets you open a command window that you can use to make further modifications to your computer.
Edit Boot.ini	The Intelligent Disaster Recovery Wizard lets you edit the boot.ini file using the Windows Notepad application to modify the boot.ini file on the root of the system partition. (Windows 2000/XP/2003 only)
Messages	Click the Messages button to review messages generated by the Intelligent Disaster Recovery Wizard during the recovery process.

- 9 Click **Finish**. Remove the CD after the computer restarts but before the start process finds the startup CD drive.

As the computer restarts, a Symantec screen momentarily appears while the Intelligent Disaster Recovery Wizard makes final recovery modifications. After the modifications are finished, the computer restarts again and the recovery process is complete.

Installing network drivers

You can install network drivers from any screen in the Intelligent Disaster Recovery Wizard where the Install Network or Configure Network buttons appear.

See [“Performing an automated restore by using the Intelligent Disaster Recovery Wizard”](#) on page 1832.

See [“Performing a manual restore by using the Intelligent Disaster Recovery Wizard”](#) on page 1841.

To install network drivers

- 1 Click **Install Network**.
- 2 On the **Network Configuration** screen, click **Next** after the Intelligent Disaster Recovery Wizard detects and binds each detected network adapter.

By default, each network adapter is assigned IP addresses from the default DHCP server. To assign a static IP address for each detected network adapter, select a network adapter and then click **Configure**.

About altering hard drive partition sizes

When you recover a Windows 2000 computer, IDR restores the hard drive partitions to the same sizes they were before the disaster. There may be unused and unallocated space. If the hard drive in the target computer is larger than the hard drive that was in place before the disaster occurred, run the Windows Disk Management program (within the Intelligent Disaster Recovery Wizard) to alter the partition sizes to reflect the larger hard drive size.

If you did not select the option Let IDR automatically partition the boot and system drive during restore when you recover a Windows 2000 computer, you must specify hard drive partitioning information during setup.

Following is an example of why the hard drive partitions should be resized:

If the pre-disaster computer hardware contained a 4 GB hard drive with two 2 GB partitions, and you replaced it with a 9 GB model, IDR (using the *.dr file) rebuilds the hard disk partition table by using the partition information that is found on the original 4 GB hard drive. As a result, only 4 GB of space is allocated on the new 9 GB hard drive, with a partition map that consists of two 2 GB partitions.

Use the Disk Management program to repartition the hard drive to include the additional space.

See [“About changing hardware in the computer to be recovered”](#) on page 1829.

Performing a manual restore by using the Intelligent Disaster Recovery Wizard

If a *.dr file is missing, you can still recover the computer by initiating a manual restore using the Intelligent Disaster Recovery Wizard. The Intelligent Disaster Recovery Wizard identifies individual backup sets by reading the backup media, so you can select the backup sets that you want to restore.

Caution: If the media that you want to restore contains both Full backup sets and Incremental or Differential backup sets, restore the Full backup sets first.

See [“About changing hardware in the computer to be recovered”](#) on page 1829.

To perform a manual restore

If you are restoring Windows 2000 computers that have utility partitions, first recreate the utility partitions using the OEM-supplied media.

- 1 Place the bootable IDR CD in the CD drive of the computer to be recovered and then start the computer.
- 2 After you read the IDR boot screen, press **Enter**.

- 3 In the initial Symantec Intelligent Disaster Recovery panel, click **Manual Recovery**.
- 4 Change the time zone to the time zone of the computer you want to recover by clicking the time zone hyperlink on the taskbar.
- 5 Select the correct time zone, and then click **Set**.
- 6 Click **Next**.
- 7 If the Intelligent Disaster Recovery Wizard fails to run and returns you to the initial IDR Recovery panel, click **Start > View Log File**. Use this log file when you contact Symantec technical support.

If SCSI or RAID controller drivers are required, the drivers are automatically installed if the Intelligent Disaster Recovery Wizard finds them in its driver database. If the SCSI drivers or RAID drivers are not found, click the Have Disk icon to install the required drivers, and then click **OK**.

Note: Each listed controller that appears in the **SCSI, RAID and USB Controllers Discovery** panel may have updated drivers available. To determine if updated drivers are available, manually select each listed controller. If updated drivers are available, then the **Update Driver** option becomes active. Symantec recommends that you update the drivers.

- 8 Do one of the following:

If BitLocker Drive Encryption is enabled on any of the existing hard drives to which you are recovering Do the following in the order listed:

- Click **Unlock**.
- In the **BitLocker Drive Recovery** panel, select the file that contains the encryption key, or enter the recovery password.
- Click **Next**.

If BitLocker Drive Encryption is not enabled Go to step 5.

- 9 On the **Modify Hard Disk Layout** screen, click **Run Disk Management**.
- 10 Recreate the hard disk partition layout to match the computer's original partition layout.
- 11 After you create the hard disk partition layout, click **Next**.

12 Select one of the following methods to access the storage device.

Use locally attached media device	Select this option if you have locally-connected backup media such as tape drives, autoloaders, or backup-to-disk folders. If you use bootable tape, you must use this option. See “Restoring from a locally attached media device” on page 1835.
Install networking and then restore from remote backup-to-disk folders	Select this option if your backup-to-disk folders are located on remote computers. See “Restoring from remote backup-to-disk folders” on page 1837.
Install networking and then restore from a remote media server	Select this option if you want to submit the restore jobs from remote media servers. See “Restoring from a remote media server” on page 1839.

13 Select the tape drive where the restore media resides.

The **Found a Backup Set** dialog box appears, showing you the first backup set found on the media.

14 To restore to a location other than the one that appears, click **Change**, and then select a location where you want to restore the data. Do not use drive C for the alternate location.**15** Click one of the following:

Click Restore Set	To restore the backup set that appears in Media Information and Set Information. IDR restores the data to the selected partition. When the restore is complete, the Found Backup Set dialog box reappears and shows the next backup set that is found on the media. If there are no more backup sets, the Select Tape Drive screen appears. To restore another backup set, click Restore Set again to restore the next backup set. Repeat this step for each backup set found on the media.
Click Skip Set	To skip the restoration of this backup set and search the media for another backup set from which to restore.
Click Skip Media	To eject the media and replace it with different media.

16 After you restore the last backup set, click **Finish** to end the recovery process and exit the Intelligent Disaster Recovery Wizard.

Microsoft SQL Server recovery notes

The Backup Exec Agent for Microsoft SQL Server option must be installed on the media server in order to perform a complete SQL Server database recovery.

After using Intelligent Disaster Recovery to recover the Windows server, IDR automatically replaces the damaged master and model databases with copies of the master and model databases. After SQL is restarted and the latest master database backup and all other system databases are restored, you must still restore all user databases after completing the IDR Recovery.

Caution: For the Intelligent Disaster Recovery Option to work with SQL 2000, copies are made of the master and model databases. Copies are made only after running non-AOFO (Advanced Open File Option) backups of the master and model databases. If you are using AOFO for SQL backups, you must make at least one backup of the master and model databases without using AOFO. If you upgrade SQL 2000, run another non-AOFO backup of the master and model databases.

See [“How to use snapshot technology with the SQL Agent”](#) on page 1262.

See [“About the Advanced Disk-based Backup Option”](#) on page 898.

Microsoft Exchange recovery notes

The Backup Exec Agent for Microsoft Exchange Server option must be installed on the media server in order to perform a complete Exchange Server database recovery.

After you use Intelligent Disaster Recovery to recover the Windows server, use Backup Exec to restore the Exchange Server databases from the most recent Exchange Server database backups.

SharePoint Portal Server recovery notes

After you use Intelligent Disaster Recovery to recover a Windows server that has SharePoint Portal Server 2001 installed (after restarting the system), the SharePoint Portal Server software is installed but is not functional; you must remove SharePoint Portal Server 2001 and reinstall it before the SharePoint data can be restored.

Citrix Metaframe recovery notes

Backup Exec supports IDR of Citrix Metaframe 1.8, XPa, XPe, and XP computers with the following exceptions:

- IDR of a remote computer is not supported if Citrix is installed on the media server and drive C on the media server is remapped.
- If other drives on a Citrix computer were remapped prior to IDR, the drives must also be remapped during the IDR process before any files are restored.
 - If you selected Automated Recovery during the IDR process and selected a *.dr file, Backup Exec automatically remaps the drives.
 - If you selected Automated Recovery or Manual Recovery during the IDR process, but did not select a *.dr file, you must manually remap the drives.

About using IDR with the Central Admin Server Option

If you have purchased and installed the Central Admin Server Option (CASO), you can perform IDR of the managed media servers in a CASO environment. To prepare recovery media for the managed media servers, you must run the Intelligent Disaster Recovery Preparation Wizard on the central administration server. The *.dr files are stored on the central administration server. During IDR recovery of a managed media server, all restore jobs are submitted from the central administration server. The central administration server will then send the restore jobs to the appropriate managed media server.

Note: You cannot select a managed media server node as a valid media server for IDR preparation if the managed media server node is active in a CAS environment. To create the IDR boot media, you must select the CAS server to which the managed media server is attached.

About using IDR with Veritas Storage Foundation for Windows

If you use VERITAS Storage Foundation for Windows on Windows 2003, IDR can restore the dynamic volumes. During backup, IDR gathers the applications and components necessary to restore the dynamic volumes and adds them to the recovery media. During recovery, the gathered applications are run as part of the Windows Automated System Recovery (ASR) process in order to bring back the dynamic volumes. After the dynamic volumes are recovered, the data recovery on the volumes proceeds as usual.

Best Practices for IDR

The following table presents best practices when using IDR.

Table R-7 Best practices for IDR

Item	Description
Remote IDR	To perform disaster recovery on a remote computer, you must purchase the Remote Agent separately, and it must be running on the remote computer.
Creating Bootable Media	<p>Review the following recommendations before you create bootable media:</p> <ul style="list-style-type: none"> ■ Always verify that the *.dr file was created in the alternate data path that you selected. ■ When creating a bootable tape, run the Intelligent Disaster Recovery Preparation Wizard and create the bootable image before you run a full backup. ■ When creating a bootable CD, run a full backup before you create the bootable media. ■ If the backup media resides on another Backup Exec media server, select the option Choose a media server that has the IDR Option installed on the first Intelligent Disaster Recovery Preparation Wizard screen. ■ For local IDR, backup-to-disk folders must be on a drive or drives that can be accessed for recovery.
Disaster Recovery	<ul style="list-style-type: none"> ■ The new partition layout must be the same size or larger than the original. ■ Have the latest RAID, SCSI, or NIC (if remote) drivers available on disks.

Symantec Backup Exec NDMP Option

This appendix includes the following topics:

- [About the NDMP Option](#)
- [Requirements for using the NDMP Option](#)
- [About installing the NDMP Option](#)
- [Adding an NDMP server to Backup Exec](#)
- [Sharing the devices on an NDMP server between multiple media servers](#)
- [Backing up NDMP resources](#)
- [About including and excluding directories and files for NDMP backup selections](#)
- [How to duplicate backed up NDMP data](#)
- [Restoring NDMP data](#)
- [About redirecting restored NDMP data](#)
- [Setting the default backup and restore options for NDMP](#)
- [Viewing NDMP server properties](#)

About the NDMP Option

The Symantec Backup Exec NDMP Option uses the Network Data Management Protocol (NDMP) to back up and restore Network Attached Storage (NAS) devices.

You can back up data from a NAS device to the following locations:

- A storage device that is directly connected to the NDMP-enabled NAS device (direct-attached)
- A storage device that is connected to another NDMP-enabled NAS device (filer-to-filer)
- A backup-to-disk device on a Backup Exec media server (remote)
- A tape device that is attached to a Backup Exec media server (remote)
- A shared tape device or backup-to-disk folder that is connected to a SAN with Backup Exec media servers.

Note: You cannot back up NDMP data to a simulated tape library or to a tape device that is attached to a Backup Exec Remote Media Agent for Linux Servers.

You can restore data from a storage device on a Backup Exec media server to a NAS device. However, you cannot redirect NDMP data to a computer that runs the Windows or Linux operating systems.

You can share tape devices between single or multiple Backup Exec media servers and NAS devices by using the Backup Exec SAN Shared Storage Option. In addition, you can mix NDMP data with non-NDMP data in the same backup job.

See [“About the SAN Shared Storage Option”](#) on page 1991.

See [“Requirements for using the NDMP Option”](#) on page 1848.

See [“About installing the NDMP Option”](#) on page 1849.

Requirements for using the NDMP Option

To use the NDMP Option, the Backup Exec media server must have the following items installed:

- Windows XP/Server 2003/Server 2008/Server 2008 R2.
- Backup Exec.

See [“Installing a custom installation of Backup Exec”](#) on page 119.

In addition, you must have an NDMP server with version 4 of the Network Data Management Protocol enabled.

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

About installing the NDMP Option

The NDMP Option is installed locally on the media server as a separate add-on component of Backup Exec. No files are copied to the NDMP server.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

Adding an NDMP server to Backup Exec

To configure Backup Exec to use the NDMP option, you must add the NDMP server to Backup Exec. If any storage devices are attached to the NDMP server, Backup Exec discovers them automatically after the services are restarted, and then adds them to the list of devices.

In a CASO environment, you can add an NDMP server only to the following servers:

- A central administration server
- A managed media server on which the device and media database is located

When you add an NDMP server, you can select the media servers that can access the devices that are attached to the NDMP server.

See [“About sharing storage”](#) on page 444.

To add an NDMP server to Backup Exec

- 1 On the navigation bar, click **Devices**.
- 2 In the task pane, under **Device Tasks**, click **Configure devices assistant**.
- 3 In the **Configure Devices Assistant** dialog box, under **NDMP Option**, click **NDMP Storage**.
- 4 If the **NDMP Server Configuration** dialog box appears, click **Add an NDMP Server**.

This step does not apply if this is the first NDMP server. The **NDMP Server Configuration** dialog box appears only if an NDMP server already exists.

- 5 On the **General** tab, enter the appropriate information.
See [“Add NDMP Server options”](#) on page 1850.
- 6 On the **Sharing** tab, select each media server that you want to use the devices that are attached to this NDMP server.

The **Sharing** tab appears only if you have the Backup Exec Central Admin Server Option or the SAN Shared Storage Option installed. You must also have more than one media server.

- 7 Click **OK** to add the NDMP server.
- 8 Restart Backup Exec services.
See “[Starting and stopping Backup Exec services](#)” on page 168.

Add NDMP Server options

When you add an NDMP server to Backup Exec, the following options are required. See “[Adding an NDMP server to Backup Exec](#)” on page 1849.

Table S-1 Add NDMP Server options

Item	Description
Server	Indicates the name of the NDMP server.
Port	Lists the port to be used for communications between the Backup Exec media server and the NDMP server.
Description	Shows the user-defined description of the server.
Enable ICMP ping operations for Backup Exec to detect the NDMP Server	Ensures that Backup Exec can use ping to locate the NDMP server.
Logon account	Indicates the name of the logon account for the NDMP server.

Sharing the devices on an NDMP server between multiple media servers

If you use the Backup Exec Central Admin Server Option or the SAN Shared Storage Option, you can select which media servers can share the devices that are attached to an NDMP server. When you add an NDMP server, the media server that you used to add the server is automatically selected for sharing.

Note: If you upgraded from an earlier version of Backup Exec, your existing configuration is preserved, so you do not have to set up sharing for existing configurations.

See “[About sharing storage](#)” on page 444.

To share the devices on an NDMP server between multiple media servers

- 1 On the navigation bar, click **Devices**.
- 2 In the **Devices** view, right-click the NDMP server that has the devices you want media servers to access.
- 3 Select **Manage sharing**.

The **Manage sharing** option appears only if you have the Backup Exec Central Admin Server Option or the SAN Shared Storage Option installed. You must also have more than one media server.
- 4 Select the NDMP server that has the devices you want to share.
- 5 Under **Media Servers**, select the media servers that you want to use with the devices that are attached to the selected NDMP server.
- 6 Click **OK**.
- 7 Restart the services on the media servers that you selected in step 5.

Backing up NDMP resources

Before you back up NDMP resources, review the following information:

- The NDMP Option does not exclude folders from the backup job if the parent folder is backed up. Instead, all items in the parent folder are backed up, even if you marked items for exclusion from the backup.
- Backup Exec cannot gather sufficient file and directory information on an NDMP backup to accurately populate the **Job Summary** and **Set Detail Information** sections of the job history. Therefore, the number of files, directories, files skipped, corrupt files, and files in use always appears as 0.
- If your environment includes NAS devices from more than one provider, you must create separate backup jobs for each provider. If you include NAS devices from multiple providers in the same backup job, the job will fail.

To back up NDMP resources

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 In the **Properties** pane, under **Source**, click **Selections**.

- 4 In the backup selections tree, expand either **NDMP Servers** or **User-defined Selections**.

If Backup Exec added the NDMP server to the backup selection list automatically, the NDMP server appears under **NDMP Servers**. If you added the NDMP server to the backup selection list manually, the NDMP server appears under **User-defined Selections**.

- 5 Select the NDMP resource you want to back up.

You may be prompted to select or create a logon account for this resource.

If you do not want to back up the entire NDMP resource, select specific files or directories to include in or exclude from the backup job.

See [“About including and excluding directories and files for NDMP backup selections”](#) on page 1854.

- 6 In the **Properties** pane, under **Destination**, click **Device and Media**.

- 7 In the Device list, select a storage device.

- 8 In the **Properties** pane, under **Settings**, click **NDMP**.

- 9 Select the tab for the type of NDMP server that you use, and then select backup method and other backup options that you want to use for this job.

See [“NDMP backup options for NetApp/IBM/Fujitsu”](#) on page 1852.

See [“NDMP backup options for EMC”](#) on page 1853.

See [“NDMP backup options for providers other than NetApp/IBM/Fujitsu and EMC”](#) on page 1854.

- 10 Do one of the following:

- Start the backup job.
- Select other backup options from the **Properties** pane, and then start the backup job.

NDMP backup options for NetApp/IBM/Fujitsu

When you create a backup job for NDMP, you can set any of the following options that are appropriate for the job.

See [“Backing up NDMP resources”](#) on page 1851.

Table S-2 NDMP backup options for NetApp/IBM/Fujitsu

Item	Description
Backup method	Specifies the backup level. Level 0 provides a full backup. Levels 1 through 9 provide various levels of incremental backups. Level 1 backup method backs up new or modified files since the level 0 backup. Level 2 backup method backs up new or modified files since the level 1 backup, and so on.
Back up Access Control Lists	(NetApp filers only). Backs up NetApp Access Control Lists.
Enable file history	Enables the generation of file history data. File history is used to optimize recovery of selected subsets of data from the backup image. File history generation and processing increase the backup time. Disabling this option improves backup time. If file history is made unavailable and you must restore data later, restore the entire backup image.

NDMP backup options for EMC

When you create a backup job for NDMP, you can set any of the following options that are appropriate for the job.

See “[Backing up NDMP resources](#)” on page 1851.

Table S-3 NDMP backup options for EMC

Item	Description
Backup method	Specifies the backup level. Level 0 provides a full backup. Levels 1 through 9 provide various levels of incremental backups. Level 1 backup method backs up new or modified files since the level 0 backup. Level 2 backup method backs up new or modified files since the level 1 backup, and so on.
Backup type	Determines the backup type for this backup job. The following backup types are available: <ul style="list-style-type: none"> ■ VBB ■ Dump

Table S-3 NDMP backup options for EMC (*continued*)

Item	Description
Back up with integrated checkpoints (SnapSure)	Enables Backup Exec to create a backup set that uses the EMC SnapSure feature. For more information about SnapSure, see your EMC documentation.
Enable file history	Allows for the recovery of selected subsets of data from the backup history. If you uncheck Enable file history , file history data is not generated, but backup time may be reduced. This option is selected by default.

NDMP backup options for providers other than NetApp/IBM/Fujitsu and EMC

On the **Other** tab, you select the NDMP provider that you use. When you select a provider, the appropriate environment variables for that provider appear. The values for most variables, such as the backup level and file history option, can be changed. Some variables can be changed to a predefined option, some variables can be changed to any value, and some variables cannot be changed at all.

Note: Symantec has tested the variables that appear for each NDMP provider. The NDMP provider may support additional variables. However, Symantec does not support them at this time.

Backup Exec does not validate the values that you enter for the variables, so you should ensure that you enter the values correctly. See the documentation from your NDMP provider for information about the values that you use for the variables that appear. If you use the **Include/Exclude Selections** dialog box to include or exclude data from a backup job, be sure that the vendor that you select on the **Other** tab matches the vendor that you selected on the **Include/Exclude Selections** dialog box.

See “[Backing up NDMP resources](#)” on page 1851.

About including and excluding directories and files for NDMP backup selections

When you create a backup job, you can do the following:

About including and excluding directories and files for NDMP backup selections

- Select specific directories to include in the backup job.
- Select specific directories and files to exclude from the backup job.

Table S-4 What you can include and exclude for NDMP backup selections

Type of NDMP backup selection	Include	Exclude
NetApp/IBM/Fujitsu	Single or multiple directories	Directories and files
EMC	Single directory	Directories and files (only if you select the "dump" backup type)
Other	Single or multiple directories, depending on the NDMP provider.	Directories and files

When you create a backup job, you can select the NDMP resources that you want to include from either the main resource selection area or from the **Include/Exclude Selections** dialog box.

The following limitations apply when you select NDMP resources from the main resource selection area instead of from the **Include/Exclude Selections** dialog box:

- You can include entire volumes for any NDMP provider.
- You can include subfolders for NetApp/IBM/Fujitsu providers only.
- You cannot exclude files or directories. You must use the **Include/Exclude Selections** dialog box to exclude NDMP resources from a backup job.
- For providers other than NetApp/IBM/Fujitsu and EMC, Backup Exec uses the "Other" provider that was selected last on either the NDMP settings on the **Backup Job Properties** dialog box or the **Include/Exclude Selections** dialog box. If no previous selections exist on those dialog boxes, then Backup Exec uses the default NDMP settings from **Tools > Options**.

See [“Including specific directories in an NDMP backup selection”](#) on page 1856.

See [“How to use patterns to exclude files and directories from an NDMP backup selection”](#) on page 1857.

See [“Excluding directories and files from a NetApp/IBM/Fujitsu backup selection”](#) on page 1859.

Including specific directories in an NDMP backup selection

When you create a backup job to back up NDMP resources, you can select specific directories to include in the backup job. However, you cannot include specific files in an NDMP backup selection.

To include specific directories in an NDMP backup selection

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 Select the resource that contains the directories you want to include in the backup selection
- 4 Click **Include/Exclude**.
- 5 In **Resource type**, verify that **NDMP** is selected.
- 6 In **NDMP type**, verify that the type of NDMP provider that you use is selected.

If you do not use NetApp/IBM/Fujitsu or EMC, select **Other**. After you select **Other**, an option appears that lets you select a specific provider.

- 7 In **Type**, click **Include**.
- 8 Do one of the following:

If you selected NetApp/IBM/Fujitsu in the **NDMP type** field

Do one of the following:

- In the **Resources** pane, navigate to the directory you want to include.

Note: If you want to select multiple directories, Symantec recommends that you select them from the Resources pane instead of typing the names of the directories.

- In **Path**, type the directory you want to include.

If you selected EMC in the **NDMP type** field

In **Path**, type the directory you want to include.

- If you selected Other in the **NDMP type** field
- Do the following:
- In **Select an NDMP provider**, select the provider you use.
 - Under **NDMP selection settings**, enter the volume and path to the directory you want to include.
 - In **Value**, enter the appropriate values for any variables that can be changed.
- Note:** See the documentation from your NDMP provider for information about what value to enter for the variable that is listed.

9 Do one of the following

- | | |
|---|--|
| To include additional directories in the backup selection | Click Apply , and then enter the additional directories that you want to include. |
| To complete this procedure | Click OK . |

See “[About including and excluding directories and files for NDMP backup selections](#)” on page 1854.

How to use patterns to exclude files and directories from an NDMP backup selection

When you exclude files and directories from a backup selection for an EMC Celerra Server or a NetApp/IBM/Fujitsu appliance, you must use patterns. You should enter patterns carefully to ensure that you exclude the correct files and directories. Backup Exec does not verify the validity of exclude patterns. If you enter an invalid pattern, the pattern is ignored and therefore the files or directories are not excluded.

For details about how to use patterns, see your NDMP vendor's documentation.

The following example shows a pattern to exclude files and directories from a backup selection for a NetApp/IBM/Fujitsu appliance:

Table S-5 Example pattern for NetApp/IBM/Fujitsu appliances

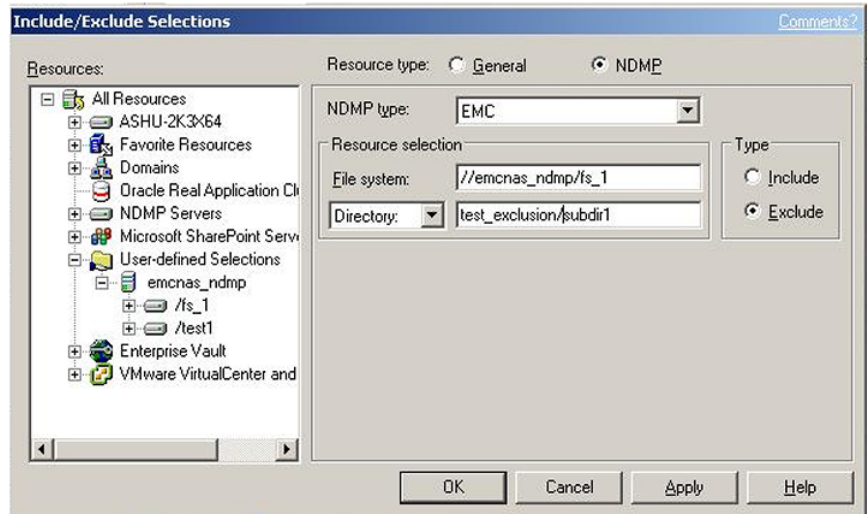
Pattern	Example
tmp	Excludes all files and directories that have the name “tmp”.

Table S-5 Example pattern for NetApp/IBM/Fujitsu appliances (*continued*)

Pattern	Example
*.core	Excludes all files and directories that end with ".core".

To exclude directories for an EMC Celerra Server, do not include the name of the EMC Celerra Server or the name of the file system in the pattern. The names of the NDMP server and the file system are already included in the **File system** text box. If you repeat the name of the NDMP server and the file system in the Directory pattern, the EMC Celerra Server ignores the exclusion. Type the path from the root directory to the directory that you want to exclude. Do not include an initial forward slash (/).

Figure S-1 Excluding EMC directories



The following example shows a pattern to exclude directories from a backup selection for an EMC Celerra Server:

Table S-6 Example pattern to exclude directories for an EMC Celerra Server

Pattern	Description
test_exclusion/subdir1	Excludes only the "subdir1" directory on the file system that is listed in the File system text box.

Table S-6 Example pattern to exclude directories for an EMC Celerra Server
(continued)

Pattern	Description
d*	Excludes all directories that start with the letter “d” on the file system that is listed in the File system text box.

The following example shows a pattern to exclude files from a backup selection for an EMC Celerra Server:

Table S-7 Example pattern to exclude files for an EMC Celerra Server

Pattern	Description
*.mp3	Excludes all files that end with “.mp3”.
temp	Excludes all files that have the name “temp”.

See [“Excluding directories and files from a NetApp/IBM/Fujitsu backup selection”](#) on page 1859.

See [“Excluding directories and files from an EMC backup selection”](#) on page 1860.

Excluding directories and files from a NetApp/IBM/Fujitsu backup selection

When you create a backup job, you can select specific files and directories that you do not want to include in the backup job.

To exclude directories and files from a NetApp/IBM/Fujitsu backup selection

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 Select the resource that contains the files you want to exclude from the backup selection.
- 4 Click **Include/Exclude**.
- 5 In **Resource type**, verify that **NDMP** is selected.
- 6 In **NDMP type**, verify that **NetApp/IBM/Fujitsu** is selected.
- 7 In **Type**, click **Exclude**.
- 8 Use one of the following methods to select the volume where the file or directory you want to exclude is located:

- Under **Resources**, navigate to the volume.
 - Under **Resource selection**, in **Volume**, type the path for the volume.
- 9 In **Pattern**, type the pattern to exclude the file or directory.
See “[How to use patterns to exclude files and directories from an NDMP backup selection](#)” on page 1857.
- 10 Do one of the following:
- | | |
|--|---|
| To exclude additional directories or files from the backup selection | Click Apply , and then repeat steps 6 through 9. |
| To complete this procedure | Click OK . |

Excluding directories and files from an EMC backup selection

When you select the “dump” backup type, you can select specific files and directories that you do not want to include in the backup job.

Note: If you select VBB as the backup method, exclusions are ignored.

To exclude directories and files from an EMC backup selection

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 Select the resource that contains the files you want to exclude from the backup selection.
- 4 Click **Include/Exclude**.
- 5 In **Resource type**, verify that **NDMP** is selected.
- 6 In **NDMP type**, verify that **EMC** is selected.
- 7 In **Type**, click **Exclude**.
- 8 Use one of the following methods to select the file system where the file or directory you want to exclude is located:
 - Under **Resources**, navigate to the file system.
Backup Exec automatically fills in the **File system** text box with the name of the EMC Celerra Server and the file system that you selected.
 - Under **Resource selection**, in **File system**, type the path for the file system if it is not already filled in.
Use the following format:

About including and excluding directories and files for NDMP backup selections

//EMC_Celerra_Server_name/file_system_name

For example, type "//emcnas_ndmp/fs_1" to indicate an EMC Celerra Server named "emcnas_ndmp" and a file system named "fs_1".

9 Do one of the following:

To exclude a file

Under **Resource selection**, in the drop-down list, click **File pattern**, and then type the pattern in the text box.

To exclude a directory

Under **Resource selection**, in the drop-down list, click **Directory**, and then type the pattern in the text box.

Do not include the name of the EMC Celerra Server or the name of the file system in the pattern. The names of the NDMP server and the file system are already included in the **File system** text box. Type the path from the root directory to the directory that you want to exclude. Do not include an initial forward slash (/).

The following example shows how to type the pattern to exclude the "/test_exclusion/subdir1" directory:
test_exclusion/subdir1

See [“How to use patterns to exclude files and directories from an NDMP backup selection”](#) on page 1857.

10 Do one of the following:

To exclude additional directories or files from the backup selection

Click **Apply**, and then repeat steps 6 through 9.

To complete this procedure

Click **OK**.

Excluding files and directories from backup selections for NDMP providers other than NetApp/IBM/Fujitsu and EMC

You can exclude specific files and directories from an NDMP backup job based on the NDMP provider's rules for data exclusion.

To exclude files and directories from backup selections for NDMP providers other than NetApp/IBM/Fujitsu and EMC

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 Select the resource that contains the files you want to exclude from the backup selection.
- 4 Click **Include/Exclude**.
- 5 In **Resource type**, verify that **NDMP** is selected.
- 6 In **NDMP type**, verify that **Other** is selected.
- 7 In **Type**, click **Exclude**.
- 8 In **Volume**, enter the volume that contains the files or directories you want to exclude.
- 9 In **Select an NDMP provider**, select the name of your provider.
- 10 In **Value**, enter the appropriate value for the variable that is listed.
See the documentation from your NDMP provider for information about the appropriate value for the variable.
- 11 Click **OK**.

See [“About including and excluding directories and files for NDMP backup selections”](#) on page 1854.

How to duplicate backed up NDMP data

You can create a job to duplicate backup data. When you create a duplicate job, you can select a device that is attached to a Backup Exec media server or to a NAS sever. You can use tape devices, backup-to-disk devices, or virtual tape libraries.

Backup Exec supports the following configurations:

- Two tape devices that are attached locally to the Backup Exec media server.
- Two tape devices that are attached locally to a NAS server.
- One tape device that is attached locally to a NAS server and one tape device that is attached locally to another NAS server.
- One tape device that is attached locally to a Backup Exec media server and one tape device that is attached locally to a NAS server.

The procedure to duplicate NDMP data is the same as the procedure to duplicate any other type of data. However, you must select the logon credential for the source NDMP server.

See [“Duplicating backed up data”](#) on page 369.

Note: If the data that you want to duplicate is hardware-encrypted, you should choose a destination device that allows hardware encryption. Otherwise, the duplicate job fails.

Restoring NDMP data

During the restore process, you can select individual files for restore if file history was enabled for the backup job.

Backup Exec cannot gather sufficient file and directory information on an NDMP restore job to accurately populate the **Job Summary** and **Set Detail Information** sections of the job history. Therefore, the number of files, directories, files skipped, corrupt files, and files in use always appears as 0.

NDMP backup sets cannot be cataloged unless the following option is selected as a catalog default:

Use storage media-based catalogs

See [“Changing application settings for the catalogs”](#) on page 601.

Note: You cannot exclude files and directories from restore jobs on NDMP servers. Excluded directories and files are restored.

To restore NDMP data

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 Select the data you want to restore.
- 4 To change or test logon credentials, on the **Properties** pane, under **Source**, click **Resource Credentials**.

See [“Resource Credentials options”](#) on page 337.

- 5 In the **Properties** pane, under **Source**, click **Device and Media**.
- 6 Select any of the appropriate options.

See [“Device options for restore jobs”](#) on page 611.

- 7 In the **Properties** pane, under **Settings**, click **NDMP**.
- 8 Select the tab for your NDMP provider, and then select the appropriate options for this job.
See “[NDMP restore options for NetApp/IBM/Fujitsu](#)” on page 1864.
See “[NDMP restore options for EMC](#)” on page 1865.
See “[NDMP restore options for providers other than NetApp/IBM/Fujitsu and EMC](#)” on page 1866.
- 9 Do one of the following:
 - Start the restore job.
 - Select other restore options from the **Properties** pane, and then start the restore job.

NDMP restore options for NetApp/IBM/Fujitsu

When you create a restore job for NDMP, you can select any of the following options.

See “[Restoring NDMP data](#)” on page 1863.

Table S-8 NDMP restore options for NetApp/IBM/Fujitsu

Item	Description
Restore Access Control Lists	Restores NetApp Access Control Lists.
Enable Direct Access Recovery	Enables Backup Exec to use Direct Access Recovery (DAR) during the restore job. With DAR-enabled recovery, Backup Exec can specify the exact location of a file in a backup data stream. The NDMP server can then read the data applicable to the single file being restored. This practice reduces the amount of information that is processed and significantly reduces recovery time. If DAR is not available, the restore may take significantly longer.
Restore without writing data to disk (Verify data without doing a restore)	Tests the validity of the data that you selected for the restore job. Backup Exec does not restore the data. For NetApp/IBM filers, you should use this option to verify data instead of Backup Exec's verify backup job option.

Table S-8 NDMP restore options for NetApp/IBM/Fujitsu (*continued*)

Item	Description
Preserve tree	<p>Restores the data with its original directory structure intact. This option is enabled by default. If you clear this option, all data in the directories and the subdirectories is restored to the path you specify in the File Redirection dialog box.</p> <p>See “File Redirection restore options” on page 634.</p> <p>Note: This option affects restore jobs for NetApp/IBM/Fujitsu data only. For EMC data, use the Preserve tree option on the EMC tab. For non-NDMP data, use the Preserve tree option on the General Restore Job Properties dialog box.</p> <p>See “General options for restore jobs” on page 612.</p>

NDMP restore options for EMC

When you create a restore job for NDMP, you can select any of the following options.

See [“Restoring NDMP data”](#) on page 1863.

Table S-9 NDMP restore options for EMC

Item	Description
Enable Direct Access Recovery	<p>Enables Backup Exec to use Direct Access Recovery (DAR) during the restore job. With DAR-enabled recovery, Backup Exec can specify the exact location of a file in a backup data stream. The NDMP server can then read the data applicable to the single file being restored. This practice reduces the amount of information that is processed and significantly reduces recovery time. If DAR is not available, the restore may take significantly longer.</p>

Table S-9 NDMP restore options for EMC (continued)

Item	Description
<p>Preserve tree</p>	<p>Restores the data with its original directory structure intact. This option is enabled by default. If you clear this option, all data in the directories and the subdirectories is restored to the path you specify in the File Redirection dialog box.</p> <p>See "File Redirection restore options" on page 634.</p> <p>Note: This option affects the restore of EMC data only. For NetApp/IBM/Fujitsu data, use the Preserve tree option in the NetApp/IBM/Fujitsu tab. For non-NDMP data, use the Preserve tree option on the General Restore Job Properties dialog box.</p> <p>See "General options for restore jobs" on page 612.</p>
<p>Restore over existing files</p>	<p>Overwrites files on the target resource that have the same name as files that are being restored. Use this option only when you are sure that you want to restore an older version of a file.</p>

NDMP restore options for providers other than NetApp/IBM/Fujitsu and EMC

On the **Other** tab, you select the NDMP provider that you use. When you select a provider, the appropriate variables for that provider appear. Backup Exec sets the values for the variables by default. However, you can change the values to meet your needs.

Variables that begin with the prefix "@@" are specific to Backup Exec rather than to a specific NDMP vendor. For example, the @@PRESERVETREE variable is specific to Backup Exec and is the same as the **Preserve tree** option on the NDMP restore options for NetApp/IBM/Fujitsu. The @@PRESERVETREE option restores the data with its original directory structure intact. This option is enabled by default. If you uncheck this option, all data in the directories and the subdirectories is restored to the path you specify in the **File Redirection** dialog box.

Note: Symantec has tested the variables that appear for each NDMP provider. The NDMP provider may support additional variables. However, Symantec does not support them at this time.

Backup Exec does not validate the values for the variables that you enter, so you should ensure that you enter the values correctly. See the documentation from your NDMP provider for information about the values that you use for the variables.

See [“Restoring NDMP data”](#) on page 1863.

About redirecting restored NDMP data

You can redirect NDMP data from one NDMP server to another NDMP server.

When you redirect NDMP data, be aware of the following limitations:

- You cannot redirect NDMP data to a computer that runs the Windows or Linux operating systems.
- You cannot redirect non-NDMP data, such as NTFS or SQL data, to an NDMP server.
- The server to which you want to redirect the restored data must be from the same vendor as the server from which the data was backed up.

See [“About redirecting restore jobs”](#) on page 634.

Setting the default backup and restore options for NDMP

You can use the defaults that Backup Exec sets during installation for all NDMP backup and restore jobs, or you can choose your own defaults. You can also change the defaults for any specific backup or restore job.

To set default backup and restore options for NDMP

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **NDMP**.

- 3 Select the appropriate options.
 See [“NDMP default options for NetApp/IBM/Fujitsu backup and restore”](#) on page 1868.
 See [“NDMP default options for EMC backup and restore”](#) on page 1869.
 See [“NDMP default backup and restore options for providers other than NetApp/IBM/Fujitsu and EMC”](#) on page 1871.
- 4 Click **OK**.

NDMP default options for NetApp/IBM/Fujitsu backup and restore

You can set up default options for all backup and restore jobs.

See [“Setting the default backup and restore options for NDMP”](#) on page 1867.

Table S-10 NDMP default options for NetApp/IBM/Fujitsu backup and restore

Item	Description
Backup method	Specifies the backup level. Level 0 provides a full backup. Levels 1 through 9 provide various levels of incremental backups. Level 1 backup method backs up new or modified files since the level 0 backup. Level 2 backup method backs up new or modified files since the level 1 backup, and so on.
Back up Access Control Lists	Backs up NetApp Access Control Lists.
Enable file history	Allows for the recovery of selected subsets of data from the backup history. If you uncheck Enable file history, file history data is not generated, but backup time may be reduced. This option is selected by default.
Restore Access Control Lists	Restores NetApp Access Control Lists.
Enable Direct Access Recovery	Enables Backup Exec to use Direct Access Recovery (DAR) during the restore job. With DAR-enabled recovery, Backup Exec can specify the exact location of a file in a backup data stream. The NDMP server can then read the data applicable to the single file being restored. This practice reduces the amount of information that is processed and significantly reduces recovery time. If DAR is not available, the restore may take significantly longer.

Table S-10 NDMP default options for NetApp/IBM/Fujitsu backup and restore (continued)

Item	Description
Restore without writing data to disk (Verify data without doing a restore)	Tests the validity of the data that you selected for the restore job. Backup Exec does not restore the data. For NetApp/IBM filers, you should use this option to verify data instead of Backup Exec's verify backup job option.
Preserve tree	Restores the data with its original directory structure intact. This option is enabled by default. If you clear this option, all data in the directories and the subdirectories is restored to the path you specify in the File Redirection dialog box. See “File Redirection restore options” on page 634. Note: This option affects restore jobs for NetApp/IBM data only. For EMC data, use the Preserve tree option in the EMC group box. For non-NDMP data, use the Preserve tree option on the General Restore Job Properties dialog box. See “General options for restore jobs” on page 612.

NDMP default options for EMC backup and restore

You can set up default options for all NDMP backup and restore jobs that use EMC providers.

See [“Setting the default backup and restore options for NDMP”](#) on page 1867.

Table S-11 NDMP default options for EMC backup and restore

Item	Description
Backup method	Specifies the backup level. Level 0 provides a full backup. Levels 1 through 9 provide various levels of incremental backups. Level 1 backup method backs up new or modified files since the level 0 backup. Level 2 backup method backs up new or modified files since the level 1 backup, and so on.

Table S-11 NDMP default options for EMC backup and restore (*continued*)

Item	Description
Backup type	<p>Determines the backup type for backup jobs.</p> <p>The following backup types are available:</p> <ul style="list-style-type: none"> ■ VBB ■ Dump
Back up with integrated checkpoints (SnapSure)	<p>Enables Backup Exec to create a backup set that uses the EMC SnapSure feature.</p> <p>For more information about SnapSure, see your EMC documentation.</p>
Enable file history	<p>Allows for the recovery of selected subsets of data from the backup history. If you uncheck Enable file history, file history data is not generated, but backup time may be reduced. This option is selected by default.</p>
Enable Direct Access Recovery	<p>Enables Backup Exec to use Direct Access Recovery (DAR) during the restore job. With DAR-enabled recovery, Backup Exec can specify the exact location of a file in a backup data stream. The NDMP server can then read the data applicable to the single file being restored. This practice reduces the amount of information that is processed and significantly reduces recovery time. If DAR is not available, the restore may take significantly longer.</p>

Table S-11 NDMP default options for EMC backup and restore (*continued*)

Item	Description
Preserve tree	<p>Restores the data with its original directory structure intact. This option is enabled by default. If you clear this option, all data in the directories and the subdirectories is restored to the path you specify in the File Redirection dialog box.</p> <p>See “File Redirection restore options” on page 634.</p> <p>Note: This option affects the restore of EMC data only. For non-NDMP data, use the Preserve tree option on the General Restore Job Properties dialog box.</p> <p>See “General options for restore jobs” on page 612.</p>
Restore over existing files	<p>Overwrites files on the target resource that have the same name as files that are being restored. Use this option only when you are sure that you want to restore an older version of a file.</p>

NDMP default backup and restore options for providers other than NetApp/IBM/Fujitsu and EMC

You can set up default options for all NDMP backup and restore jobs that use providers other than NetApp/IBM/Fujitsu and EMC. On the **Other** tab, you select the NDMP provider that you use for backups and restores.

Note: You must select a provider for backup settings and a provider for restore settings.

When you select a provider, the appropriate variables for that provider appear. Backup Exec sets the values for the variables by default. However, you can change the values to meet your needs.

Note: Symantec has tested the variables that appear for each NDMP provider. The NDMP provider may support additional variables, but Symantec does not support them at this time.

Backup Exec does not validate the values that you enter, so you should ensure that you enter the values correctly. See the documentation from your NDMP provider for information about the values that you can use for the variables.

Some variables are specific to Backup Exec rather than to a specific NDMP vendor. For example, the @@PRESERVETREE variable is specific to Backup Exec and is the same as the **Preserve tree** option on the NDMP restore options for NetApp/IBM/Fujitsu. The @@PRESERVETREE option restores the data with its original directory structure intact. This option is enabled by default. If you uncheck this option, all data in the directories and the subdirectories is restored to the path you specify in the **File Redirection** dialog box.

See [“Setting the default backup and restore options for NDMP”](#) on page 1867.

Viewing NDMP server properties

You can view details for all NDMP servers that appear on the **View by Resources** tab.

See [“About the Agent for VMware”](#) on page 1386.

To view NDMP server properties

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 On the **View by Resource** tab, expand **NDMP Servers**.
- 4 Right-click the name of an NDMP server.
- 5 Click **Properties**.

See [“NDMP server properties”](#) on page 1872.

NDMP server properties

You can view the following properties for NDMP servers.

See [“Viewing NDMP server properties”](#) on page 1872.

Table S-12 NDMP server properties

Item	Description
Comment	Displays the user-defined text for this NDMP server.
Domain	Indicates the name of the network domain for the Backup Exec media server.
Type	Indicates that this is an NDMP server.
Installed	Indicates if Backup Exec or the Remote Agent is installed.
Status	Indicates if the installed software is running.
Version	Indicates the version of Backup Exec or the Remote Agent that is installed.

Symantec Backup Exec Remote Agent for Linux or UNIX Servers

This appendix includes the following topics:

- [About the Remote Agent for Linux or UNIX Servers](#)
- [Requirements for the Remote Agent for Linux or UNIX Servers](#)
- [About installing the Remote Agent for Linux or UNIX Servers](#)
- [About configuring the Remote Agent for Linux or UNIX Servers](#)
- [About publishing Linux, UNIX, and Macintosh computers to media servers](#)
- [About excluding files and directories from backup jobs for Linux, UNIX, and Macintosh computers](#)
- [Editing configuration options for Linux, UNIX, and Macintosh computers](#)
- [About backing up data by using the Remote Agent for Linux or UNIX Servers](#)
- [Restoring data to Linux, UNIX, and Macintosh computers](#)
- [Edit the default backup and restore job options for Linux, UNIX, and Macintosh computers](#)
- [Uninstalling the Remote Agent for Linux or UNIX Servers](#)
- [Starting the Remote Agent for Linux or UNIX Servers daemon](#)
- [Stopping the Remote Agent for Linux or UNIX Servers daemon](#)
- [Troubleshooting the Remote Agent for Linux or UNIX Servers](#)

About the Remote Agent for Linux or UNIX Servers

The Backup Exec Remote Agent for Linux or UNIX Servers (Remote Agent) is installed as a separate add-on component. The Remote Agent enables network administrators to perform backup and restore operations on Linux or UNIX servers that are connected to the network. The Remote Agent must be installed on the Linux or UNIX servers before you can perform backup or restore operations.

See “[About installing the Remote Agent for Linux or UNIX Servers](#)” on page 1877.

See “[Requirements for the Remote Agent for Linux or UNIX Servers](#)” on page 1876.

Requirements for the Remote Agent for Linux or UNIX Servers

The following items are required to install the Remote Agent for Linux or UNIX Servers (Remote Agent):

- The media server must have TCP/IP installed.
- You must have a root logon account on the Linux or UNIX servers.
- You must have the Backup Exec installation media.
- You must enter a license key for the Remote Agent on the media server.

Note: Some versions of Linux may require that you install the libstdc++.so.5 package.

See “[Troubleshooting the Remote Agent for Linux or UNIX Servers](#)” on page 1908.

Symantec recommends that you use the Secure Shell (SSH) protocol when you push-install the Remote Agent to remote servers. You must enable SSH before you push-install the Remote Agent.

Backup Exec automatically installs the Remote Media Agent for Linux Servers when it installs the Remote Agent for Linux or UNIX Servers on a Linux server. However, you must enter a separate license key for the Remote Media Agent for Linux Servers before it is available for use.

See “[About the Remote Media Agent for Linux Servers](#)” on page 1966.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

See “[About installing the Remote Agent for Linux or UNIX Servers](#)” on page 1877.

See [“Requirements for backing up Novell Open Enterprise Server on SUSE Linux Enterprise Server”](#) on page 1895.

About installing the Remote Agent for Linux or UNIX Servers

Use the Backup Exec installation media to do the following:

- Install the Remote Agent for Linux or UNIX Servers (Remote Agent) on a local Linux server.
- Push-install the Remote Agent to one or more remote Linux servers.
If you push-install the Remote Agent, the RSH (Remote Shell) is used by default. Symantec recommends that you use SSH (Secure Shell) instead. To use SSH, you must enable it before you install the Remote Agent. Refer to your operating system documentation for more information about SSH.

When you install the Remote Agent, Backup Exec creates the beoper group and adds root as a member. The beoper group contains the names of the users who have permission to back up and restore the Linux or UNIX servers. However, if Backup Exec detects an NIS server during the Remote Agent installation, then the beoper group is not created. You must create the beoper group manually on the Linux or UNIX servers on which you want to install the Remote Agent.

When the installation is complete, Backup Exec saves the install log file to the following location on the server on which the Remote Agent is installed:

`/var/tmp/vxif/installralus<summary file number>/installralus.log`

See [“Installing the Remote Agent for Linux or UNIX Servers”](#) on page 1877.

See [“Troubleshooting the Remote Agent for Macintosh Systems”](#) on page 1926.

Installing the Remote Agent for Linux or UNIX Servers

You can install the Remote Agent for Linux or UNIX Servers (Remote Agent) on a local Linux or UNIX server. You can also push-install it to one or more remote Linux or UNIX servers.

Note: You must unzip the RALUS_RMALS_RAMs_<version number>.gz file on a Linux or UNIX server. The installation does not run if it is unzipped on a computer that runs the Windows operating system.

See [“About installing the Remote Agent for Linux or UNIX Servers”](#) on page 1877.

To install the Remote Agent for Linux or UNIX Servers

- 1 At a Linux or UNIX server, place the Backup Exec installation media in the appropriate drive.
- 2 Log on as root on the server on which you want to install the Remote Agent.
- 3 Navigate to the following directory on the installation media.

<LinuxUnixMac>

- 4 Copy the **RALUS_RMALS_RAM**S_<version number>.gz file in this directory to a directory on the local computer.
- 5 Unzip the file.

For example:

```
gunzip RALUS_RMALS_RAM<version number>.gz
```

- 6 Untar the file.

For example:

```
tar -xf RALUS_RMALS_RAM<version number>.tar
```

- 7 Start the **installralus** script.

For example:

```
./installralus
```

- 8 Do one of the following:

To install the Remote Agent on the local server Press **Enter**.

To install the Remote Agent to one remote server Type the name, IP address, or fully qualified domain name of a Linux or UNIX server.

To install the Remote Agent to multiple remote servers Type the names, IP addresses, or fully qualified domain names of the Linux or UNIX servers. Leave a space between each identifier.

- 9 After the installer checks for a valid Linux or UNIX operating system during the initial system check, press **Enter**.
- 10 Review the package installation summary, and then press **Enter**.
- 11 After the system installation requirements check completes, press **Enter**.
- 12 Start the prerequisites check by pressing **Enter**.

- 13 Type the name, IP address, or fully qualified domain name of the media server (directory host) that you want to back up the Remote Agent.
- 14 Type any additional names, IP addresses, or fully qualified domain names of media servers that you want to back up this Remote Agent.
- 15 Do one of the following:

If the name, IP address, or fully qualified domain name is correct Press **Enter** to continue the installation.

If you want to change the name, IP address, or fully qualified domain name Type **N**, press **Enter**, and then change the information.

- 16 Start the NIS server scan by pressing **Enter**.
- 17 Examine the results of the NIS server scan, and then do one of the following:

If an NIS server is detected The Remote Agent installer cannot create the beoper group. You must create it manually after the Remote Agent installation is complete.
Continue with the next step.

If an NIS server is not detected Use the installer to create the beoper group.
Do the following in the order listed:

- To let the installer create the beoper group, type **y**.
- To select the next available Group ID, type **n**.
- To add the root user account to the beoper group, type **y**.
- Continue with the next step.

- 18 Start the installation by pressing **Enter**.
- 19 After the installation completes, press **Enter** to start the configuration process.
- 20 After the configuration process completes, press **Enter** to save the installation log to the following file:

/var/tmp/vxif/installralussummary file number/installralus.log

- 21 If the Remote Agent installer did not create a beoper group, you must create it.
See [“Creating the Backup Exec operators group manually”](#) on page 1880.
- 22 Start the Remote Agent for Linux or UNIX Servers daemon.
See [“Starting the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.
- 23 Configure the Remote Agent for Linux or UNIX Servers as appropriate.
See [“About configuring the Remote Agent for Linux or UNIX Servers”](#) on page 1881.

About the Backup Exec operators group for the Remote Agent for Linux or UNIX Servers

The Backup Exec operators (**beoper**) group contains the names of the users who have permission to back up and restore the Linux or UNIX servers.

When you install the Remote Agent for Linux or UNIX Servers (Remote Agent), Backup Exec creates the **beoper** group and adds root as a member. Any Linux or UNIX user that you add to the **beoper** group gets the necessary permissions to back up and restore the servers.

However, if an NIS server is detected during the Remote Agent installation, Backup Exec cannot create the **beoper** group. You must create the **beoper** group manually on the Linux or UNIX servers on which you want to install the remote agent. You must create the **beoper** group before you start backup and restore operations. Otherwise, connections fail between the Linux or UNIX servers and the media server.

Before the members of the **beoper** group can perform backup or restore operations, they must have a Backup Exec logon account.

See [“Creating the Backup Exec operators group manually”](#) on page 1880.

See [“Creating a Backup Exec logon account”](#) on page 187.

Creating the Backup Exec operators group manually

You must create a beoper group on each server on which you want to install the Remote Agent for Linux or UNIX Servers.

See [“About the Backup Exec operators group for the Remote Agent for Linux or UNIX Servers”](#) on page 1880.

Note: Ensure that you understand how to set security for groups on Linux or UNIX servers before you assign a Group ID for the beoper group.

Table T-1 How to manually create the beoper group

Step	Action	More Information
Step 1	Navigate to the Linux or UNIX server on which you want to install the Remote Agent. If the Linux or UNIX server is in an NIS domain, navigate to the NIS domain's group file.	Refer to the NIS documentation for information on how to add a group to an NIS domain group file.
Step 2	Create a group with the following case-sensitive name: beoper	See the operating system's documentation for more information about how to create a group.
Step 3	In the beoper group, add the users that you want to have permission to back up and restore the Linux or UNIX server.	See the operating system's documentation for more information about how to add users to a group.
Step 4	Create a Backup Exec logon account for each user that you add to the beoper group.	See “Creating a Backup Exec logon account” on page 187.

About configuring the Remote Agent for Linux or UNIX Servers

Backup Exec creates a file named `ralus.cfg` on each Linux or UNIX server on which the Remote Agent is installed. You can edit the strings, identifiers, and variables in this file to add or edit options for the Remote Agent.

Options that you can edit in the `ralus.cfg` file include the following:

- The port to which the Remote Agent must send publishing messages.
- The logging level for Oracle and DB2 database operations that use the Backup Exec Remote Agent Utility, and for NDMP information.
- The settings to allow the Remote Agent to publish to one or more media servers.

- The files and directories on Linux and UNIX servers that you want to exclude from backups.
- The setting for a Target Service Agent File System backup for a Novell OES.

The ralus.cfg format contains three components. The first component (A) in the following example is a required string.

The second component (B) is a unique identifier followed by an equal sign (=). A unique identifier can consist of sequential numbers, letters, or alpha-numeric characters. For example, 1, 2, 3 or A, B, C. You can also use AA, BB, CC, or A1, A2, B1, B2.

The third component of the ralus.cfg format is the NetBIOS name, fully qualified domain name, or IP address of the media server.

Figure T-1 Example of the ralus.cfg file

```
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List 1=svr.mycompany.com
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List 2=datasrv
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List 3=66.35.250.151
```

- A = Required string
- B = Required and unique identifier (the order or appearance is irrelevant)
- C = File or directory to be excluded

See [“Editing configuration options for Linux, UNIX, and Macintosh computers”](#) on page 1884.

See [“Configuration options for Linux, UNIX, and Macintosh computers”](#) on page 1885.

About publishing Linux, UNIX, and Macintosh computers to media servers

The Remote Agent for Linux or UNIX Servers and the Remote Agent for Macintosh Systems must publish to a media server to get backed up. (Both of these options are referred to the Remote Agent.) During installation, you identify the media server to which you want to publish the Remote Agent. Backup Exec adds this information to the ralus.cfg file. The Remote Agent publishes information to that media server. When the media server receives the published information, the remote Linux, UNIX, and Macintosh computers appears in the media server's backup selections. It is listed under **Favorite Resources**.

The Remote Agent publishes to all of the media servers that are listed in the ralus.cfg file. For each media server that the Remote Agent publishes to, you can specify a local backup network for operations. This backup network is between the media server and the Linux, UNIX, and Macintosh computers. Jobs are then

directed to that local network rather than to a corporate network, so that the backup data traffic is isolated. As a result, other connected networks are not affected when operations are performed between the media server and the Linux, UNIX, and Macintosh computers.

Note: You can delegate Remote Agent jobs to a managed media server when the Central Admin Server Option is installed. To do so, you must publish the Remote Agent to the managed media server.

The Remote Agent publishes the following information to media servers:

- The version of the Remote Agent.
- The IP address and name of the Linux, UNIX, and Macintosh computers.
- Configuration information.

You can edit the `ralus.cfg` file to configure the following settings for publishing:

- Add, edit, or delete media servers to which the Remote Agent can publish.
- Start a new publishing cycle.
- Stop the Remote Agent from publishing.
- Edit the publishing interval.

See [“Adding media servers to which the Remote Agent for Linux, UNIX, and Macintosh can publish information”](#) on page 1883.

See [“About the Favorite Resources node in the backup selections list”](#) on page 284.

See [“About specifying backup networks”](#) on page 402.

Adding media servers to which the Remote Agent for Linux, UNIX, and Macintosh can publish information

You can specify media servers to which the Remote Agent for Linux or UNIX Servers or the Remote Agent for Macintosh Systems can publish information. (Both options are referred to as the Remote Agent.)

See [“About publishing Linux, UNIX, and Macintosh computers to media servers”](#) on page 1882.

Each media server to which the Remote Agent publishes information displays the remote computer in its backup selections.

See [“About the Favorite Resources node in the backup selections list”](#) on page 284.

To add the media servers to which the Remote Agent for Linux, UNIX, and Macintosh can publish information

- 1 Use a text editor to open the following file:

`/etc/VRTSralus/ralus.cfg`

- 2 Add the following string:

`Software\Symantec\Backup Exec For Windows\Backup
Exec\Engine\Agents\Agent Directory List unique identifier number = IP
address or DNS name of media server`

- 3 Save and close the file.

About excluding files and directories from backup jobs for Linux, UNIX, and Macintosh computers

You can exclude specific files and directories on the Linux, UNIX, and Macintosh computers from all backup jobs. Edit the `ralus.cfg` file to specify the excluded files.

Following is an example of strings in the `ralus.cfg` file that excludes files and directories from all backup jobs.

Figure T-2 Example of file and directory exclusions in the `ralus.cfg` format

```
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude1=Adev/*.*  
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude2=Bproc/*.*  
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude3=Cmnt/ns3/pools/  
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude4=Cmnt/ns3/.pools/
```

A = Required string

B = Required and unique identifier (the order or appearance is irrelevant)

C = File or directory to be excluded

To exclude files and directories for specific backup jobs, specify the exclusions in the backup job properties.

See [“Excludes Properties options”](#) on page 305.

See [“Editing configuration options for Linux, UNIX, and Macintosh computers”](#) on page 1884.

Editing configuration options for Linux, UNIX, and Macintosh computers

You can edit configuration options for the Remote Agent for Linux or UNIX Servers or the Remote Agent for Macintosh Systems.

To edit configuration options for Linux, UNIX, and Macintosh computers

- 1 Use a text editor to open the following file:

/etc/VRTSralus/ralus.cfg

- 2 Change the appropriate string in the file.

See [“Configuration options for Linux, UNIX, and Macintosh computers”](#) on page 1885.

Configuration options for Linux, UNIX, and Macintosh computers

You can edit options to configure the Remote Agent for Linux or UNIX Servers or the Remote Agent for Macintosh Systems. (Both options are referred to as the Remote Agent.)

See [“Editing configuration options for Linux, UNIX, and Macintosh computers”](#) on page 1884.

Table T-2 Configuration options for Linux, UNIX, and Macintosh computers

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Agent Browser\TcpIp\AdvertisementPort=6101	Lists the port to which the Remote Agent must send publish and purge messages.
Software\Symantec\Backup Exec for Windows\Backup Exec\Debug\AgentConfig=0	<p>Enables logging for the Remote Agent Utility that Oracle operations use.</p> <p>Values include the following:</p> <ul style="list-style-type: none"> ■ 0 Logging is not enabled. ■ 1 Logging is enabled. Backup Exec automatically generates the log file. <p>This option does not apply to the Remote Agent for Macintosh Systems.</p>

Table T-2 Configuration options for Linux, UNIX, and Macintosh computers
(continued)

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Debug\VXBSAlevel=0	<p>Enables logging for the Remote Agent for Oracle operations.</p> <p>Values include the following:</p> <ul style="list-style-type: none"> ■ 0 Logging is not enabled. ■ 5 Normal logging is enabled. ■ 6 Advanced logging is enabled. Large log files may be created. <p>This option does not apply to the Remote Agent for Macintosh Systems.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\Agents\Advertise All=1	<p>Enables the Remote Agent to publish information to all of the media servers that are listed in the \Agents\Agent Directory List strings.</p> <p>Values include the following:</p> <ul style="list-style-type: none"> ■ 1 The Remote Agent publishes information to every media server in the Agent Directory List. ■ 0 The Remote Agent publishes information to the first media server in the Agent Directory List. If the attempt is successful, the Remote Agent does not publish information to any other media servers. If the attempt is not successful, the Remote Agent attempts to publish information to the next media server in the list. Attempts continue until the Remote Agent reaches the end of the list.

Table T-2 Configuration options for Linux, UNIX, and Macintosh computers
(continued)

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Agents\Advertise Now=0	<p data-bbox="803 352 1239 439">Enables the Remote Agent to start a new publishing cycle after you add or edit any settings in the <code>ralus.cfg</code> file.</p> <p data-bbox="803 453 1085 479">Values include the following:</p> <ul data-bbox="803 496 1239 1147" style="list-style-type: none"><li data-bbox="803 496 1239 696">■ 0 The Remote Agent publishes information according to its regular cycle, set in the string <code>\Agents\Advertising Interval Minutes</code>. Any changes to the <code>ralus.cfg</code> file take effect when a new publishing cycle begins.<li data-bbox="803 708 1239 1147">■ 1 The Remote Agent starts a new publishing cycle. Any changes to the <code>ralus.cfg</code> file take effect immediately. If the media server does not receive the publishing information, the Remote Agent makes ten more attempts. Each attempt to publish information to the media server is one minute apart. If the information is not sent at the end of the ten attempts, the Remote Agent skips that media server until the next publishing cycle. The publishing cycle is the number of minutes set in the string <code>\Agents\Advertising Interval Minutes</code>.

Table T-2 Configuration options for Linux, UNIX, and Macintosh computers
(continued)

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Agents\Advertisement Purge=0	Lets the Remote Agent send a purge message to all of the media servers in the string \Agents\Advertisement Purge. When a media server receives a purge message, it removes the Remote Agent from its Favorite Resources list. The Remote Agent continues to function. Values include the following: <ul style="list-style-type: none"> ■ 0 Do not purge the Remote Agent from any media servers that are listed in the \Agents\Advertisement Purge string. ■ 1 Purge the Remote Agent from one or more media servers in the \Agents\Advertisement Purge string.
Software\Symantec\Backup Exec for Windows\Backup Exec\Agents\Advertising Disabled=0	Enables the Remote Agent to publish to media servers. Values include the following: <ul style="list-style-type: none"> ■ 0 The Remote Agent attempts to publish information to the media servers that are listed in the string \Agents\Agent Directory List. ■ 1 The Remote Agent does not publish information to media servers.
Software\Symantec\Backup Exec for Windows\Backup Exec\Agents\Advertising Interval Minutes=240	Sets the number of minutes that the Remote Agent must wait between publishing cycles. The default number of minutes is 240. The range of minutes is from 1 minute to 720 minutes.

Table T-2 Configuration options for Linux, UNIX, and Macintosh computers
(continued)

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Agents\Agent Directory List 1=<media server name>	<p>Displays the list of NetBIOS names, fully qualified domain names, or IP addresses to which the Remote Agent publishes information.</p> <p>The media server from which the Remote Agent is push installed is added to the Agent Directory List by default.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Agents\Auto Discovery Enabled=1	<p>Adds a media server to the string \Agents\Agent Directory List if the media server performs a backup job with which the Remote Agent is associated.</p> <p>Values include the following:</p> <ul style="list-style-type: none"> ■ 1 Adds the media server that performs the backup job to the Agent Directory List. The Remote Agent can publish information to the media server. ■ 0 The media server that performs the backup job is not added to the Agent Directory List.
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\Logging\RANT NDMP Debug Level=0	<p>Displays the level of verbosity for logging NDMP information for the Remote Agent.</p> <p>Values include the following:</p> <ul style="list-style-type: none"> ■ 0 Logs only the NDMP errors. ■ 1 Logs the NDMP errors and warnings. ■ 2 Logs the NDMP errors, warnings, and message information that is sent between the remote computer and the media server.
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\Encoder=	<p>Displays the encoder that you can add if the default encoder incorrectly displays characters on the user interface.</p>

Table T-2 Configuration options for Linux, UNIX, and Macintosh computers
(continued)

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\ShowTSAFS=	<p>Lets you perform a Target Service Agent file system (TSAFS) backup for applications on Novell Open Enterprise Services. By default, this option is not enabled.</p> <p>The Remote Agent backs up all file systems using the Root object. If ShowTSAFS is enabled, the Novell Open Enterprise Services resource appears in the backup selection list. If you select the whole computer for backup, then redundant backups are performed. Symantec recommends that you do not enable this option.</p> <p>Values include the following:</p> <ul style="list-style-type: none"> ■ Blank or 0 The file system TSA does not appear for backup selection. ■ 1 The file system TSA resource appears for backup selection. <p>This option does not apply to the Remote Agent for Macintosh Systems.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\SystemExclude1=	<p>Lists the files that you want to exclude from all Remote Agent backup jobs.</p> <p>See “About excluding files and directories from backup jobs for Linux, UNIX, and Macintosh computers” on page 1884.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\SystemFSTypeExclude1	<p>Lists the type of file system that you want to exclude from the Remote Agent backup.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\vfmpath= /opt/VRTSralus/VRTSvxms	<p>Displays the path to the Veritas Mapping Service libraries that the Remote Agent uses.</p>

Table T-2 Configuration options for Linux, UNIX, and Macintosh computers
(continued)

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RMAL\DisableRMAL=0	<p>Lets you use the Remote Media Agent for Linux Servers to back up the Linux server on which it is installed. By default, this option is not enabled.</p> <p>If you install the Remote Media Agent to an unsupported version of Linux, the Remote Media Agent is unavailable for use. You cannot create the jobs that run on the devices that are attached to the Linux server. However, you can back up the Linux server by using the Remote Agent for Linux or UNIX Servers component. This component is installed with the Remote Media Agent. You must change the value of this string to 1 to use the Remote Agent for Linux or UNIX Servers component.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 0 You can create backup, restore, and utility jobs on the media server that run on the Linux server's storage devices.■ 1 You can only use the Remote Agent for Linux or UNIX Servers component to back up the Linux server on which it is installed. <p>See "Troubleshooting the Remote Media Agent for Linux Servers" on page 1987.</p>

About backing up data by using the Remote Agent for Linux or UNIX Servers

The following backup methods are supported when you use the Remote Agent for Linux or UNIX Servers (Remote Agent) to back up data:

- Full - Using modified time
- Differential - Using modified time
- Incremental - Using modified time
- Working set

When you use the **Backup Wizard** to specify backup job settings for the Remote Agent for Linux or UNIX Servers, only full backups are supported. If you select any other backup method in the **Backup Wizard**, a full backup runs instead.

See [“Backing up Linux, UNIX, and Macintosh computers”](#) on page 1892.

See [“Backup job options for Linux, UNIX, and Macintosh computers”](#) on page 1918.

See [“Backing up Novell Open Enterprise Server \(OES\) components”](#) on page 1896.

Backing up Linux, UNIX, and Macintosh computers

You can edit the job properties for backing up Linux, UNIX, and Macintosh computers.

Note: Only the backup methods that use the modified date and timestamp are supported for Linux and UNIX servers.

To back up Linux, UNIX, and Macintosh computers

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 On the backup selections list, under **Favorite Resources**, expand **Linux/Unix Servers**.
- 4 In the **Properties** pane, under **Source**, click **Selections**.
- 5 Select the data that you want to back up.
See [“Creating selection lists”](#) on page 296.
- 6 In the **Properties** pane, under **Settings**, click **Linux, UNIX, and Macintosh**.
- 7 Complete the appropriate options.
See [“Backup job options for Linux, UNIX, and Macintosh computers”](#) on page 1918.
- 8 Complete the remaining backup job properties as necessary.
See [“Creating a backup job by setting job properties”](#) on page 332.

Backup job options for Linux, UNIX, and Macintosh computers

You can set backup job options for Linux, UNIX, and Macintosh computers.

See [“Backing up Linux, UNIX, and Macintosh computers”](#) on page 1892.

Table T-3 Backup job options for Linux, UNIX, and Macintosh computers

Item	Description
Preserve change time	<p>Prevents the Remote Agent from changing an object's attributes when a backup occurs. An object is a file or a directory.</p> <p>This option is not selected by default.</p> <p>During a backup, Backup Exec preserves an object's last access timestamp by resetting the last access timestamp to the value before the backup occurred. When Backup Exec modifies the object's last access timestamp, the operating system internally updates the object's ctime.</p> <p>An object's ctime is the time when an object's attributes such as permissions and timestamps are modified. If the Remote Agent does not change the attributes after a backup, the object's ctime is not changed.</p> <p>This option does not affect the object attributes that are set during restore operations.</p>
Follow local mount points	<p>Lets Backup Exec follow local mount points when it backs up data.</p> <p>This option is enabled by default.</p> <p>For more information on local mount points, see your operating system's documentation.</p>

Table T-3 Backup job options for Linux, UNIX, and Macintosh computers
(continued)

Item	Description
<p>Follow remote mount points</p>	<p>Lets Backup Exec follow remote mount points when it backs up data.</p> <p>This option is not selected by default.</p> <p>When you use this option, the following limitations apply:</p> <ul style="list-style-type: none"> ■ The data that is mounted must reside on a computer type that Backup Exec supports. You can find a list of supported operating systems, platforms, and applications at the following URL: http://entsupport.symantec.com/umi/V-269-1 ■ If the mount point leads to an operating system that Backup Exec does not support, contact the operating system's vendor to resolve any issues. <p>For more information on remote mount points, see your operating system's documentation.</p>
<p>Back up contents of soft-linked directories</p>	<p>Backs up the contents of directories that are linked using soft links.</p> <p>You must select the directory that contains the soft links. If you select only the soft link, then only the link is backed up. The data to which the link points is not backed up. You can create a single directory that contains the soft links to data that you want to back up. Then select this option to back up the single directory.</p> <p>Caution: Linux, UNIX, and Macintosh computers use many soft-links, some of which may point to parent directories. In these cases, this option can cause data to be backed up twice, and can cause backup jobs to continue indefinitely.</p> <p>For more information on soft-linked directories, see your operating system documentation.</p>

Table T-3 Backup job options for Linux, UNIX, and Macintosh computers
(continued)

Item	Description
Lock remote files	Lets the Remote Agent have exclusive access to the files on the remote servers that are connected through Network File System (NFS). Locking remote files prevents other applications from modifying the data during the backup.
Backup method for eDirectory	Displays a backup method for backing up eDirectory data for Novell OES on SUSE Linux Enterprise Server. See “Backing up Novell Open Enterprise Server (OES) components” on page 1896. Note: This option is not available for Macintosh computers.

Requirements for backing up Novell Open Enterprise Server on SUSE Linux Enterprise Server

Backup Exec requires the following to back up Novell OES:

- Novell OES must have Service Pack 1 installed.
- Novell OES 2 must have the Target Service Agent for NDS (TSANDS) loaded. The TSANDS protects eDirectory on Novell Open Enterprise Server 2. By default, TSANDS is not loaded on Novell Open Enterprise Server 2. You must manually load TSANDS before eDirectory can appear as a resource that is available for backup. See your Novell documentation for information about how to load TSANDS.
- The Target Service Agents must be enabled for the following:
 - Novell eDirectory
 - Novell iFolder
 - Novell Group Wise
- A local UNIX user name that is the equivalent of the admin-level eDirectory user in the beoper group. Backup Exec does not support eDirectory users. See [“About the Backup Exec operators group for the Remote Agent for Linux or UNIX Servers”](#) on page 1880.

- A Backup Exec logon account that contains the credentials for the equivalent admin-level eDirectory user must exist before you can perform backup jobs for eDirectory.

See [“Backing up Novell Open Enterprise Server \(OES\) components”](#) on page 1896.

See [“Backing up Linux, UNIX, and Macintosh computers”](#) on page 1892.

Novell Open Enterprise Server components that are supported for backup

Backup Exec supports the following Novell Open Enterprise Server (OES) components:

- Novell iFolder
- Novell eDirectory
- Novell Group Wise
- Novell Storage Services (NSS)

See [“Requirements for backing up Novell Open Enterprise Server on SUSE Linux Enterprise Server”](#) on page 1895.

See [“Backing up Novell Open Enterprise Server \(OES\) components”](#) on page 1896.

Backing up Novell Open Enterprise Server (OES) components

The Remote Agent for Linux or UNIX Servers must be installed on the server on which the Novell OES components reside.

See [“Requirements for backing up Novell Open Enterprise Server on SUSE Linux Enterprise Server”](#) on page 1895.

To back up Novell OES components

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 On the backup selections list, under **Favorite Resources**, expand **Linux/Unix Servers**.
- 4 Double-click the Linux or UNIX server that you want to back up.
- 5 If necessary, select a Backup Exec logon account to access the Linux or UNIX server, and then click **OK**.
- 6 Select the appropriate data to back up.

See [“Novell Open Enterprise Server components that are supported for backup”](#) on page 1896.

- 7 In the **Properties** pane, under **Settings**, click **Linux, Unix, and Macintosh**.
- 8 Select the appropriate backup options.
See [“Backing up Linux, UNIX, and Macintosh computers”](#) on page 1892.
- 9 To back up the eDirectory database, in the list for **Backup method for eDirectory**, select a backup method.
See [“About backup methods”](#) on page 274.
- 10 Complete the remaining backup job properties as necessary.
See [“Creating a backup job by setting job properties”](#) on page 332.

Restoring data to Linux, UNIX, and Macintosh computers

You can specify restore job options to restore Linux, UNIX, and Macintosh computers.

Note: You cannot perform a cross-platform restore of HP/UX file system backups for which compression or encryption is enabled. You must restore these backups to their respective platforms.

To restore Linux, UNIX, and Macintosh computers

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **Properties** pane, under **Source**, click **Selections**.
- 4 In the restore selections list, select the appropriate data to restore.
See [“Selections options for restore jobs”](#) on page 609.
- 5 In the **Properties** pane, under **Settings**, click **Linux, UNIX, and Macintosh**.
- 6 Select the appropriate restore options.
See [“Restore job options for Linux, UNIX, and Macintosh computers”](#) on page 1898.
- 7 Start the restore job, or select other restore options from the **Properties** pane.
See [“Restoring data by setting job properties”](#) on page 605.

About restoring Novell OES components

When Backup Exec restores Novell OES components, it restores the entire Novell NDS database to a set of on-disk DIB files. Then, the NDS database is taken offline. The DIB files are renamed to NDS, which overwrites the offline NDS database.

See [“Restoring data to Linux, UNIX, and Macintosh computers”](#) on page 1897.

Restore job options for Linux, UNIX, and Macintosh computers

The following are restore job options for Linux, UNIX, and Macintosh computers.

See [“Restoring data to Linux, UNIX, and Macintosh computers”](#) on page 1897.

Table T-4 Restore job options for Linux, UNIX, and Macintosh computers

Item	Description
Lock remote files	Lets Backup Exec have exclusive access to the files on the remote computers that are connected through the Network File System (NFS). This option is enabled by default.
Restore DIB set	Restores the Directory Information Base (DIB), also known as the Novell directory services (NDS) database.
Activate DIB after verify	Lets Backup Exec rename the database from .RST to .NDS after the verification process completes successfully. If the verify operation fails, the .RST file is deleted and the original .NDS file is kept intact. If you do not select this option, after the database is restored, the .RST file is available for you to perform manual activation or manual disaster recovery.
Open database when finished	Lets Backup Exec open the database after the restore completes. If you want to perform maintenance tasks before the database opens, do not select this option.

Table T-4 Restore job options for Linux, UNIX, and Macintosh computers
(continued)

Item	Description
Verify database after restore	Lets Backup Exec verify the database after the restore completes.
Roll forward log directory	Displays the location of the roll forward log directory.
Leave backup file on disk	Keeps the Novell DIB fileset on the hard drive. See “ About restoring Novell OES components ” on page 1898.

Edit the default backup and restore job options for Linux, UNIX, and Macintosh computers

You can edit the existing default options for all backup and restore jobs for Linux, UNIX, and Macintosh systems.

To edit default backup and restore job options for Linux, UNIX, and Macintosh systems

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Linux, Unix, and Macintosh**.
- 3 Set the appropriate options.

See “[Default backup and restore job options for Linux, UNIX, and Macintosh computers](#)” on page 1899.

Default backup and restore job options for Linux, UNIX, and Macintosh computers

You can set default backup and restore job properties for all jobs on Linux, UNIX, and Macintosh computers.

See “[Edit the default backup and restore job options for Linux, UNIX, and Macintosh computers](#)” on page 1899.

You can find a list of supported operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Table T-5 Default backup and restore job options for Linux, UNIX, and Macintosh computers

Item	Description
<p>Preserve change time</p>	<p>Prevents the Remote Agent from changing an object's attributes when a backup occurs. An object is a file or a directory.</p> <p>This option is not selected by default.</p> <p>During a backup, Backup Exec preserves an object's last access timestamp by resetting the last access timestamp to the value before the backup occurred. When Backup Exec modifies the object's last access timestamp, the operating system internally updates the object's ctime.</p> <p>An object's ctime is the time when an object's attributes such as permissions, timestamps, etc., have been modified. If the Remote Agent does not change the attributes after a backup, the object's ctime is not changed.</p> <p>This option does not affect the object attributes that are set during restore operations.</p>
<p>Follow local mount points</p>	<p>Lets Backup Exec follow local mount points when it backs up data.</p> <p>This option is enabled by default.</p> <p>For more information on local mount points, see your operating system's documentation.</p>

Table T-5 Default backup and restore job options for Linux, UNIX, and Macintosh computers (*continued*)

Item	Description
Follow remote mount points	<p>Lets Backup Exec follow remote mount points when it backs up data.</p> <p>This option is not selected by default.</p> <p>When you use this option, the following limitations apply:</p> <ul style="list-style-type: none"> ■ The data that is mounted must reside on an operating system that Backup Exec supports. ■ If the mount point leads to an operating system that Backup Exec does not support, contact the operating system's vendor to resolve any issues. <p>For more information on remote mount points, see your operating system's documentation.</p>

Table T-5 Default backup and restore job options for Linux, UNIX, and Macintosh computers (*continued*)

Item	Description
<p>Back up contents of soft-linked directories</p>	<p>Backs up the contents of directories that are linked using soft links.</p> <p>You must select the directory that contains the soft links. If you select only the soft link, then only the link is backed up. The data to which the link points is not backed up. You can create a single directory that contains the soft links to data that you want to back up. Then select this option to back up the single directory.</p> <p>Caution: Linux, UNIX, and Macintosh computers use many soft-links, some of which may point to parent directories. In these cases, use of this option can cause data to be backed up twice, and can cause backup jobs to continue indefinitely.</p> <p>For more information on soft-linked directories, see your operating system's documentation.</p>
<p>Backup method for eDirectory</p>	<p>Displays a backup method for backing up eDirectory data for Novell OES on SUSE Linux Enterprise Server.</p> <p>Note: This option is not supported for Macintosh computers.</p> <p>See “Backing up Novell Open Enterprise Server (OES) components” on page 1896.</p>

Table T-5 Default backup and restore job options for Linux, UNIX, and Macintosh computers (*continued*)

Item	Description
Lock remote files	Lets the Remote Agent have exclusive access to the files on the remote servers that are connected through Network File System (NFS). Locking remote files prevents other applications from modifying the data during the backup or restore job.

Uninstalling the Remote Agent for Linux or UNIX Servers

An automated uninstall process for the Remote Agent for Linux or UNIX Servers (Remote Agent) is available on the Backup Exec installation media.

`/opt/VRTS/install/logs/uninstallralus<summary file number>.summary`

To uninstall the Remote Agent for Linux or UNIX Servers

- 1 On the Linux or UNIX server, place the Backup Exec installation media in the appropriate device.
- 2 Log on as root to the server from which you want to uninstall the Remote Agent.
- 3 Navigate to the following directory on the Backup Exec installation media:
<LinuxUnixMac>
- 4 Start the **uninstallralus** script.

For example:

```
./uninstallralus
```

5 Do one of the following:

- | | |
|---|--|
| To uninstall the Remote Agent from one server | Type the name, IP address, or fully qualified domain name of a Linux or UNIX server. |
| To uninstall the Remote Agent from multiple servers | Type the names, IP addresses, or fully qualified domain names of the Linux or UNIX servers. Leave a space between each identifier. |

6 Press **Enter**.

7 After the Remote Agent package check completes successfully, press **Enter**.

8 When you are prompted to uninstall the RALUS packages, press **Enter**.

9 To save the uninstall summary to the following location on the Linux or UNIX server, press **Enter**:

`/opt/VRTS/install/logs/uninstallralus<summary file number>.summary`

Manually uninstalling the Remote Agent for Linux or UNIX Servers

You can manually uninstall the Remote Agent for Linux or UNIX Servers (Remote Agent).

To manually uninstall the Remote Agent for Linux or UNIX Servers

1 Use a terminal session to connect to the Linux or UNIX server as the root user.

2 Change to the following directory:

`/opt/VRTSralus/bin`

For example:

```
cd /opt/VRTSralus/bin
```

3 Delete the following line if it is found in the `/etc/inittab` file:

`/opt/VRTSralus/bin/VRTSralus.init`

For example:

```
rm -r /opt/VRTSralus/bin/VRTSralus.init
```

4 Stop the Remote Agent daemon.

See [“Stopping the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.

5 Remove the Remote Agent package from the Linux or UNIX server.

For example:

Debian GNU/Linux, Ubuntu	<code>dpkg -r VRTSralus</code>
Linux	<code>rpm -e VRTSralus</code>
Sun Solaris	<code>pkgrm VRTSralus</code>

6 Change back to the root directory.

For example:

```
cd /
```

7 Remove the following files:

```
/etc/VRTSralus
```

```
/opt/VRTSralus
```

```
/var/VRTSralus
```

For example:

```
rm -r /etc/VRTSralus /opt/VRTSralus /var/VRTSralus
```

8 Type y if you are prompted to descend into the directories.**9 Type y if you are prompted to delete a directory.****10 Remove runtime scripts if they are present.**

See [“Runtime scripts to remove when manually uninstalling the Remote Agent for Linux or UNIX Servers”](#) on page 1905.

Runtime scripts to remove when manually uninstalling the Remote Agent for Linux or UNIX Servers

When you manually uninstall the Remote Agent for Linux or UNIX Servers (Remote Agent), remove the following runtime scripts if they are present.

Table T-6 Runtime scripts to remove when manually uninstalling the Remote Agent

Operating system	Runtime scripts to remove
Debian, Ubuntu	<p>/etc/rc5.d/S95VRTSralus.init</p> <p>/etc/rc3.d/S95VRTSralus.init</p> <p>/etc/rc2.d/S95VRTSralus.init</p> <p>/etc/init.d/VRTSralus.init</p> <p>For example:</p> <pre>rm /etc/rc5.d/S95VRTSralus.init</pre>
Red Hat Linux, Asianux	<p>/etc/rc.d/rc5.d/S95VRTSralus.init</p> <p>/etc/rc.d/rc3.d/S95VRTSralus.init</p> <p>/etc/rc.d/rc2.d/S95VRTSralus.init</p> <p>/etc/rc.d/init.d/VRTSralus.init</p> <p>For example:</p> <pre>rm /etc/rc.d/rc5.d/S95VRTSralus.init</pre>
Novell Open Enterprise Server 1.0/ SUSE Linux Enterprise Server 9 (32-bit only)	<p>/etc/init.d/rc5.d/SxxVRTSralus.init</p> <p>/etc/init.d/rc3.d/SxxVRTSralus.init</p> <p>/etc/init.d/rc2.d/SxxVRTSralus.init</p> <p>/etc/init.d/VRTSralus.init</p> <p>For example:</p> <pre>rm /etc/init.d/rc5.d/SxxVRTSralus.init</pre>
Novell Open Enterprise Server 2.0/ SUSE Linux Enterprise Server 10 (32-bit and 64-bit)	<p>/etc/init.d/VRTSralus.init,start=2,3,5</p> <p>/etc/init.d/VRTSralus.init</p> <p>For example:</p> <pre>rm /etc/init.d/VRTSralus.init</pre>
Solaris	<p>/etc/rc2.d/S95VRTSralus.init</p> <p>/etc/rc2.d/S91VRTSralus.init</p> <p>/etc/init.d/VRTSralus.init</p> <p>For example:</p> <pre>rm /etc/rc2.d/S95VRTSralus.init</pre>

Starting the Remote Agent for Linux or UNIX Servers daemon

If necessary, you can start the Remote Agent for Linux or UNIX Servers (Remote Agent) daemon after the operating system starts.

To start the Remote Agent for Linux or UNIX Servers daemon

- 1 Use a terminal session to connect to the Linux or UNIX server as the root user.
- 2 Navigate to the following directory:

```
/etc/init.d/
```

For example:

```
cd /etc/init.d/
```

- 3 Start the Remote Agent daemon.

For example:

```
/etc/init.d/VRTSralus.init start
```

Stopping the Remote Agent for Linux or UNIX Servers daemon

You can stop the Remote Agent for Linux or UNIX Servers (Remote Agent) daemon.

See [“Starting the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.

To stop the Remote Agent for Linux or UNIX Servers daemon

- 1 Use a terminal session to connect to the Linux or UNIX server as the root user.
- 2 Navigate to the following directory:

```
/etc/init.d/
```

For example:

```
cd /etc/init.d/
```

- 3 Stop the Remote Agent daemon:

For example:

```
/etc/init.d/VRTSralus.init stop
```

- 4 Restart the daemon when necessary.

Troubleshooting the Remote Agent for Linux or UNIX Servers

If you experience problems with the Remote Agent for Linux or UNIX Servers (Remote Agent) review the following questions and answers.

Table T-7 Troubleshooting the Remote Agent

Question	Answer
Some characters do not appear correctly in the terminal session during the installation. What should I do?	This error occurs when the system location uses a non-English language character-set on the computer on which you install the Remote Agent. You can switch to another location setting of the same language to try to resolve this issue.
The Remote Agent installer is unable to install the Remote Agent. The following error is reported in the installralus log file: VxIF::Error:: Unable to compress files. Hash(0x8711e8)->({GUNZIP}not found on <hostname>	To support the uncompressing of the Remote Agent platform-specific packages, you can install the GNU data compression utility. Install this utility on the computer on which you want to install the Remote Agent. The utility is available at the following URL: http://www.gzip.org
The Remote Agent for Linux or UNIX Servers is installed on a UNIX or Linux server in an NIS domain. Backup Exec is unable to browse resources on the server. What should I do?	Verify if the group line and the password line in the nsswitch.conf file are set to compatibility mode. If they are, then you must configure the /etc/passwd and /etc/group files. Refer to the nsswitch.conf man pages for additional information on how to configure the nsswitch.conf to use compatibility mode. Alternatively, change the password line and the group line to NIS files so that the UNIX or Linux server validates the user through NIS. If the NIS server is unavailable or if the user is not found, the local files are used for validation.

Table T-7 Troubleshooting the Remote Agent (*continued*)

Question	Answer
<p>I cannot load the Remote Agent. When I attempt to load the Remote Agent in console mode, /beremote --log-console shows the following message:</p> <p>ACE_SV_Semaphore_Complex: no space left on device.</p> <p>What should I do?</p>	<p>This issue occurs when the computer reaches its maximum limit on allowable semaphores. It can occur after an unexpected termination of the Remote Agent. When the Remote Agent unexpectedly terminates, it is unable to clean up some of the semaphore resources that it used. Other processes may have caused the use of semaphores to reach the limit. You must restart the computer to safely recover it from this condition.</p> <p>If other processes are running, it may not be feasible to restart the computer. Instead, you can use the commands that let you list and then remove all semaphores in use by the operating system. Be careful when you select semaphores to remove. Semaphores that are in use by the Remote Agent cannot be identified. If you remove semaphores of other programs that are in use, those programs can become unstable.</p> <p>To list semaphores, you can type the following command:</p> <pre>ipcs -a</pre> <p>To remove semaphores for each identifier that is listed, you can type the following command:</p> <pre>ipcrm -s <id></pre>
<p>I cannot load the Remote Agent. When I attempt to load the Remote Agent in console mode, /beremote --log-console shows the following message:</p> <p>Error while loading shared libraries: libstdc++.so.5: cannot open shared object file: No such file or directory.</p> <p>What should I do?</p>	<p>This error indicates that the libstdc++.so.5 library is not in the /usr/lib directory. This library is necessary to let the Remote Agent start and function. To resolve this issue, install the libstdc++5 package.</p> <p>You can install this package from the media on which your copy of Linux was provided. Or, you can run the following command from a computer that has Internet access:</p> <pre>apt-get install libstdc++5</pre> <p>For SUSE Linux Enterprise Server 11, run the following command:</p> <pre>zypper install libstdc++5</pre>

Symantec Backup Exec Remote Agent for Macintosh Systems

This appendix includes the following topics:

- [About the Remote Agent for Macintosh Systems](#)
- [Requirements for the Remote Agent for Macintosh Systems](#)
- [About the Backup Exec admin group on Macintosh systems](#)
- [About installing the Remote Agent for Macintosh Systems](#)
- [About configuring the Remote Agent for Macintosh Systems](#)
- [About backing up data by using the Remote Agent for Macintosh Systems](#)
- [Macintosh restore options](#)
- [Editing the default backup and restore options for Macintosh systems](#)
- [Uninstalling the Remote Agent for Macintosh Systems](#)
- [Troubleshooting the Remote Agent for Macintosh Systems](#)

About the Remote Agent for Macintosh Systems

The Remote Agent for Macintosh Systems (Remote Agent) is installed as a separate add-on component. The Remote Agent enables Windows Servers network administrators to perform backup and restore operations on Macintosh systems

that are connected to the network. The Remote Agent must be installed on the Macintosh systems before you can perform backup or restore operations.

See “[Requirements for the Remote Agent for Macintosh Systems](#)” on page 1912.

See “[About installing the Remote Agent for Macintosh Systems](#)” on page 1914.

Requirements for the Remote Agent for Macintosh Systems

The following are required to install the Remote Agent for Macintosh Systems (Remote Agent):

- The media server must have TCP/IP installed.
- You must be a member of the admin group on the Macintosh system on which you want to install the Remote Agent.
- You must have the Backup Exec installation media.
- You must enter a license key for the Remote Agent on the media server.

Symantec recommends that you use the Secure Shell (SSH) protocol when you push-install the Remote Agent to remote Macintosh systems. You must enable SSH before you install the Remote Agent.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

See “[Installing the Remote Agent for Macintosh Systems](#)” on page 1914.

See “[About the Backup Exec admin group on Macintosh systems](#)” on page 1912.

About the Backup Exec admin group on Macintosh systems

The Backup Exec admin group contains the name of the users who have permission to back up and restore the Macintosh systems.

When you install the Remote Agent for Macintosh Systems, Backup Exec creates the admin group and adds root as a member. Any Macintosh user that you add to the admin group gets the necessary permission to back up and restore the Macintosh systems.

However, if an NIS server is detected during the Remote Agent installation, Backup Exec cannot create the admin group. After the installation, you must create the

admin group manually on the Macintosh system. You must create the admin group before you start backup and restore operations. Otherwise, connections fail between the Macintosh system and the media server.

Before the members of the admin group can perform backup or restore operations, they must have a Backup Exec logon account.

See [“Creating the Backup Exec admin group manually on Macintosh systems”](#) on page 1913.

See [“Creating a Backup Exec logon account”](#) on page 187.

Creating the Backup Exec admin group manually on Macintosh systems

You must create an admin group on each Macintosh system on which you want to install the Remote Agent.

See [“About the Backup Exec admin group on Macintosh systems”](#) on page 1912.

Note: Ensure that you understand how to set security for groups on Macintosh systems before you assign a Group ID for the admin group.

Table U-1 How to manually create the admin group

Step	Action	More Information
Step 1	Navigate to the Macintosh system on which you want to install the Remote Agent. If the Macintosh system is in a NIS domain, navigate to the NIS domain's group file.	See the NIS documentation for information on how to add a group to a NIS domain group file.
Step 2	Create a group with the following case-sensitive name: admin	See the Macintosh documentation for more information about how to create a group.
Step 3	In the admin group, add the users that you want to have permission to back up and restore the Macintosh system.	See the Macintosh documentation for more information about how to add users to a group.
Step 4	Create a Backup Exec logon account for each user that you add to the admin group.	See “Creating a Backup Exec logon account” on page 187.

About installing the Remote Agent for Macintosh Systems

Use the Backup Exec installation media to do the following:

- Install the Remote Agent for Macintosh Systems (Remote Agent) on a local Macintosh system.
- Push-install the Remote Agent to one or more remote Macintosh systems. If you push-install the Remote Agent, the RSH (Remote Shell) is used by default. Symantec recommends that you use SSH (Secure Shell) instead. To use SSH, you must enable it before you install the Remote Agent. See your Macintosh documentation for more information about SSH.

When the installation is complete, Backup Exec saves the install log file to the following location on the system on which the Remote Agent is installed:

`/var/tmp/vxif/installrams <unique identifier number> for installs`

Note: Some characters may not appear correctly in the terminal session during the installation. This error occurs when the system location uses a non-English language character-set on the computer on which you install the Remote Agent. You can switch to another location setting of the same language to try to resolve this issue.

See [“Installing the Remote Agent for Macintosh Systems”](#) on page 1914.

Installing the Remote Agent for Macintosh Systems

You can install the Remote Agent for Macintosh Systems (Remote Agent) on a local Macintosh system. You can also push-install the Remote Agent to one or more remote Macintosh systems.

See [“About installing the Remote Agent for Macintosh Systems”](#) on page 1914.

Note: You must unzip the RALUS_RMALS_RAMs_<version number>.gz file on a Linux, UNIX, or Macintosh computer. The installation does not run if it is unzipped on a computer that runs the Windows operating system.

To install the Remote Agent for Macintosh Systems

- 1 At a Macintosh system, place the Backup Exec installation media in the appropriate drive.
- 2 Navigate to the following directory on the installation media:
<LinuxUnixMac>
- 3 Copy the RALUS_RMALS_RAMs_<version number>.gz file in this directory to a directory on the local system.
- 4 Unzip the file.
For example:

```
gunzip RALUS_RMALS_RAMs_<version number>.gz
```
- 5 Untar the file.
For example:

```
tar -xf RALUS_RMALS_RAMs_<version number>.tar
```
- 6 Open **Finder**, and then browse to **Applications>Utilities**.
- 7 Open **Terminal**.
- 8 Start the **installrams** script.
For example:

```
sudo ./installrams
```
- 9 Enter the password for the user name that is currently logged on.
- 10 Do one of the following:

To install the Remote Agent on a local system	Press Enter .
To install the Remote Agent to one remote system	Type the name, IP address, or fully qualified domain name of a Macintosh system.
To install the Remote Agent to multiple remote systems	Type the names , IP addresses, or fully qualified domain names of the Macintosh systems. Leave a space between each identifier.
- 11 Press **Enter**.
- 12 After the installer checks for a valid Macintosh system operating system during the initial system check, press **Enter**.

- 13 Review the package installation summary, and then press **Enter**.
- 14 After the system installation requirements check completes, press **Enter**.
- 15 Start the prerequisites check by pressing **Enter**.
- 16 Type the name, IP address, or fully qualified domain name of the media server that you want to back up the Remote Agent.
- 17 Press **Enter**.
- 18 Type any additional names, IP addresses, or fully qualified domain names of the media servers that you want to back up this Remote Agent.
- 19 Do one of the following:

If the name, IP address, or fully qualified domain name is correct Press **Enter** to continue the installation.

If you want to change a name, IP address, or fully qualified domain name Type **N**, press **Enter**, and then change the information.

- 20 Start the NIS server scan by pressing **Enter**.
- 21 Examine the results of the NIS server scan, and then do one of the following:

If an NIS server is detected

The Remote Agent installer cannot create the admin group for Backup Exec operators. You must create it manually after the Remote Agent installation is complete.

Continue with the next step.

If a NIS server is not detected

Use the installer to create the admin group.

Do the following in the order listed:

- To let the installer create the admin group, type **y**.
- To select the next available Group ID, type **n**.
- To add the root user account to the admin group, type **y**.
- Continue with the next step.

- 22 Press **Enter** to begin the installation.

- 23 After a message appears stating that the installation has completed successfully, press **Enter**.
- 24 Start the Remote Agent.
See [“Starting the Remote Agent for Macintosh Systems”](#) on page 1924.
- 25 Create the admin group if the installation did not create it automatically.
See [“Creating the Backup Exec admin group manually on Macintosh systems”](#) on page 1913.
- 26 Perform additional configuration as appropriate.
See [“About configuring the Remote Agent for Macintosh Systems”](#) on page 1917.

About configuring the Remote Agent for Macintosh Systems

Backup Exec creates a file named `ralus.cfg` on each Macintosh system on which the Remote Agent for Macintosh Systems (Remote Agent) is installed.

You can edit the following strings, identifiers, and variables for the Remote Agent in the `ralus.cfg` file:

- The port to which the Remote Agent must send publishing messages.
- The settings to allow the Remote Agent to publish to one or more media servers.
- The files and directories on Macintosh systems that you want to exclude from backups.

See [“Editing configuration options for Linux, UNIX, and Macintosh computers”](#) on page 1884.

See [“About publishing Linux, UNIX, and Macintosh computers to media servers”](#) on page 1882.

See [“Adding media servers to which the Remote Agent for Linux, UNIX, and Macintosh can publish information”](#) on page 1883.

See [“About excluding files and directories from backup jobs for Linux, UNIX, and Macintosh computers”](#) on page 1884.

About backing up data by using the Remote Agent for Macintosh Systems

When you use the Remote Agent for Macintosh Systems (Remote Agent) to back up data, only the following backup methods are supported for Macintosh systems:

- Full - Using modified time
- Differential - Using modified time
- Incremental - Using modified time
- Working set

When you use the **Backup Wizard** to specify backup job settings for the Remote Agent, only full backups are supported. If you select any other backup method in the **Backup Wizard**, a full backup runs.

See [“Backing up Macintosh systems”](#) on page 1918.

Backing up Macintosh systems

You can edit the default options for backing up Macintosh systems.

Note: Only the backup methods that use the modified date and timestamp are supported for Macintosh systems.

To back up Macintosh systems

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 On the backup selections list, under **Favorite Resources**, expand **Macintosh Systems**.
- 4 In the **Properties** pane, under **Source**, click **Selections**.
- 5 Select the data that you want to back up.
See [“Creating selection lists”](#) on page 296.
- 6 In the **Properties** pane, under **Settings**, click **Linux, UNIX, and Macintosh**.
- 7 Complete the appropriate options.
See [“Backup job options for Linux, UNIX, and Macintosh computers”](#) on page 1918.
- 8 Complete the remaining backup job properties as necessary.
See [“Creating a backup job by setting job properties”](#) on page 332.

Macintosh restore options

When you restore Macintosh systems, you can enable the option to **Lock remote files**. This option allows exclusive access to the files on the remote systems that

are connected through the Network File System (NFS). This option is enabled by default.

See [“Restoring Macintosh systems”](#) on page 1919.

Restoring Macintosh systems

You can specify restore job properties to restore Macintosh systems.

To restore Macintosh systems

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the **Properties** pane, under **Source**, click **Selections**.
- 4 In the restore selections list, select the appropriate data to restore.
See [“Selections options for restore jobs”](#) on page 609.
- 5 In the **Properties** pane, under **Settings**, click **Linux, UNIX, and Macintosh**.
- 6 To give the Remote Agent exclusive access to the files on the remote system that are connected through the NFS, select **Lock remote files**.
- 7 Start the restore job, or select other restore options from the **Properties** pane.
See [“Restoring data by setting job properties”](#) on page 605.

Editing the default backup and restore options for Macintosh systems

You can use the existing defaults for all backup and restore jobs for Macintosh systems, or you can edit the defaults.

To edit default backup and restore options for Macintosh computers

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Job Defaults**, click **Linux, Unix, and Macintosh**.
- 3 Set the appropriate options.
See [“Default backup and restore job options for Macintosh systems”](#) on page 1919.

Default backup and restore job options for Macintosh systems

You can set default backup and restore job options for all jobs on Macintosh systems.

See [“Editing the default backup and restore options for Macintosh systems”](#) on page 1919.

Table U-2 Default backup and restore job options for Macintosh systems

Item	Description
Preserve change time	<p>Prevents the Remote Agent for Macintosh Systems (Remote Agent) from changing an object's attributes when a backup occurs. An object is a file or a directory.</p> <p>This option is not selected by default.</p> <p>During a backup, Backup Exec preserves an object's last access timestamp by resetting the last access timestamp to the value before the backup occurred. When Backup Exec modifies the object's last access timestamp, the operating system internally updates the object's ctime.</p> <p>An object's ctime is the time when an object's attributes, such as permissions and timestamps, have been modified. If the Remote Agent does not change the attributes after a backup, the object's ctime is not changed.</p> <p>This option does not affect the attributes of the object that are set during restore operations.</p>
Follow local mount points	<p>Lets Backup Exec follow local mount points to back up data.</p> <p>This option is enabled by default.</p> <p>For more information on local mount points, see your operating system's documentation.</p>

Table U-2 Default backup and restore job options for Macintosh systems
(continued)

Item	Description
Follow remote mount points	<p>Lets Backup Exec follow remote mount points to back up data.</p> <p>This option is not selected by default.</p> <p>When you use this option, the following limitations apply:</p> <ul style="list-style-type: none"> ■ The data that is mounted must reside on a system that Backup Exec supports. You can find a list of supported operating systems, platforms, and applications at the following URL: http://entsupport.symantec.com/umi/v-269-1 ■ If the mount point leads to an operating system that Backup Exec does not support, contact the operating system's vendor to resolve any issues. <p>For more information on remote mount points, see your operating system's documentation.</p>

Table U-2 Default backup and restore job options for Macintosh systems
(continued)

Item	Description
<p>Back up contents of soft-linked directories</p>	<p>Backs up the contents of directories that are linked using soft links.</p> <p>You must select the directory that contains the soft links. If you select only the soft link, then only the link is backed up. The data to which the link points is not backed up. You can create a single directory that contains the soft links to data that you want to back up. Then, select this option to back up this single directory.</p> <p>Caution: Linux, UNIX, and Macintosh computers use many soft-links, some of which may point to parent directories. In these cases, use of this option can cause data to get backed up twice, and can cause backup jobs to continue indefinitely.</p> <p>For more information on soft-linked directories, see your operating system's documentation.</p>
<p>Backup method for eDirectory</p>	<p>Displays a backup method for backing up eDirectory data for Novell OES on SUSE Linux Enterprise Server.</p> <p>Note: This option is not supported for Macintosh systems.</p> <p>See “Novell Open Enterprise Server components that are supported for backup” on page 1896.</p>

Table U-2 Default backup and restore job options for Macintosh systems
(continued)

Item	Description
Lock remote files	Lets the Remote Agent have exclusive access to the files on the remote servers that are connected through Network File System (NFS). Locking remote files prevents other applications from modifying the data during the backup or restore job.

Uninstalling the Remote Agent for Macintosh Systems

An automated uninstall process for the Remote Agent for Macintosh Systems (Remote Agent) is available on the Backup Exec installation media.

You can also manually uninstall the Remote Agent.

See [“Manually uninstalling the Remote Agent for Macintosh Systems”](#) on page 1925.

The uninstall summary is saved to the following location on the Macintosh system:

`/var/tmp/vxif/uninstallrams<unique identifier>.summary`

The uninstall log file is saved to the following location on the Macintosh system:

`/opt/VRTS/install/logs/uninstallrams<summary file number>.log`

After the log files are saved, the uninstall process is complete.

To uninstall the Remote Agent for Macintosh Systems

- 1 On a Macintosh system, place the Backup Exec installation media in the appropriate drive.
- 2 On the Macintosh system from which you want to uninstall the Remote Agent, log on using Admin privileges.
- 3 Navigate to the following directory on the Backup Exec installation media:
<LinuxUnixMac>
- 4 Start the **uninstallrams** script.

For example:

```
./uninstallrams
```

5 Do one of the following:

- | | |
|---|--|
| To uninstall the Remote Agent from one system | Type the name, IP address, or fully qualified domain name of the Macintosh system. |
| To uninstall the Remote Agent from multiple systems | Type the names, IP addresses, or fully qualified domain names of the Macintosh systems. Leave a space between each identifier. |

6 Press **Enter**.

7 After the Remote Agent package check completes successfully, press **Enter**.

8 When you are prompted to uninstall the RALUS packages, press **Enter**.

9 When the uninstall process is complete, press **Enter**.

Starting the Remote Agent for Macintosh Systems

You can manually start the Remote Agent for Macintosh Systems.

See [“Stopping the Remote Agent for Macintosh Systems”](#) on page 1924.

To manually start the Remote Agent for Macintosh Systems

- 1 Use a terminal session to connect to the target Macintosh system as the root user.
- 2 From the root prompt, start the VRTSrams service.

For example:

```
SystemStarter start VRTSrams
```

Stopping the Remote Agent for Macintosh Systems

You can manually stop the Remote Agent for Macintosh Systems.

See [“Starting the Remote Agent for Macintosh Systems”](#) on page 1924.

To manually stop the Remote Agent for Macintosh Systems

- 1 Use a terminal session to connect to the target Macintosh system as the root user.
- 2 From the root prompt, stop the VRTSrams service:

For example:

```
SystemStarter stop VRTSrams
```


Manually uninstalling the Remote Agent for Macintosh Systems

You can manually uninstall the Remote Agent for Macintosh Systems (Remote Agent) from Macintosh systems.

You can also use the Backup Exec installation media to uninstall the Remote Agent.

See [“Uninstalling the Remote Agent for Macintosh Systems”](#) on page 1923.

To manually uninstall the Remote Agent for Macintosh Systems

1 Use a logon account with Admin privileges to log on to a terminal session to connect to the Macintosh system.

2 Change to the following directory:

```
/opt/VRTSralus/bin
```

For example:

```
cd /opt/VRTSralus/bin
```

3 Delete the following line if it is found in the `/etc/inittab` file:

```
/opt/VRTSralus/bin/VRTSralus.init
```

For example:

```
rm -r /opt/VRTSralus/bin/VRTSralus.init
```

4 Stop the Remote Agent daemon.

See [“Stopping the Remote Agent for Macintosh Systems”](#) on page 1924.

5 Remove the Remote Agent package from the Linux or UNIX server.

6 Change back to the root directory.

For example:

```
cd /
```

7 Remove the following files:

```
/etc/VRTSralus
```

```
/opt/VRTSralus
```

```
/var/VRTSralus
```

For example:

```
rm -r /etc/VRTSralus /opt/VRTSralus /var/VRTSralus
```

8 Type **y** if you are prompted to descend into the directories.

9 Type **y** if you are prompted to delete a directory.

10 Remove the /Library/StartupItems/VRTSrams folder.

For example:

```
rm -r /Library/StartupItems/VRTSrams
```

11 Type y if you are prompted to delete a directory.

Troubleshooting the Remote Agent for Macintosh Systems

If you experience problems with the Remote Agent for Macintosh Systems (Remote Agent), read the following questions and answers.

Table U-3 Troubleshooting the Remote Agent

Question	Answer
The Remote Agent is installed on a Macintosh system in a NIS domain, but Backup Exec is unable to browse resources on the system. What do I do?	If the group line and the password line in the nsswitch.conf file are set to compatibility mode, additional configuration is necessary. Refer to the nsswitch.conf man pages for additional information on configuring nsswitch.conf to use compatibility mode. Alternately, change the password line and the group line to NIS files so that the Macintosh system validates the user through NIS. If the NIS server is unavailable or the user is not found, the local files are used for validation.

Table U-3 Troubleshooting the Remote Agent *(continued)*

Question	Answer
<p>I cannot load the Remote Agent. When I attempt to load the Remote Agent in console mode, <code>"/beremote --log-console"</code> shows the following message:</p> <p>"ACE_SV_Semaphore_Complex:no space left on device."</p> <p>What should I do?</p>	<p>This issue occurs when the computer reaches its maximum limit on allowable semaphores. It can occur after an unexpected termination of the Remote Agent. When the Remote Agent unexpectedly terminates, it is unable to clean up some of the semaphore resources that it used. Other processes may have caused the use of semaphores to reach the limit. You must restart the computer to safely recover it from this condition.</p> <p>If other processes are running, it may not be feasible to restart the computer. Instead, you can use the commands that let you list and then remove all semaphores in use by the operating system. Be careful when you select semaphores to remove. Semaphores that are in use by the Remote Agent cannot be identified. If you remove semaphores of other programs that are in use, those programs can become unstable.</p>

Symantec Backup Exec Remote Agent for NetWare Systems

This appendix includes the following topics:

- [About the Remote Agent for NetWare Systems](#)
- [Requirements for installing the Remote Agent for NetWare Systems on a NetWare server](#)
- [About installing the Remote Agent for NetWare Systems](#)
- [About backing up NetWare servers](#)
- [About restoring NetWare servers](#)
- [About default options for the Remote Agent for NetWare Systems](#)
- [Saving configuration information for the NetWare server](#)

About the Remote Agent for NetWare Systems

The Symantec Backup Exec Remote Agent for NetWare Systems (Remote Agent) is installed as a separate, add-on component that must be used for the backup and restore of remote NetWare resources.

The Remote Agent allows network administrators for Windows servers to perform backup and restore operations on NetWare servers that are connected to their network.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Because the Remote Agent is also a Client Access License (CAL), it enables you to fully protect your NetWare data and to perform accelerated backups of NetWare data. You cannot select data and special files on resources for a remote NetWare server for backup until a Remote Agent has been installed.

The Remote Agent is a NetWare Loadable Module (NLM) installed on the NetWare server. The Remote Agent provides faster backups by locally performing the tasks that require extensive network interaction. The Remote Agent processes backup data into continuous streams that use Novell's Storage Management Services (SMS). The media server then processes the backup data as a single task. The Remote Agent is fully Novell-SMS compliant.

See “[Requirements for installing the Remote Agent for NetWare Systems on a NetWare server](#)” on page 1930.

See “[About installing the Remote Agent for NetWare Systems](#)” on page 1931.

See “[About backing up NetWare servers](#)” on page 1934.

See “[Setting default options for the Remote Agent for NetWare Systems](#)” on page 1941.

See “[About backing up the NetWare Directory Services \(NDS\)](#)” on page 1936.

Requirements for installing the Remote Agent for NetWare Systems on a NetWare server

The following are required to install the Remote Agent on a NetWare server:

- The Backup Exec media server must have network access by the TCP/IP protocol to the remote NetWare server.
- The computer from which the installation program is running must be able to access the NetWare server.
- You must have administrative privileges on the NetWare server on which you install the agent.

See “[About installing the Remote Agent for NetWare Systems](#)” on page 1931.

About installing the Remote Agent for NetWare Systems

When you install the Remote Agent on the NetWare server, you must do the following:

- Create a directory named BKUPEXEC in the SYS: volume. You can then copy the latest tested versions of the NLMs needed by Backup Exec to process NetWare-specific requests to a Bkupexec/Nlms directory.
- Create the Bestart.ncf and the Bestop.ncf files and place them in the SYS:SYSTEM directory. These files contain commands to load the appropriate NLMs that allow the NetWare server to be backed up.

Note: If you have previously installed the Remote Agent on NetWare servers, run Bestop from each NetWare console before you install the Remote Agent on those servers again.

When Backup Exec is installed, the TCP/IP protocol is selected for use by default. However, you can change the default settings through **Tools>Options>Network and Security** dialog box.

See [“Setting default options for the Remote Agent for NetWare Systems”](#) on page 1941.

See [“Adding BESTART to the Autoexec.ncf file on the NetWare server”](#) on page 1934.

Installing the Remote Agent for NetWare Systems

You can install the Remote Agent for NetWare Systems (Remote Agent) on a local NetWare server.

To install the Remote Agent for NetWare Systems

- 1 At the Backup Exec media server or a workstation that can access the NetWare server, place the Backup Exec installation media in the appropriate drive.
- 2 Log on to the NetWare server on which you want to install the Remote Agent.
- 3 Create a directory named BKUPEXEC in the SYS: volume.
- 4 Navigate to the following directory on the installation media:
BE\Winnt\Install\Netware\Nwagtacc
- 5 Copy the contents of the directory to the SYS:BKUPEXEC directory on the NetWare server.

- 6 Navigate to the following directory on the installation media:
BE\Winnt\Install\Netware\Netware\Nksfiles\en
- 7 Copy the Novell.nks file to the SYS:BKUPEXEC directory on the NetWare server.
- 8 Navigate to the SYS:SYSTEM directory on the NetWare server.
- 9 With a text editor, create a file named Bestart.ncf and add the following:
SEARCH ADD SYS:/BKUPEXEC/NLMS
Load SYS:/BKUPEXEC/NLMS/BKUPEXEC.NLM -!x -tr -to %1 %2

Note: You can remove the -to switch if you do not have a license for the Advanced Open File Option.

- 10 Save the Bestart.ncf file.
- 11 With a text editor, create a file named Bestop.ncf and add the following:
Load SYS:/BKUPEXEC/NLMS/BESTOP.NLM %1 %2
- 12 Save the Bestop.ncf file.
- 13 Do one of the following:

If you have a license for the Advanced Open File Option

Do the following in the order listed.

- Navigate to the SYS:BKUPEXEC directory on the NetWare server.
- With a text editor, create a file named License.ofo.
- Enter the Advanced Open File Option license key with no dashes or spaces.
- Save the License.ofo file.

If you do not have a license for the Advanced Open File Option

Go to step 14.

- 14 Create the Advrtms.dat file.
See [“Creating the Advrtms.dat file”](#) on page 1933.

About publishing NetWare servers to the NetWare agents list

In order for Backup Exec to display a NetWare server in the NetWare Agents list, the agent must publish its existence. Or, you must manually add the servers running the Remote Agent.

When you install the Remote Agent, the Advrtms.dat file in SYS:BKUPEXEC is not included. The Remote Agent cannot publish information to a media server. You must create the Advrtms.dat file in the SYS:BKUPEXEC directory on the NetWare server. Ensure that this file contains all the names or the IP addresses of the Backup Exec media servers that you want to back up the NetWare server.

Running BESTART at the NetWare server automatically loads Novell's SMDR.NLM component, which publishes the availability of the server for backups using the TCP/IP protocol. This protocol must be enabled on the network and in Backup Exec's Network and Security dialog box in order for the servers to be automatically added to the NetWare Agents list.

See [“Setting default backup network and security options”](#) on page 404.

If your network cannot run this protocol, you must manually add the NetWare server names to Backup Exec's server list.

See [“Setting default options for the Remote Agent for NetWare Systems”](#) on page 1941.

To manually add a NetWare server to the Backup Exec User-defined Selections node, you must do one of the following:

- Configure name resolution for your network.
- Publish the NetWare servers on which the Remote Agent is installed on the media server. To configure publishing, you must edit the Advrtms.dat file on each NetWare server you want to protect.

See [“About the User-defined Selections node in the backup selections list”](#) on page 290.

See [“Creating the Advrtms.dat file”](#) on page 1933.

See [“Adding BESTART to the Autoexec.ncf file on the NetWare server”](#) on page 1934.

Creating the Advrtms.dat file

To configure publishing, you must create the Advrtms.dat file. Ensure that this file contains all the names or the IP addresses of the Backup Exec media servers that you want to back up the NetWare server.

See [“About publishing NetWare servers to the NetWare agents list”](#) on page 1933.

To create the Advrtms.dat file

- 1 Navigate to the SYS:BKUPEXEC directory on the NetWare server.
- 2 With a text editor, create a file named Advrtms.dat.
- 3 Add the name or IP address of the NetWare server.
- 4 Save the Advrtms.dat file.

Adding BESTART to the Autoexec.ncf file on the NetWare server

After installing the Remote Agent on the NetWare server, you should load the latest Novell patches. You can also add the BESTART command, which loads the Remote Agent whenever the server is started, to the Autoexec.ncf file.

See [“About publishing NetWare servers to the NetWare agents list”](#) on page 1933.

To add BESTART to the Autoexec.ncf file on the NetWare server

- 1 Add the command BESTART as the last line in the Autoexec.ncf file so that the Remote Agent is automatically started each time the NetWare server is started.
- 2 After you save the Autoexec.ncf file, restart the NetWare server in order for the changes to take effect.

See [“Unloading the Remote Agent for NetWare Systems”](#) on page 1934.

Unloading the Remote Agent for NetWare Systems

If you added BESTART as the last line in the Autoexec.ncf file on the NetWare server, the Remote Agent is automatically loaded whenever the Autoexec.ncf file is executed on the NetWare server. You can unload it by typing a command.

See [“About publishing NetWare servers to the NetWare agents list”](#) on page 1933.

To unload the Remote Agent for NetWare Systems

- 1 At the NetWare system console prompt, type:

```
bestop
```
- 2 Press ENTER.

All NLMs associated with the Agent are unloaded.

About backing up NetWare servers

The first time you access the NetWare servers for backup, you may be prompted for a username and password. The usernames and passwords you enter to gain

initial access to remote servers and workstations are kept in a password database. This database prevents you from having to type usernames and passwords each time you need to access remote devices. It also allows Backup Exec to log on to servers and attach to agent workstations for unattended jobs.

To back up and restore the NetWare file system, you must have an account on the NetWare server with the following rights:

Table V-1 Necessary rights for NFS backup and restore

To do this:	You need these rights:
Backup	Read files File scan Modify file attributes Access control Erase files (only required if you select the full backup method of Back up and delete the files)
Restore	Write files Create files File scan Modify file attributes Access control

To back up, and restore when necessary, the NDS Tree, you must have a user account on the NetWare server that has the following rights to the [Root] object of the NDS Tree:

Table V-2 Necessary rights for NDS backup and restore

To do this:	You need these rights:
Backup	Object rights Supervisor Browse Create Delete Rename Inheritable

Table V-2 Necessary rights for NDS backup and restore (*continued*)

To do this:	You need these rights:
Restore	Property rights, All properties Supervisor Compare Read Write Add Self Inheritable

Note: White check boxes for these rights display with black check marks in the Trustees of [Root] dialog box. With default rights only, these check boxes are gray with gray check marks.

See [“Setting default options for the Remote Agent for NetWare Systems”](#) on page 1941.

About backing up the NetWare Directory Services (NDS)

Novell recommends using replication to provide the first line of protection for NDS in a multi-server installation. Additionally, back up the NDS database on a regular basis in case it is needed to replace objects that have been accidentally deleted.

Note that if you have multiple servers in the NDS tree, the entire NDS can be backed up from any of those servers. You do not need to back up all of NDS from all of the NDS TSAs in the tree unless you are doing it for redundancy purposes.

Depending on your environment (single-server, multi-server, single administrator, or multi-administrator), you must perform replication of partitions and backups to provide protection for NDS.

Following are some backup strategies that can be applied:

- **Single-server strategy.** NDS installations that consist of a single network server must rely completely on Backup Exec to provide protection for the directory database, since the built-in replication feature cannot be used.
You should back up the entire NDS database whenever any type of backup (either full or modified) is performed. If the NDS database rarely changes, that is, if the objects stored within and/or their properties and values are seldom modified, then less frequent backups may be performed.

As with file system backups, you must consider what might be lost if a disaster occurs on the day the next full backup is to be performed. Be sure to figure in the time it takes to rebuild the changes to the directory manually, if such a disaster occurs.

- **Single administrator - multiple servers strategy.** NDS installations that have a single network administrator (a single object with supervisor rights to the entire directory database), and multiple servers should rely almost entirely on the built-in replication features of NDS for fault tolerance. If a disaster occurs on a specific server, NDS remains intact and available from replicas that are stored on other servers. When the failed server is repaired, NDS is reinstalled using Novell's NWCONFIG.NLM on NetWare 5.x or later. Replicas are then placed back onto the server, if required.

The NDS database should still be backed up regularly in case it is needed to replace the objects that have been accidentally deleted.

- **Multiple administrator strategy.** NDS installations that have multiple network administrators, each with access to only a portion of the directory tree, are faced with additional challenges when designing a backup strategy. Within this type of installation, it is rare that an object has full rights to the entire directory tree, as is the case with many smaller- to medium-sized networks. Instead, the tree is logically broken into smaller components. For example, partitions with specific administrators assigned the responsibility to manage each component. While this type of installation offers the highest level of network security, it brings with it the most complicated level of disaster recovery.

The best method for implementing fault tolerance should remain partition replication. Because it is likely that Inherited Rights Filters (IRFs) are applied at the container level, a properly replicated directory offers a much quicker restoration in the event of a disaster. If possible, you should create an object that has full rights as a trustee of the root of the NDS tree, and perform full backups on the NDS tree, instead of partial backups. Doing so reduces the complexity of rebuilding NDS in the event of a disaster.

You should refer to your Novell documentation for more information on configuring and managing NDS replicas and partitions.

See [“About backing up NetWare servers”](#) on page 1934.

See [“Setting default options for the Remote Agent for NetWare Systems”](#) on page 1941.

Backing up NetWare servers

The following procedure provides details on how to back up NetWare servers. Full, differential, and incremental jobs that specify using the modified time reverts to using the archive bit for NetWare servers that are included in the job.

Note: Backup Exec does not support backing up double-byte character sets for NetWare servers that have a double-byte code page loaded. Software encryption is also not supported.

See [“About backing up NetWare servers”](#) on page 1934.

See [“About the User-defined Selections node in the backup selections list”](#) on page 290.

See [“About encryption”](#) on page 415.

See [“About restoring NetWare servers”](#) on page 1939.

To back up a NetWare server

- 1 On the navigation bar, click the arrow next to Backup.
- 2 Click **New Backup Job**.
- 3 On the **Properties** pane, under **Source**, click **Selections**.
- 4 In the backup selections tree, expand **User-defined Selections**.

When logging on to the NetWare server, you may need to provide a fully distinguished and a typeless name, such as .admin.novell.

A fully distinguished, or complete, name consists of different object types, such as common name (CN), Organizational Unit (OU) objects, and Organization (O) objects. When the abbreviations for these objects are not included as part of the object’s complete name, the naming is referred to as a typeless name. For more information about complete, partial, typeful, or typeless names, refer to your Novell NetWare documentation.

- 5 Select the NetWare resource you want to back up.
NetWare File System and NetWare Directory Services (Novell Directory) are listed separately. Each directory that you want to back up must be selected.
- 6 If you want to use hardware encryption, do the following steps in order:
 - On the **Properties** pane, under **Settings**, click **Network and Security**.
 - Select **Hardware** as the Encryption type.
 - Select or create an encryption key.

- 7 If you want to change the backup default, on the **Properties** pane, under **Settings**, click **NetWare SMS**.
- 8 Select or clear **Back up compressed files in decompressed form**.

If you select this option, Backup Exec decompresses, or expands, compressed files as they are backed up. If you select this option, the server may run out of memory or disk space. Also, the backup job takes longer due to the extra time that is involved in file decompression.
- 9 After selecting job options, start the backup job or select other backup options from the **Properties** pane, and then start the backup job.

See [“Creating a backup job by setting job properties”](#) on page 332.

NetWare SMS backup options

The **Back up compressed files in decompressed form** option lets you decompress, or expand, compressed files as they are backed up. If you select this option, the server may run out of memory or disk space. Also, the backup job takes longer due to the extra time that is involved in file decompression.

About restoring NetWare servers

Before restoring your NetWare server, you may want to read about restore operations in general.

See [“Restoring data by setting job properties”](#) on page 605.

If you have more than one server in the NDS tree, it is not necessary to restore NDS since a replica should be available from another server. The only time an NDS restore operation needs to be done is to replace the objects that have been deleted accidentally.

Because information about partitions and replicas would probably change between a backup of NDS and any subsequent restores, this information is not saved by SMS when a backup of NDS is performed. Thus, when NDS is backed up, it appears as though all objects are stored in a single partition.

However, if information about partitions is available when the restore operation is performed, objects are restored to the proper partition.

See [“Restoring NetWare servers”](#) on page 1939.

Restoring NetWare servers

You can restore data to the NetWare server from which the data was backed up, or to another server.

If you are redirecting a restore operation, note the following:

- Only data can be included in a redirected restore operation; NDS objects cannot be redirected.
- The data that is backed up from a Novell server can be restored to a Windows volume; again, NDS objects cannot be redirected.

See [“About restoring NetWare servers”](#) on page 1939.

To restore NetWare servers

- 1 On the navigation bar, click the arrow next to Restore.
- 2 Click **New Restore Job**.
- 3 Select the data you want to restore.
See [“About selecting data to restore”](#) on page 626.
- 4 If you want to change the Restore option default, on the **Properties** pane, under **Settings**, click **NetWare SMS**.
- 5 Select or clear **Restore volume restriction**.

If you select this option, Backup Exec restores NetWare volume restrictions. Restoring volume restrictions is not recommended unless you are performing disaster recovery.

- 6 (Optional) If you want to redirect the restore to another server, under Destination, click **File Redirection** and complete the options.

See [“File Redirection restore options”](#) on page 634.

If you restore NetWare data to a Windows volume, trustee data that is associated with the files is not restored. If the file was compressed by NetWare and was backed up in a compressed format, you cannot restore it to a Windows volume.

- 7 Start the restore job or select other restore options from the **Properties** pane.

If you clear the Preserve tree option on the General Restore Job Properties and the target directory is the volume root, Backup Exec still uses the Preserve tree option and data is restored with its original directory intact.

See [“Restoring data by setting job properties”](#) on page 605.

About default options for the Remote Agent for NetWare Systems

By default, Backup Exec detects NetWare servers that are publishing using the TCP/IP protocol. If these protocols are made unavailable, the NetWare remote

agents are not detected. And, the NetWare Agents node does not appear under Favorite Resources in the backup selections tree.

Note: If a protocol is not installed on the system, it is not available in this dialog. For example, if the TCP/IP protocol is not installed on the media server, the TCP/IP protocol check box is grayed out and made unavailable.

You can also set network defaults for all backup and restore operations that are performed on the NetWare servers by Backup Exec. For example, you can specify a dynamic port range to be used by the remote agent. You can override some of these defaults each time you create a backup or restore job.

See [“Setting default options for the Remote Agent for NetWare Systems”](#) on page 1941.

See [“Specifying TCP dynamic port ranges on the media server”](#) on page 1943.

Setting default options for the Remote Agent for NetWare Systems

The following procedure provides details on how to set backup and restore options for NetWare.

See [“About default options for the Remote Agent for NetWare Systems”](#) on page 1940.

See [“Saving configuration information for the NetWare server”](#) on page 1943.

To change backup and restore defaults for the NetWare server

- 1 On the Tools menu, click **Options**.
- 2 On the Properties pane, under Job Defaults, click **NetWare SMS**.
- 3 Select the appropriate options.
See [“NetWare SMS default options”](#) on page 1941.
- 4 Click **OK**.

NetWare SMS default options

You can set default options for all backup and restore jobs that use the NetWare Agent.

See [“Setting default options for the Remote Agent for NetWare Systems”](#) on page 1941.

The following table describes the NetWare SMS default options:

Table V-3 NetWare SMS default options

Item	Description
Display the following servers	<p>Lets you select which servers Backup Exec displays. Backup Exec checks the registry for a list of NetWare servers. If the list does not exist, Backup Exec creates it using the wildcard (*) default. The wildcard lets all servers that are published using the Service Location Protocol (TCP/IP protocol) be seen. Backup Exec displays these servers in this field.</p>
Add	<p>Lets you add a media server to the servers list.</p> <p>If you add a server name to the list, Backup Exec must be able to resolve the name to a TCP/IP address. If Backup Exec cannot resolve the name to a TCP/IP address, the server name appears in the servers list. However, Backup Exec is unable to connect to it. Backup Exec can resolve the name if the NetWare server names and IP addresses are in your network's Domain Naming Services (DNS) database. If these names and IP addresses are not in DNS, you must manually add the names and IP addresses to the media server's HOSTS file. This file usually is found in the \WINDOWS\SYSTEM32\Drivers\ETC directory.</p> <p>For these changes to take effect, you must restart the Backup Exec administration console.</p>
Delete	<p>Lets you remove a media server from the servers list.</p>
Back up compressed files in decompressed form	<p>Decompresses, or expands, compressed files as they are backed up. If you select this option, the server may run out of memory or disk space. Also, the backup job takes longer due to the extra time that is involved in file decompression.</p> <p>In most cases, this option should not be selected.</p>

Table V-3 NetWare SMS default options (*continued*)

Item	Description
Restore volume restrictions	Restores NetWare volume restrictions. Restoring volume restrictions is not recommended unless you are performing a disaster recovery job.

Specifying TCP dynamic port ranges on the media server

The following procedure provides details on how to specify TCP dynamic port ranges on the media server.

See [“About default options for the Remote Agent for NetWare Systems”](#) on page 1940.

To specify TCP dynamic port ranges on the media server

- 1 On the **Tools** menu, click **Options**.
- 2 On the **Properties** pane, under **Job Defaults**, click **Network and Security**.
- 3 Specify a TCP dynamic port range by clicking **Enable remote agent TCP dynamic port range** and enter the port ranges.
- 4 Click **OK**.
- 5 Restart Backup Exec.

Saving configuration information for the NetWare server

Use the Bediag.nlm utility to create an ascii file called Bediag.fax that includes useful configuration information for your server.

The information in this file includes the following:

- The contents of your Config.sys and Autoexec.bat files.
- Contents of the Startup.ncf file.
- The amount of memory available.
- The contents of your Autoexec.ncf file.
- A listing of the NLMs that are currently loaded on your server, including the version numbers and the date stamp.
- Configuration settings for your server, including volumes and individual namespace support.

Keep a copy of the Bediag.fax available so that if you have to contact Technical Support, you can quickly provide system configuration information.

To save configuration information for the NetWare server

- 1 At the NetWare system console prompt, type:

```
load SYS:BKUPEXEC/NLMS/BEDIAG
```

The Bediag.fax file is created.

You can use the following options when loading Bediag.nlm: (for example,

```
load bediag /c)
```

/c - outputs the file to the screen

/s - gathers information for SCSI devices only

/n - exclude information for SCSI devices

- 2 View the Bediag.fax file with a text editor or word processor.
- 3 On the print-out of the Bediag.fax, write the Supervisor user and password.
Keep this print-out locked in a safe place.

Symantec Backup Exec Remote Agent for Windows Systems

This appendix includes the following topics:

- [About the Remote Agent for Windows Systems](#)
- [Requirements for the Remote Agent for Windows Systems](#)
- [Stopping and starting the Remote Agent for Windows Systems](#)
- [About the Remote Agent Utility for Windows Systems](#)
- [Configuring database access](#)
- [About the Remote Agent Utility Command Line Applet](#)

About the Remote Agent for Windows Systems

The Backup Exec Remote Agent for Windows Systems (Remote Agent) is installed as a separate add-on component. The Remote Agent enables Windows Servers network administrators to perform backup and restore operations on Windows resources that are connected to the network.

The Remote Agent is a system service that runs on remote Windows servers and workstations. The Remote Agent provides faster backup processing by locally performing tasks that in typical backup technologies, require extensive network interaction. The Remote Agent processes backup data into a continuous stream that the media server then processes as a single task. This method provides better

data transfer rates over traditional technologies, which require multiple requests and acknowledgments between the media server and the remote server.

The Remote Agent enables you to do the following:

- Back up and restore in firewall environments.
- Back up and restore using a specified local network if the media server and the remote computer are on the same subnet.
- Display the remote computer in the media server's Favorite Resources node.
- Attain significant performance increases when running modified backups (for example, differential and incremental). This occurs because file selection is performed locally by the Remote Agent instead of across the network as performed by traditional network backup applications.

Note: Network hardware has a major impact on performance. Performance is directly related to the capabilities of the networking hardware in the media server and the remote device. Higher network bandwidth ratings also contribute to faster operation processing.

See [“Requirements for the Remote Agent for Windows Systems”](#) on page 1946.

See [“About installing the Remote Agent for Windows Systems”](#) on page 139.

See [“Setting default backup network and security options”](#) on page 404.

See [“About using Backup Exec with firewalls”](#) on page 409.

See [“About the Backup Exec Shadow Copy Components file system”](#) on page 320.

See [“About the Remote Agent Utility for Windows Systems”](#) on page 1948.

Requirements for the Remote Agent for Windows Systems

Because a Remote Agent is also a Client Access License (CAL), you must install the Remote Agent on any remote Windows computer that you want to back up. You cannot fully protect resources on a remote server until a Remote Agent has been installed.

At the Backup Exec media server, you must enter Remote Agent license keys for each remote Windows computer that you want to protect. To back up a remote Windows computer from more than one media server, you must enter the same Remote Agent license key on each media server.

Backup Exec database agents also include a Remote Agent that allows you to protect one remote Windows computer. The Remote Agent license is enabled when you install the database agents on the media server.

To protect the Workstation versions of the supported Windows platforms, you must install the Remote Agent on each platform.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Note: If a previous version of the Remote Agent is installed, it is automatically upgraded when you initiate a new Remote Agent installation. Previous versions of the Remote Agent are automatically detected on the remote computers and replaced with the new version during installation of the new Remote Agent. The name of the system service may have changed when the upgrade is complete.

You can install the Remote Agent for Windows Systems using many methods, depending on your environment.

See “[About installing the Remote Agent for Windows Systems](#)” on page 139.

Stopping and starting the Remote Agent for Windows Systems

The Remote Agent is automatically started as a service when Windows is started on the remote computer.

To stop or start the Remote Agent for Windows Systems

1 Do one of the following:

On Windows 7/Vista/Server 2008
R2/Server 2008 computers

Right-click **Computer**.

On a Windows Server 2003 computer

Right-click **My Computer**.

2 Click **Manage**.

3 Do one of the following:

On a Windows Server 2008 R2/Server 2008 computer

On the **Server Manager** dialog box, expand **Configuration**.

On Windows 7/Vista/Server 2003 computers

On the **Computer Management** dialog box, double-click **Services and Applications**.

4 Click **Services**.

5 In the Results pane, right-click **Backup Exec Remote Agent for Windows Systems**.

6 Do one of the following:

To stop the Remote Agent

Click **Stop** to stop the Remote Agent.

To start the Remote Agent

Click **Start** to start the Remote Agent.

About the Remote Agent Utility for Windows Systems

The Remote Agent Utility is installed when the Remote Agent is installed on a remote Windows computer.

You can perform the following tasks with the Remote Agent Utility:

- Start the Remote Agent Utility each time you log on.
See [“Starting the Remote Agent Utility”](#) on page 1949.
- View current activity on the remote Windows computer.
See [“Viewing the activity status of the remote computer in the Remote Agent Utility”](#) on page 1949.
- Configure the Remote Agent to send information about itself, such as its version and IP address, to a media server.
See [“About publishing the Remote Agent for Windows Systems to media servers”](#) on page 1951.
- Configure the Remote Agent Utility for backup and restore operations of Oracle instances.
See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1318.
- Configure the Remote Agent Utility for backup and restore operations of DB2 instances.

See [“Configuring the DB2 Agent on Windows computers”](#) on page 955.

- Configure the Remote Agent Utility for media server database access for Oracle and DB2 operations.

See [“Configuring database access”](#) on page 1955.

Starting the Remote Agent Utility

You access the Remote Agent Utility from the Windows taskbar.

See [“Viewing the activity status of the remote computer in the Remote Agent Utility”](#) on page 1949.

See [“About publishing the Remote Agent for Windows Systems to media servers”](#) on page 1951.

To start the Remote Agent Utility

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

When the Remote Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 To open the Registry Editor, the Services window, and the Event Viewer on the remote Windows computer, right-click the Remote Agent Utility icon in the system tray, and then click **Tools**.

Viewing the activity status of the remote computer in the Remote Agent Utility

You can use the Remote Agent Utility to view the activity status of the remote Windows computer.

To view the activity status of the remote computer in the Remote Agent Utility

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

If the Remote Agent Utility is already running, you can double-click its icon in the system tray.

- 2 Click the **Status** tab.

See [“Status options for the Remote Agent Utility”](#) on page 1950.

- 3 Click **OK**.

Status options for the Remote Agent Utility

You can set the following status options for the Remote Agent Utility.

See [“Viewing the activity status of the remote computer in the Remote Agent Utility”](#) on page 1949.

Table W-1 Remote Agent Utility Status options

Item	Description
Start the Remote Agent Utility every time you log on	Indicates if the Remote Agent Utility displays when you log on to this computer.
Refresh interval	Displays the number of seconds for the Remote Agent Utility to wait before refreshing the status of the computer. The default setting is to refresh every 5 seconds.
Media server	Displays the name of the media server that is processing the current operation.
Source	Displays the media or share that is being processed.
Current folder	Displays the name of the current directory, folder, or database (depending on the specific agent) that is being processed.
Current file	Displays the name of the current file that is being processed.

Viewing the activity status of the remote computer from the system tray

You can view the activity status for a remote computer.

Possible statuses are as follows:

- A backup job is running
- A restore job is running
- A backup and a restore job are running
- Snapshot in progress
- The Backup Exec client service, Beremote.exe, is not running on the computer
- Idle

To view the activity status of a remote computer

- ◆ Position the cursor over the Remote Agent icon in the system tray.

Starting the Remote Agent Utility automatically on the remote computer

You can start the Remote Agent Utility automatically each time you log on to the remote computer.

See [“Status options for the Remote Agent Utility”](#) on page 1950.

To start the Remote Agent Utility automatically on the remote computer

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

If the Remote Agent Utility is already running, you can double-click its icon in the system tray.

- 2 Click the **Status** tab.
- 3 Check the **Start the Remote Agent Utility every time you log on** check box.
- 4 Click **OK**.

Setting the refresh interval on the remote computer

You can display the number of seconds for the Remote Agent Utility to wait before refreshing the status of the computer.

See [“Status options for the Remote Agent Utility”](#) on page 1950.

To set the refresh interval on the remote computer

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

If the Remote Agent Utility is already running, you can double-click its icon in the system tray.

- 2 Click the **Status** tab.
- 3 In the **Refresh interval** box, type the number of seconds to refresh the status.
- 4 Click **OK**.

About publishing the Remote Agent for Windows Systems to media servers

Use the Remote Agent Utility to add, change, or delete the media server names or IP addresses that this remote Windows computer publishes to. Each media

server that you add to the list on the Publishing tab displays this remote computer in its backup selection tree, under Favorite Resources.

This information that the Remote Agent publishes includes the version of the Remote Agent and the remote computer's IP addresses. Because the remote computer's IP address is published to the media server, the media server can connect to and display the remote computer even if it is in an unknown domain.

For each media server that is published to, you can specify a local backup network for operations between the media server and the remote computer. Directing jobs to a specified local network rather than to a corporate network isolates the backup data traffic so that other connected networks are not affected when operations are performed between the media server and the remote computer.

See [“About specifying backup networks”](#) on page 402.

See [“Adding media servers that the Remote Agent for Windows Systems can publish to”](#) on page 1952.

See [“Editing media server information that the Remote Agent for Windows Systems publishes to”](#) on page 1954.

See [“Removing media servers that the Remote Agent for Windows Systems can publish to”](#) on page 1955.

Adding media servers that the Remote Agent for Windows Systems can publish to

You can use the Remote Agent Utility to add a media server to which the Remote Agent can publish information.

See [“About publishing the Remote Agent for Windows Systems to media servers”](#) on page 1951.

See [“About the Favorite Resources node in the backup selections list”](#) on page 284.

See [“Adding a Windows system to the Favorite Resources node in the backup selections list”](#) on page 285.

See [“Deleting a Windows system from the Favorite Resources node in the backup selections list”](#) on page 286.

See [“Viewing the activity status of the remote computer in the Remote Agent Utility”](#) on page 1949.

To add media servers that the Remote Agent can publish to

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

When the Remote Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 Click the **Publishing** tab.
- 3 Complete the appropriate options.
See [“Remote Agent Utility Publishing options”](#) on page 1953.
- 4 Click **OK**.

Remote Agent Utility Publishing options

You can set the following publishing options for the Remote Agent Utility.

See [“Adding media servers that the Remote Agent for Windows Systems can publish to”](#) on page 1952.

Table W-2 Remote Agent Utility Publishing options

Item	Description
<p>Enable the Remote Agent to publish information to the media servers in the list</p>	<p>Indicates if the remote agent sends information about itself, such as its version and IP address, to all of the media servers in the list. The media servers display the remote agent in the backup selections tree, under Favorite Resources and under Domains.</p> <p>By default, the name of the media server that push-installed this remote agent is displayed in this list. If the remote agent is also a media server, the name is displayed as 127.0.0.1.</p> <p>To stop the information from being sent to all of the media servers, uncheck Enable the Remote Agent to publish information to the media servers in the list. The list of media servers is preserved, but the Remote Agent does not send any information about itself to the media servers.</p>
<p>Publishing interval</p>	<p>Displays an interval, in minutes, for the Remote Agent to send information about its status to the media servers in the list. The default interval is 240 minutes. This is the recommended setting to appropriately balance system responsiveness with network traffic. The maximum interval allowed is 720 minutes.</p>

Table W-2 Remote Agent Utility Publishing options (*continued*)

Item	Description
Change Settings	<p>Enables the settings to let you add, edit, or remove media servers in the media servers list.</p> <p>This option appears the first time you start the Remote Agent Utility.</p>
Add	<p>Lets you add the media server name or IP address to the media servers list.</p>
Edit	<p>Lets you edit a name or address in the media servers list.</p>
Remove	<p>Lets you delete a media server name or IP address from the media servers list. The Remote Agent no longer publishes information to the media server. You cannot select the remote computer for backup from the media server's Favorite Resources node.</p>
Published names for this agent	<p>Shows the names that are used when this remote computer is published. The names appear under a media server's Favorite Resources.</p> <p>These names can include the following:</p> <ul style="list-style-type: none"> ■ The fully qualified domain name. ■ The computer name. ■ The NetBIOS computer name. ■ Virtual Service names, which are the names given to clustered resources that are hosted by the remote computer. ■ Oracle RMAN Real Application Cluster (RAC) name, which is the virtual name used by computers in a RAC for the computer that hosts the Oracle application. This name is displayed in a media server's backup selection list under the Oracle RAC node.

Editing media server information that the Remote Agent for Windows Systems publishes to

You can use the Remote Agent Utility to edit a media server name or IP address to which the Remote Agent can publish information.

See [“About publishing the Remote Agent for Windows Systems to media servers”](#) on page 1951.

To edit media server information

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

When the Remote Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 Click the **Publishing** tab.
- 3 Select the media server that you want to edit from the list.
- 4 Click **Edit**.
- 5 Edit the media server name or IP address.
- 6 Click **OK**.

Removing media servers that the Remote Agent for Windows Systems can publish to

You can use the Remote Agent Utility to remove a media server so that the Remote Agent no longer publishes information to it.

See [“About publishing the Remote Agent for Windows Systems to media servers”](#) on page 1951.

To remove media servers that the Remote Agent can publish to

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.

When the Remote Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 Click the **Publishing** tab.
- 3 Select the media server that you want to remove from the list.
- 4 Click **Remove**.
- 5 Click **OK**.

Configuring database access

You can configure database access to enable the media server to authenticate Oracle and DB2 operations.

See [“Setting authentication credentials on the media server for Oracle operations”](#) on page 1330.

See [“Adding information about DB2 servers to the media server's authentication credentials list”](#) on page 955.

To configure database access

- 1 On the computer on which the Remote Agent is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec 2010 Remote Agent Utility**.
- 2 Click the **Database Access** tab.
- 3 To make changes, click **Change Settings**.
- 4 Complete the appropriate options.
See [“Database access options for the Remote Agent Utility”](#) on page 1956.
- 5 Click **OK**.
- 6 On the media server, add the name of the Oracle or DB2 server and the user name that you entered on the Database Access tab to the media server's list of authentication credentials.

Database access options for the Remote Agent Utility

You can set the following database access options for the Remote Agent Utility.

See [“Configuring database access”](#) on page 1955.

Table W-3 Remote Agent Utility database access options

Item	Description
Enable media server authentication for Oracle and DB2 operations	<p>Specifies the credentials that the media server will use for all operations on the Oracle and DB2 servers, including DBA-initiated operations. The media server also uses these credentials for authentication of the Oracle and DB2 servers.</p> <p>You must check this to enable DB2 and Oracle operations between the media server and this computer.</p>

Table W-3 Remote Agent Utility database access options (*continued*)

Item	Description
User name	<p>Specifies a user name that has administrative rights to this computer. This logon account is what the media server uses when it connects to this computer.</p> <p>If you specify an IP address or a fully qualified computer name as part of the user name, the Remote Agent Utility may not be able to verify the user account. If the credentials entered are incorrect, the error “cannot attach to a resource” may be displayed when you run a backup or restore job.</p> <p>You must add this computer name and logon account to the media server’s list of authentication credentials for Oracle and DB2 servers. If the authentication fails when the Oracle resources are backed up, the backup job fails. If the authentication fails when you are browsing the backup sets for a restore job, then the backup sets become unavailable, and you must run a DBA-initiated restore job to restore data.</p>
Password	<p>Specifies the password for this logon account.</p> <p>Note: For security reasons, the logon credentials are not stored on the remote computer.</p>
Confirm Password	<p>Specifies the password again to confirm it.</p>
Use a custom port to connect to the media server during Oracle and DB2 operations	<p>Designates the port that is used for communications between this computer and the media server during Oracle or DB2 operations. By default, port 5633 is used.</p> <p>If you change the port number on this computer, you must also change it on the media server, and then restart the Backup Exec Job Engine Service on the media server.</p>

Table W-3 Remote Agent Utility database access options (*continued*)

Item	Description
Port number	Specifies the port number to use for operation requests that are sent to the media server.

About the Remote Agent Utility Command Line Applet

You can use the Remote Agent Utility Command Line Applet from any Windows operating system command prompt to access the Remote Agent Utility. The Remote Agent Utility Command Line Applet is installed when you install the Remote Agent. If you run the command line utility on a Windows 7/Vista/Server 2008 R2/Server 2008 computer, you must run it in elevated command prompt.

Note: To run the Remote Agent Utility Command Line Applet on a Microsoft Windows Server 2008 R2/Server 2008 computer, you must use Server Core.

You can run the following Remote Agent Utility functions with the Remote Agent Utility Command Line Applet:

- Set the publishing interval (in minutes).
- List the published name for the agent.
- List the media server names to which the agent is publishing.
- Add a media server to the publishing list.
- Remove a media server from the publishing list.
- View the following status information:
 - Activity status
 - Current source
 - Current folder
 - Current file
 - Currently attached media server

See [“Using the Remote Agent Utility Command Line Applet”](#) on page 1958.

Using the Remote Agent Utility Command Line Applet

Use the following steps to use the Remote Agent Utility Command Line Applet.

See [“About the Remote Agent Utility Command Line Applet”](#) on page 1958.

To use the Remote Agent Utility Command Line Applet:

- 1 Open a command prompt.
- 2 From the Backup Exec installation directory, type `ramcmd.exe` followed by a series of command switches.

The default installation location is `c:\Program Files\Symantec\Backup Exec\RAWS`

See [“Remote Agent Utility Command Line Applet switches”](#) on page 1959.

Remote Agent Utility Command Line Applet switches

The following table describes the switches that you can use with the Remote Agent Utility Command Line Applet.

See [“About the Remote Agent Utility Command Line Applet”](#) on page 1958.

Table W-4 Remote Agent Utility Command Line Applet switches

Switch	Description
<code>status:[n]</code>	Status output is repeated every <n> seconds, with a range of 1 - 86400. Press Q to stop the output from running. <code>ramcmd /status:[n]</code> When you use the <code>/status</code> switch without a time value, the Remote Agent status appears in the command window and then the applet exits.

Table W-4 Remote Agent Utility Command Line Applet switches (*continued*)

Switch	Description
<p>/publish:[on off add remove interval][/ms:<media server>] [/t:<x>]</p>	<p>Use the following parameters with the /publish switch:</p> <ul style="list-style-type: none"> ■ No parameter specified- Displays the publishing status and then exits. ■ [on] - Turns publishing on. Lets the Remote Agent send information about itself, such as its version and IP address. ■ [off] - Turns publishing off. ■ [add], [remove] - Used with /ms. You can use this parameter to add or remove media servers from the Remote Agent's publish list. ■ [interval] - Used with /t. Specifies the time interval that the Remote Agent sends information about itself to the media server. <p>You can set the time interval in minutes using the /t:<x> parameter.</p> <p>Note: The [interval] switch must be used with the /t: switch. Using [interval] alone on the command line is not supported.</p> <pre>ramcmd /publish:[on off add remove interval] [/ms<media server>] [/t:<x>]</pre>

Table W-4 Remote Agent Utility Command Line Applet switches (*continued*)

Switch	Description
/oracle: [new edit delete]	Use the following parameters with the /oracle switch:
/in:[<instance name>]	<ul style="list-style-type: none"> ■ No parameter specified- Displays the existing Oracle instances and then exits.
/ms:[<media server address>]	<ul style="list-style-type: none"> ■ [new], [edit], [delete] - Used with switch /in.
/jt:[<job template>]	<ul style="list-style-type: none"> ■ /in:[<instance name>] - Used to add, edit, and delete Oracle instance names from the Oracle instance list.
/user:[<username>]	<ul style="list-style-type: none"> ■ /ms:[<media server name address>] - Sets the media server name or its IP address.
/password:[<password> *]	<ul style="list-style-type: none"> ■ /jt:[<job template>] - Sets a Backup Exec job template.
/rc: [yes no]	<ul style="list-style-type: none"> ■ /user:[<username>] - Sets a username.
/tns:[<TNS name>]	<ul style="list-style-type: none"> ■ /password:[<password> *] - Sets a password to be used with /user:[<username>]. If you omit the password, or you use an asterisk [*], you do not need to enter the password on the command line. After the command runs, a prompt appears asking you for a password. ■ /rc:[yes no] - Turns the Use recover catalog setting on or off. If /rc appears without a parameter, then the current status for that instance is displayed. ■ /tns:[TNS name] - Sets the TNS name alias of an available Oracle database and the server it resides on in the Oracle TNSNAMES file. <pre>ramcmd.exe /oracle:edit /in:<instance name> /rc: [yes no] [/tns:<TNS name>] [/user:<username>] [/password:password *]</pre>

Table W-4 Remote Agent Utility Command Line Applet switches (*continued*)

Switch	Description
/db2:[new edit delete] /in:[<instance name>] /ms:[<media server address>] /jt:[<job template>]/user:[<username>] /password:[<password> *] /al:<archive log template> /tns:[<TNS name>]	Use the following parameters with the /db2 switch: <ul style="list-style-type: none"> ■ No parameter specified- Displays the existing DB2 instances and then exits. ■ [new], [edit], [delete] - Used with switch /in. ■ /in:[<instance name>] - Used to add, edit, and delete DB2 instance names from the DB2 instance list. ■ /ms:[<media server name address>] - Sets the media server name or its IP address. ■ /jt:[<job template>] - Sets a Backup Exec job template. ■ /user:[<username>] - Sets a username. ■ /password:[<password> *] - Sets a password to be used with /user:[<username>]. If you omit the password, or you use an asterisk [*], you do not need to enter the password on the command line. After the command runs, a prompt appears asking you for a password. ■ /al:<archive log template>- Sets the archive log template name to <archive log template>. ■ /tns:[TNS name] - Sets the TNS name alias of an available Oracle database and the server it resides on in the Oracle TNSNAMES file. <pre> ramcmd .exe /db2:new /in:<instance name> /ms:<media server address> [/jt:<job template>] [/al:<archive log template>] /user:<username> [/password:<password> *] </pre>
/auth:[on off] [/user:<username>] [/password:<password> *]	Enables or disables media server authentication for Oracle and DB2 operations. <ul style="list-style-type: none"> ■ /auth:on - Turns the state on. Requires /user parameter. ■ /auth:off - Turns the state off. Requires /user parameter. ■ /user:<username> - Sets a username. ■ /password:<password> - Sets a password to be used with /user:<username>. If you enter an asterisk for the password or omit the password, you are prompted for the password.

Table W-4 Remote Agent Utility Command Line Applet switches (*continued*)

Switch	Description
/port:[<port>]	<p>Displays or sets a custom port that is used to connect to the media server during Oracle and DB2 operations.</p> <ul style="list-style-type: none">■ /port - Displays the current port number. If the port is the default port, displays "(default)".■ /port:<port> - Sets the port number to <port>. To change the port to the default port number, type [/port:0].
/log_path:[<log path>]	<p>Displays or sets a custom path for debug logs.</p> <ul style="list-style-type: none">■ /log_path - Displays the log directory path and then exits.■ /log_path:<"logs path"> - Creates the directory <"logs path">. If the path has a space in the name, enclose the path in quotes. For example, "C:\Program files\LogFolder".

See [“Using the Remote Agent Utility Command Line Applet”](#) on page 1958.

Symantec Backup Exec Remote Media Agent for Linux Servers

This appendix includes the following topics:

- [About the Remote Media Agent for Linux Servers](#)
- [How the Remote Media Agent for Linux Servers works](#)
- [Requirements for the Remote Media Agent for Linux Servers](#)
- [About installing the Remote Media Agent for Linux Servers](#)
- [Adding a Linux server as a Remote Media Agent](#)
- [Editing properties for the Remote Media Agent for Linux Servers](#)
- [Sharing a Remote Media Agent between multiple media servers](#)
- [About creating device pools for devices attached to the Remote Media Agent for Linux Servers](#)
- [Deleting a Remote Media Agent for Linux Servers from a media server](#)
- [Backing up data by using the Remote Media Agent for Linux Servers](#)
- [Restoring data by using the Remote Media Agent for Linux Servers](#)
- [About the Tape Library Simulator Utility](#)
- [Uninstalling the Remote Media Agent for Linux Servers](#)
- [Troubleshooting the Remote Media Agent for Linux Servers](#)

About the Remote Media Agent for Linux Servers

The Remote Media Agent for Linux Servers lets you back up data from remote computers to the following devices:

- The storage devices that are directly attached to a Linux server.
- A simulated tape library on a Linux server.

You can add a Linux server to a media server as a Remote Media Agent. Then, you can back up data from the Linux server or from supported remote computers to the devices that are attached to the Linux server. You can also create a virtual device on a server on which the Remote Media Agent for Linux Servers is installed. This virtual device emulates a SCSI tape library.

The Remote Media Agent supports operations for the following remote agents:

- Remote Agent for Windows Systems
- Remote Agent for NetWare Systems
- Remote Agent for Macintosh Systems
- Remote Agent for Oracle on Linux or Windows Systems
- Agent for DB2 on Windows Servers
- Agent for SAP Applications

See [“How the Remote Media Agent for Linux Servers works”](#) on page 1966.

See [“Requirements for the Remote Media Agent for Linux Servers”](#) on page 1967.

See [“About the Tape Library Simulator Utility”](#) on page 1979.

How the Remote Media Agent for Linux Servers works

From the Backup Exec media server, you can add a Linux server as a Remote Media Agent. The Remote Media Agent establishes a data connection to the remote computer on which a supported remote agent is installed. Then, you can create backup, restore, and utility jobs on the media server that run on the Linux server's storage devices.

If you use Backup Exec Central Admin Server Option or the SAN Shared Storage Option, you can share a Remote Media Agent between multiple media servers. Sharing can be enabled when you add a Remote Media Agent. You can select new media servers to share a Remote Media Agent or remove the sharing ability from the media servers at any time.

See [“About sharing storage”](#) on page 444.

Job performance increases because data travels from the remote computers to the devices that are attached to the Linux server. This increase is especially apparent if the media server is located at a different site than the Remote Media Agent and the remote computers.

The Remote Media Agent does not have a user interface. You use the administration console on the media server to manage the jobs and devices on the Remote Media Agent. The Backup Exec media server maintains job logs, catalogs, job histories, alerts, and notifications.

See “[Requirements for the Remote Media Agent for Linux Servers](#)” on page 1967.

See “[About installing the Remote Media Agent for Linux Servers](#)” on page 1968.

See “[Adding a Linux server as a Remote Media Agent](#)” on page 1972.

See “[About the Tape Library Simulator Utility](#)” on page 1979.

Requirements for the Remote Media Agent for Linux Servers

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

You must have superuser privileges on the Linux servers before you install the Remote Media Agent for Linux Servers.

Note: The Remote Media Agent does not support the Backup Exec File System Archiving Option or the Exchange Mailbox Archiving Option.

Symantec recommends that you use the Secure Shell (SSH) protocol when you push-install the Remote Media Agent to remote servers. You must enable SSH before you install the Remote Media Agent.

Note: Some versions of Linux may require that you install the `libstdc++.so.5` package.

See “[Troubleshooting the Remote Media Agent for Linux Servers](#)” on page 1987.

See “[About installing the Remote Media Agent for Linux Servers](#)” on page 1968.

About installing the Remote Media Agent for Linux Servers

Use the Backup Exec installation media to do the following:

- Install the Remote Media Agent for Linux Servers on a local Linux server.
- Push-install the Remote Media Agent for Linux Servers to one or more remote Linux servers.

If you push-install the Remote Media Agent for Linux Servers, the RSH (Remote Shell) is used by default. Symantec recommends that you use SSH (Secure Shell) instead. To use SSH, you must enable it before you install the Remote Media Agent for Linux Servers. Refer to your operating system documentation for more information about SSH.

When you install the Remote Media Agent for Linux Servers, Backup Exec creates the beoper group and adds root as a member. Any Linux user that you add to the beoper group gets the permissions necessary to back up and restore the Linux servers.

However, if Backup Exec detects an NIS server during the Remote Media Agent for Linux Servers installation, then the beoper group is not created. You must create the beoper group manually on the Linux servers.

After the installation completes, you must add the Linux server as a Remote Media Agent on the media server. Then, you can send jobs to the devices that are attached to the Linux server.

See [“Adding a Linux server as a Remote Media Agent ”](#) on page 1972.

See [“Installing the Remote Media Agent for Linux Servers”](#) on page 1968.

See [“About the Backup Exec operators group for the Remote Media Agent for Linux Servers”](#) on page 1971.

Installing the Remote Media Agent for Linux Servers

You can install the Remote Media Agent on a local Linux server or push-install it to one or more remote Linux servers.

See [“About installing the Remote Media Agent for Linux Servers”](#) on page 1968.

Note: You must unzip the RALUS_RMALS_RAMMS_<version number>.gz file on a Linux or UNIX server. The installation does not run if it is unzipped on a computer that runs the Windows operating system.

To install the Remote Media Agent for Linux Servers

- 1 At a Linux server, place the Backup Exec installation media in the appropriate drive.
- 2 Log on as root on the server on which you want to install the Remote Media Agent for Linux Servers.

- 3 Navigate to the following path on the installation media:

<LinuxUnixMac>

- 4 Copy the RALUS_RMALS_RAMs_<version number>.gz file in this directory to a directory on the local server.

- 5 Unzip the file.

For example:

```
gunzip RALUS_RMALS_RAMs_<version number>.gz
```

- 6 Untar the file.

For example:

```
tar -xf RALUS_RMALS_RAMs_<version number>.tar
```

- 7 Start the **installrml** script.

For example:

```
./installrml
```

- 8 Do one of the following:

To install on a local server

Press **Enter**.

To install to one remote server

Type the name, IP address, or fully qualified domain name of a Linux server.

To install to multiple remote servers

Type the names, IP addresses, or fully qualified domain names of the Linux servers. Leave a space between each identifier.

- 9 After the installer checks for a valid Linux operating system during the initial system check, press **Enter**.
- 10 Review the package installation summary, and then press **Enter**.
- 11 After the system installation requirements check completes, press **Enter**.
- 12 Start the prerequisites check by pressing **Enter**.

- 13 Type the name, IP address, or fully qualified domain name of the media server (directory host) that you want to use this Remote Media Agent.
- 14 Type any additional names, IP addresses, or fully qualified domain names of media servers that you want to use this Remote Media Agent.
- 15 Do one of the following:

If the server name, IP address, or fully qualified domain name is correct Press **Enter** to continue the installation.

If you want to change a server name, IP address, or fully qualified domain name Type **N**, press **Enter**, and then change the information.

- 16 Start the NIS server scan by pressing **Enter**.
- 17 Examine the results of the NIS server scan, and then do one of the following:

If an NIS server is detected The Remote Media Agent installer cannot create the beoper group. You must create it manually after the Remote Media Agent installation is complete.
Continue with the next step.

If an NIS server is not detected Use the installer to create the beoper group.
Do the following in the order listed:

- To let the installer create the beoper group, type **y**.
- To select the next available Group ID, type **n**.
- To add the root user account to the beoper group, type **y**.
- Continue with the next step.

- 18 Start the installation by pressing **Enter**.
- 19 After the installation completes, press **Enter** to start the configuration process.
- 20 After the configuration process completes, press **Enter** to save the installation log to the following file:

/var/tmp/vxif/installrma/summary file number/installrma.log

- 21 If the Remote Media Agent installer did not create a beoper group, you must create it.
See [“Creating the Backup Exec operators group manually”](#) on page 1880.
- 22 Start the Remote Agent for Linux or UNIX Servers daemon.
See [“Starting the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.
- 23 Add the Linux server as a Remote Media Agent.
See [“Adding a Linux server as a Remote Media Agent ”](#) on page 1972.

About the Backup Exec operators group for the Remote Media Agent for Linux Servers

The Backup Exec operators (**beoper**) group contains the names of the users who have permission to back up and restore the Linux servers.

When you install the Remote Media Agent for Linux Servers, Backup Exec creates the **beoper** group and adds root as a member. Any Linux user that you add to the beoper group gets the necessary permission to back up and restore the Linux servers.

However, if an NIS server is detected during the Remote Media Agent installation, Backup Exec cannot create the **beoper** group. You must create the **beoper** group manually on the Linux servers on which you want to install the Remote Media Agent. You must create the **beoper** group before you start backup and restore operations. Otherwise, connections fail between the Linux servers and the media server.

Before the members of the **beoper** group can perform backup or restore operations, they must have a Backup Exec logon account.

See [“Creating the Backup Exec operators group manually for the Remote Media Agent for Linux Servers”](#) on page 1971.

See [“Creating a Backup Exec logon account”](#) on page 187.

Creating the Backup Exec operators group manually for the Remote Media Agent for Linux Servers

You must create a beoper group on each Linux server on which you want to install the Remote Media Agent for Linux Servers.

See [“About the Backup Exec operators group for the Remote Media Agent for Linux Servers”](#) on page 1971.

Note: Ensure that you understand how to set security for groups on Linux servers before you assign a Group ID for the beoper group.

Table X-1 How to manually create the beoper group

Step	Action	More Information
Step 1	Navigate to the Linux server on which you want to install the Remote Media Agent. If the Linux server is in a NIS domain, navigate to the NIS domain's group file.	See the NIS documentation for information on how to add a group to a NIS domain group file.
Step 2	Create a group with the following case-sensitive name: beoper	See the operating system documentation for more information about how to create a group.
Step 3	In the beoper group, add the users that you want to have permission to back up and restore the Linux server.	See the operating system documentation for more information about how to add users to a group
Step 4	Create a Backup Exec logon account for each user that you add to the beoper group.	See “Creating a Backup Exec logon account” on page 187.

Adding a Linux server as a Remote Media Agent

When you add the Linux server as a Remote Media Agent, you can select the media servers that can access the devices that are attached to the Linux server.

See [“About sharing storage”](#) on page 444.

To add a Linux server as a Remote Media Agent

- 1 On the navigation bar, click **Devices**.
- 2 In the task pane, under **Device Tasks**, click **Configure devices assistant**.
- 3 On the Configure Devices Assistant dialog box, under **Remote Media Agent Option**, click **Remote Media Agent Storage**.
- 4 Complete the options on the **General** tab.

See [“Add Remote Media Agent options”](#) on page 1973.

- 5 On the **Sharing** tab, select the media servers that you want to use with this Remote Media Agent
- 6 Click **OK** to add the Remote Media Agent.
- 7 On the media server, restart the Backup Exec services so that the Remote Media Agent and the storage devices that are directly attached to it appear in the **Devices** view.
See [“Starting and stopping Backup Exec services”](#) on page 168.

Add Remote Media Agent options

You must provide information when you add a Linux server as a Remote Media Agent to a media server.

See [“Adding a Linux server as a Remote Media Agent”](#) on page 1972.

Table X-2 Add Remote Media Agent options

Item	Description
Server	Specifies the name of the Linux server that you want to add as a Remote Media Agent. If the Backup Exec SAN Shared Storage Option is installed in your environment, use the host name or fully qualified domain name of the Linux server. That is, use the name of the Linux computer that appears when you browse for backup selections. If you use the IP address, Backup Exec cannot distinguish which device path to use for jobs.

Table X-2 Add Remote Media Agent options (*continued*)

Item	Description
Port	<p>Lists the port to use for communications between the media server and the Remote Media Agent. If you change the port number, you must edit the services file in the /etc directory on the Linux server, and update the NDMP entry.</p> <p>See “Changing the port for communications between the media server and the Remote Media Agent” on page 1975.</p> <p>Ensure that this port is open in any firewalls that exist between the Remote Media Agent and the media server. Use a port number that is not in use by another application or service.</p> <p>The default port is 10000.</p>
Description	Displays a description that you choose.
Logon Account	<p>Indicates the logon account for the Remote Media Agent.</p> <p>The default logon account is the system logon account for the media server.</p>
Enable ICMP ping operations for Backup Exec to detect the server	<p>Lets the media server use ICMP ping operations to locate the Linux server. You can turn off this option in environments where ping requests are blocked.</p> <p>This option is enabled by default.</p>
Backup Exec logon account	<p>Indicates the Backup Exec logon account that you want to use to log on this server.</p> <p>See “About configuring logon accounts” on page 184.</p>

See [“About creating device pools for devices attached to the Remote Media Agent for Linux Servers ”](#) on page 1977.

See [“Backing up data by using the Remote Media Agent for Linux Servers”](#) on page 1978.

Changing the port for communications between the media server and the Remote Media Agent

You can change the port that Backup Exec uses to communicate with the Remote Media Agent.

To change the port for communications between the media server and the Remote Media Agent

- 1 On the computer on which the Remote Media Agent is installed, use a text editor to open the services file in the `/etc` directory.

For example:

```
vi /etc/services
```

- 2 Search the file for an entry that is similar to the following:

ndmp 10000/tcp

- 3 Do one of the following:

If this entry exists

Change the port number to the port number that you want to use.

If this entry does not exist

Do the following in the order listed:

- At the end of the file, type `ndmp`, and then press **Tab**.
- Type the port number that you want NDMP to use, and then type `/tcp`.
- Press **Enter**.

- 4 Save the file, and then exit the editor.
- 5 Restart the Remote Agent for Linux or UNIX Servers daemon.

See [“Starting the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.

Editing properties for the Remote Media Agent for Linux Servers

You can edit the properties for a Remote Media Agent for Linux Servers.

To edit properties for the Remote Media Agent for Linux Servers

- 1 On the navigation bar, click **Devices**.
- 2 Select a Remote Media Agent.
- 3 In the task pane, under **General Tasks**, click **Properties**.
See “[Remote Media Agent properties](#)” on page 1976.
- 4 Click **OK**.

Remote Media Agent properties

You can view properties for the Remote Media Agent for Linux Servers.

See “[Editing properties for the Remote Media Agent for Linux Servers](#)” on page 1975.

The following table lists the Remote Media Agent properties:

Table X-3 Remote Media Agent properties

Item	Description
Name	Displays the name, IP address, or fully qualified domain name of the Remote Media Agent.
Port	Displays the port that is used for communications between the media server and the Remote Media Agent.
Media server status	Displays the status of the media server. Media server status includes Online, Pause, Unavailable, and Offline.
Description	Displays a description of the Remote Media Agent. You can edit this description.
Enable ICMP ping operations for Backup Exec to detect the Remote Media Agent	Lets Backup Exec communicate with the Remote Media Agent. You can turn off this option in environments where ping requests are blocked. This option is enabled by default.
Host ID	Displays the identifier number that the Remote Media Agent generates.
System version	Displays the version of the operating system that runs on the Remote Media Agent.

Table X-3 Remote Media Agent properties (*continued*)

Item	Description
Logon account	Indicates the logon account for the Remote Media Agent. Click Change to select or create another logon account.

Sharing a Remote Media Agent between multiple media servers

If the Central Admin Server Option or the SAN Shared Storage Option is installed, you can select media servers to share a Remote Media Agent. When you add a Remote Media Agent, the media server that you used to add the device is automatically selected for sharing.

See “[About sharing storage](#)” on page 444.

To share a Remote Media Agent between multiple media servers

- 1 On the navigation bar, click **Devices**.
- 2 In the **Devices** view, right-click the Remote Media Agent that you want media servers to access.
- 3 Select **Manage sharing**.
- 4 Select the Remote Media Agent that you want to share.
- 5 Under **Media Servers**, select the media servers that you want to use with the Remote Media Agent.
- 6 Click **OK**.
- 7 Restart the Backup Exec services on the media servers that you selected in step 5.

About creating device pools for devices attached to the Remote Media Agent for Linux Servers

Backup Exec does not include the devices that are attached to a Remote Media Agent in the **All Devices** device pool. You cannot add these devices to the **All Devices** device pool.

Remote Media Agents may reside in different physical locations. To reduce network traffic and increase job performance, create separate device pools for Remote Media Agents that are located at different sites.

See [“Creating device pools”](#) on page 516.

See [“About the Tape Library Simulator Utility”](#) on page 1979.

Deleting a Remote Media Agent for Linux Servers from a media server

You can delete a Remote Media Agent from a media server.

To delete a Remote Media Agent for Linux Servers from a media server

- 1 On the navigation bar, click **Devices**.
- 2 Select the Remote Media Agent that you want to delete.
- 3 In the task pane, under **General Tasks**, click **Delete**.
- 4 Restart the Backup Exec services at a convenient time.

See [“Starting and stopping Backup Exec services”](#) on page 168.

Backing up data by using the Remote Media Agent for Linux Servers

Create a backup job for the Remote Media Agent from the media server.

To back up data by using the Remote Media Agent for Linux Servers

- 1 On the navigation bar, click the arrow next to **Backup**.
- 2 Click **New Backup Job**.
- 3 On the backup selections list, select the data that you want to back up.
See [“Creating selection lists”](#) on page 296.
- 4 In the task pane, under **Destination**, click **Device and Media**.
- 5 Select the Remote Media Agent to which you want to send the backup.
- 6 Complete the remaining backup job properties as necessary.

See [“Creating a backup job by setting job properties”](#) on page 332.

Restoring data by using the Remote Media Agent for Linux Servers

Create a restore job for the Remote Media Agent from the Backup Exec media server.

Note: Use devices that are attached to the Backup Exec media server to restore data from the tapes that other applications created. The Remote Media Agent supports only Microsoft Tape Format (MTF) media.

To restore data by using the Remote Media Agent for Linux Servers

- 1 On the navigation bar, click the arrow next to **Restore**.
- 2 Click **New Restore Job**.
- 3 In the restore selections list, on the **View by Resource** tab, select the appropriate data to restore.
See [“Selections options for restore jobs”](#) on page 609.
- 4 In the task pane, under **Source**, click **Device**.
- 5 Select a device pool that contains Remote Media Agent devices.
- 6 Complete other restore job properties as necessary.
See [“Restoring data by setting job properties”](#) on page 605.

About the Tape Library Simulator Utility

The Tape Library Simulator Utility lets you create a virtual device on a hard disk or on any mounted volume on a Linux server. This virtual device emulates a SCSI tape library. The Remote Media Agent for Linux Servers must be installed on the server.

When you run the Tape Library Simulator Utility, you are prompted for the following information:

- The number of slots that you want to allocate to this library.
- The location or path for the library.

The Tape Library Simulator Utility then creates the media for the simulated tape library. To ensure that each media has a unique name, the Tape Library Simulator Utility creates a bar code label for each media. You cannot rename these bar code labels. However, you can add a unique media description.

See [“General properties for media”](#) on page 261.

The simulated tape library emulates an Advanced Intelligent Tape (AIT) media type. This media type is seldom used, so it helps you distinguish between a physical robotic library and a simulated tape library. The simulated media also has an AIT media type label.

The format of the files that are written to the simulated tape library is similar to the file format of backup-to-disk files. However, you cannot copy or move files between simulated tape libraries and backup-to-disk folders.

Backup Exec does not include simulated tape libraries in the **All Devices** device pool. You cannot add a simulated tape library to the **All Devices** device pool. You can add the simulated tape library to another device pool.

To use the Tape Library Simulator Utility, you must have a minimum of 500 MB of available space on the Linux server. The available space includes hard disk space, flash drives, and USB drives. If there is not enough space, the jobs fail with an end-of-media error. You must either create available disk space or you must direct the jobs to another volume, and then start the jobs again.

A simulated tape library does not support all of the tasks that are available for physical robotic libraries.

See [“Utility jobs for virtual tape libraries and simulated tape libraries”](#) on page 482.

See [“Creating a simulated tape library”](#) on page 1980.

Creating a simulated tape library

Create a simulated tape library on a server on which the Remote Media Agent for Linux Servers is installed. You must create the simulated tape library on a hard disk or on a mounted volume.

See [“About the Tape Library Simulator Utility”](#) on page 1979.

To create a simulated tape library

- 1 At the Remote Media Agent, stop the Remote Agent for Linux or UNIX Servers daemon.

See [“Stopping the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.

- 2 Navigate to the following path that contains the Tape Library Simulator Utility:

```
</opt/VRTSralus/bin>
```

For example:

```
cd /opt/VRTSralus/bin
```

- 3 Start the **mktls** utility.

For example:

```
./mktls
```

- 4 Select **Create a new simulated tape library**, and then press **Enter**.

- 5 Enter the appropriate information.
See [“Simulated Tape Library options”](#) on page 1981.
- 6 Exit the utility.
- 7 Restart the Remote Agent for Linux or UNIX Servers daemon.
See [“Starting the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.
- 8 On the media server, restart the Backup Exec services.
See [“Starting and stopping Backup Exec services”](#) on page 168.

Simulated Tape Library options

When you create a simulated tape library, you must provide a directory path and the number of slots for the library.

See [“Creating a simulated tape library”](#) on page 1980.

Table X-4 Simulated Tape Library options

Item	Description
Directory Path	Type the path of the directory for the simulated tape library. You can enter up to 512 characters. If the path does not exist, the Tape Library Simulator Utility creates it.
Number of Slots	Select the number of slots for this simulated tape library. The number of slots can range from 1 to 50. The default number of slots is 20.

See [“Viewing simulated tape libraries properties”](#) on page 1981.

Viewing simulated tape libraries properties

You can use the Symantec Tape Library Simulator Utility to view information about a simulated tape library and its contents.

To view simulated tape library properties

- 1 On the Remote Media Agent, stop the Remote Agent for Linux or UNIX Servers daemon.

See [“Stopping the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.

- 2 Navigate to the following directory that contains the Tape Library Simulator Utility:

```
/opt/VRTSralus/bin
```

For example:

```
cd /opt/VRTSralus/bin
```

- 3 Start the **mktls** utility.

For example:

```
./mktls
```

- 4 Select **View an existing simulated tape library**.

- 5 Move your cursor to the simulated tape library that you want to view, and then press **Enter**.

- 6 Press **Enter** again to view the simulated tape library properties.

See [“Simulated tape library properties”](#) on page 1982.

- 7 Type **Q** to exit the utility.

- 8 Restart the Remote Agent for Linux or UNIX Servers daemon.

See [“Starting the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.

Simulated tape library properties

You can view the properties of a simulated tape library.

See [“Viewing simulated tape libraries properties”](#) on page 1981.

Table X-5 Simulated tape library properties

Item	Description
Number of drives	Displays the number of drives for this simulated tape library. A simulated tape library can have only drive. This drive is not configurable.

Table X-5 Simulated tape library properties (*continued*)

Item	Description
Number of slots	Displays the number of slots for this simulated tape library. The number of slots can range from 1 to 50. The default number of slots is 20.
Tape capacity	Displays the tape capacity. The default capacity is 100 gigabytes.
Directory path	Displays the directory path where the simulated tape library exists.

Deleting a simulated tape library

You can use the Tape Library Simulator Utility to delete a simulated tape library. You must then manually delete the content of the simulated tape library files, and then delete the directories that contain these files.

To delete a simulated tape library

- 1 At the Remote Media Agent, stop the Remote Agent for Linux or UNIX Servers daemon.

See [“Stopping the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.

- 2 Navigate to the following directory that contains the Tape Library Simulator:

```
/opt/VRTSralus/bin/
```

For example:

```
cd /opt/VRTSralus/bin/
```

- 3 Start the **mktls** utility:

For example:

```
./mktls
```

- 4 Select **View an existing simulated tape library**.
- 5 Select the simulated tape library that you want to delete.
- 6 When you are prompted, delete the simulated tape library.
- 7 Exit the utility.

- 8 Restart the Remote Agent for Linux or UNIX Servers daemon.
See [“Starting the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.
- 9 Find the simulated tape library files, and then manually delete them.
See [“About the Tape Library Simulator Utility”](#) on page 1979.
- 10 On the media server, restart the Backup Exec services when it is convenient.
See [“Starting and stopping Backup Exec services”](#) on page 168.

Managing simulated tape libraries from the command line

You can use the command line to create a simulated tape library. Create a simulated tape library on a hard disk or on any mounted volume on the Remote Media Agent. From the command line, you can also view and delete simulated tape libraries.

To manage simulated tape libraries from the command line

- 1 At the Remote Media Agent server, stop the Remote Agent for Linux or UNIX Servers daemon.
See [“Stopping the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.
- 2 Navigate to the following directory that contains the Tape Library Simulator Utility:

```
/opt/VRTSralus/bin
```

For example:

```
cd /opt/VRTSralus/bin
```
- 3 Start the **mktls** utility with the appropriate parameter switches.
See [“Command line switches for the Tape Library Simulator Utility”](#) on page 1984.
- 4 Start the Remote Agent for Linux or UNIX Servers daemon.
See [“Starting the Remote Agent for Linux or UNIX Servers daemon”](#) on page 1907.

Command line switches for the Tape Library Simulator Utility

You can use command line switches to manage simulated tape libraries. For example, the following command line creates a simulated tape library with 10 slots that is located at /TLS2/Testing.

```
./mktls -s10 -p/TLS2/Testing
```

See [“Managing simulated tape libraries from the command line”](#) on page 1984.

Table X-6 Command line switches for the Tape Library Simulator Utility

Switch	Description
-p<path>	Specifies the path to the directory for the simulated tape library. If the path does not exist, the utility creates it. The maximum path size is 512 characters.
-s<number of slots>	Specifies the number of slots for this simulated tape library. The number of slots can range from one to 50. The default number is 20.
-r	Prevents the information from displaying.
-l	Lists the simulated tape libraries that exist for the Remote Media Agent.
-d -p<path>	Specifies the path of the simulated tape library that you want to delete.
-h	Displays the online Help.

Uninstalling the Remote Media Agent for Linux Servers

Before you uninstall the Remote Media Agent for Linux Servers, you should note the location of the simulated tape library files. Then, you can delete all of the simulated tape library files after the uninstall operation completes. When you delete these files, you delete the backup data that you stored on the Linux server.

See [“Finding simulated tape library files”](#) on page 1986.

Note: You must have the Backup Exec installation media to uninstall the Remote Media Agent for Linux Servers.

To uninstall the Remote Media Agent for Linux Servers

- 1 On the Linux server, place the Backup Exec installation media in the appropriate device.
- 2 Log on as root to the server from which you want to uninstall the Remote Media Agent for Linux Servers.

- 3 Navigate to the following path on the installation media:
<LinuxUnixMac>
- 4 Start the **uninstallrmal** script.
For example:

```
./uninstallrmal
```
- 5 Do one of the following:

To uninstall the Remote Media Agent from one server	Type the name, IP address, or fully qualified domain name of the Linux server.
To uninstall the Remote Media Agent from multiple servers	Type the names, IP addresses, or the fully qualified domain names of the Linux servers. Leave a space between each identifier.
- 6 Press **Enter**.
- 7 After the Remote Media Agent package check completes successfully, press **Enter**
- 8 When you are prompted to uninstall the Remote Media Agent packages, press **Enter** to save the uninstall summary and log to the following location:

```
/var/tmp/vxif/uninstallrmalsummary file number.log
```
- 9 Manually delete the simulated tape library files.

Finding simulated tape library files

Before you uninstall the Remote Media Agent for Linux Servers, you should note the location of the simulated tape library files. Then, after you uninstall the Remote Agent, you can delete all of the simulated tape library files. When you delete these files, you delete the backup data that you stored on the Linux server.

See [“Uninstalling the Remote Media Agent for Linux Servers”](#) on page 1985.

See [“About the Tape Library Simulator Utility”](#) on page 1979.

To find simulated tape library files

- 1 Log on as root to the server on which you want to find the simulated tape library files.
- 2 Navigate to the following directory that contains the Tape Library Simulator:
`/opt/VRTSralus/bin`
For example:

```
cd /opt/VRTSRALus/bin
```
- 3 Start the **mktls** utility to list the simulated tape library files and folders.
For example:

```
/opt/VRTSralus/bin/mktls -l
```
- 4 Write down the locations of the directories for the simulated tape library files.

Troubleshooting the Remote Media Agent for Linux Servers

If there are issues with the Remote Media Agent, review the following questions and answers.

Table X-7 Troubleshooting the Remote Media Agent for Linux Servers

Question	Answer
The Remote Media Agent does not detect my attached device. What should I do?	<p>First, ensure that Backup Exec and the Remote Media Agent for Linux Servers support the device.</p> <p>You can find a list of compatible devices at the following URL:</p> <p>http://entsupport.symantec.com/umi/V-269-2</p> <p>If the device is listed on the hardware compatibility list, ensure the following:</p> <ul style="list-style-type: none">■ The operating system detects the device■ The device is listed in <code>/proc/scsi/scsi</code> <p>If the operating system can detect the device, ensure that the device is listed in <code>/etc/VRTSralus/TILDBG.TXT</code>.</p>

Table X-7 Troubleshooting the Remote Media Agent for Linux Servers
(continued)

Question	Answer
<p>My Backup Exec media server does not display the devices that are attached to my Remote Media Agent. What should I do?</p>	<p>Try the following procedures:</p> <ul style="list-style-type: none"> ■ Ensure that the Remote Agent for Linux or UNIX Servers daemon is running. If it is not running, start the daemon, and check that power for the server is enabled, and that all cables are properly attached. ■ Ensure that the Remote Media Agent properties are set to the correct port, and that ICMP ping operations are enabled. ■ Ensure that the Backup Exec services are restarted after a Remote Media Agent is added to the media server. The available devices should be displayed under the Remote Media Agent node. <p>See “Editing properties for the Remote Media Agent for Linux Servers” on page 1975.</p> <p>See “Starting the Remote Agent for Linux or UNIX Servers daemon” on page 1907.</p>
<p>Why don't my remote devices appear in the All Devices device pools?</p>	<p>By default, Backup Exec does not include remote devices in the All Devices device pool. Symantec recommends that you create a separate device pool for the devices that are attached to each Remote Media Agent.</p> <p>See “About creating device pools for devices attached to the Remote Media Agent for Linux Servers” on page 1977.</p>

Table X-7 Troubleshooting the Remote Media Agent for Linux Servers
(continued)

Question	Answer
<p>The Remote Media Agent won't run on the remote computer. What should I do?</p>	<p>Ensure that the Remote Media Agent is installed on a supported version of Linux.</p> <p>You can find a list of compatible operating systems, platforms, and applications at the following URL: http://entsupport.symantec.com/umi/V-269-1</p> <p>If you install the Remote Media Agent to an unsupported version of Linux, the Remote Media Agent is unavailable for use. You cannot create the jobs that run on the devices that are attached to the Linux server. However, you can back up the Linux server by using the Remote Agent for Linux or UNIX Servers component. This component is installed with the Remote Media Agent.</p> <p>To use the Remote Agent for Linux or UNIX Servers component to back up the Linux server, do the following:</p> <ul style="list-style-type: none"> ■ Edit the <code>ralus.cfg</code> file. ■ In the string <code>Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RMAL\DisableRMAL=0</code>, change 0 to 1. See “Editing configuration options for Linux, UNIX, and Macintosh computers” on page 1884. See “Running the begather utility to troubleshoot Backup Exec components on Linux servers” on page 809.
<p>I cannot load the Remote Media Agent. When I attempt to load the Remote Media Agent in console mode, <code>/beremote --log-console</code> shows the following message:</p> <p>Error while loading shared libraries: libstdc++.so.5: cannot open shared object file: No such file or directory.</p> <p>What should I do?</p>	<p>This error indicates that the libstdc++.so.5 library is not in the <code>/usr/lib</code> directory. This library is necessary to let the Remote Media Agent start and function. To resolve this issue, install the libstdc++.so.5 package. You can install this package from the media on which your copy of Linux was provided. Or, you can run the following command from a computer that has Internet access:</p> <pre>apt-get install libstdc++5</pre> <p>For SUSE Linux Enterprise Server 11, run the following command:</p> <pre>zypper install libstdc++5</pre>

Symantec Backup Exec SAN Shared Storage Option

This appendix includes the following topics:

- [About the SAN Shared Storage Option](#)
- [Requirements for the SAN Shared Storage Option](#)
- [About installing the SAN Shared Storage Option](#)
- [About devices in the SAN Shared Storage Option](#)
- [About designating a new primary database server and setting up servers in the SAN Shared Storage Option](#)
- [Troubleshooting failed components in the SAN Shared Storage Option](#)
- [Best practices for the SAN Shared Storage Option](#)

About the SAN Shared Storage Option

The Symantec Backup Exec SAN Shared Storage Option (SSO) enables multiple media servers to share secondary storage devices, such as robotic libraries, in a SAN. The secondary storage devices are not directly connected to a single server by SCSI, but are connected to a Fibre Channel Switched Fabric (FC-SW) or iSCSI.

To allow for sharing of storage devices and media between multiple media servers, a shared Advanced Device and Media Management (ADAMM) database resides on one media server called the primary database server or primary server. All media servers on the SAN connect to this database to obtain a single, unified view of all shared devices and media. Backup Exec uses this shared database to arbitrate all device and media requests with comprehensive overwrite protection policies to prevent accidental media overwrites. Multiple media servers can share devices

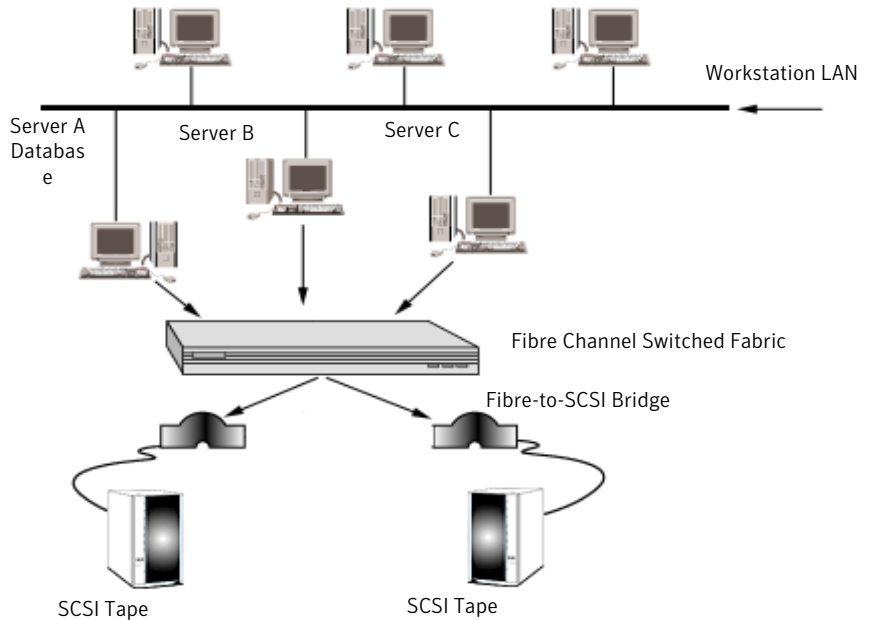
when sharing is enabled for those devices. To resolve potential access conflicts between multiple media servers, Backup Exec "reserves" robotic arms and tape devices while operations are being performed.

See "About sharing storage" on page 444.

Media catalogs are shared between media servers, so restore operations can be performed using any media servers that share catalogs. If the media must be moved from one device to another on the SAN, the media does not have to be cataloged again. In a shared storage environment, secondary storage devices also can be attached to the local SCSI, SATA, SAS, USB, or ATAPI buses of any media server. However, these local devices, disk or tape drives connected to a server, are only available to the server to which they are attached. Storage devices used with the SAN Shared Storage Option must have a vendor-assigned serial number.

In the following example of a shared storage environment, the primary database server and the media servers transmit data over the FC-SW through a fibre to SCSI bridge (router) to secondary storage devices (SCSI robotic libraries).

Figure Y-1 Example of the SAN Shared Storage Option using FC-SW



You can use any media server to change the names of the robotic libraries and drives to names that are more descriptive of your operations.

If you have multiple SANs, it is recommended that you treat each SAN independently, with each SAN having its own Backup Exec database server for the shared ADAMM database and catalogs for that loop. Using a single Backup Exec database server for more than one SAN increases the number of single-point failures that can affect the system.

Job completion statistics or errors can be viewed through any administration console attached to the server that ran the job.

See “[About sharing media in the SAN Shared Storage Option](#)” on page 1997.

See “[About scheduling and viewing jobs in the SAN Shared Storage Option](#)” on page 1997.

See “[Requirements for the SAN Shared Storage Option](#)” on page 1993.

Requirements for the SAN Shared Storage Option

The following are the minimum system requirements for running this release of the SAN Shared Storage Option:

- Windows 2003/2008 must be installed. You cannot use SAN SSO with the Windows Server Core installation option of Windows Server 2008.
- Physical Memory Available, as shown in the Windows Task Manager, plus the File Cache should exceed 256 MB.
- The SAN Shared Storage Option must be locally installed at each server that will be sharing secondary storage devices.
- The devices in your SAN must be on the supported device list. You can find a list of compatible devices at the following URL: <http://entsupport.symantec.com/umi/V-269-2>
- All hardware drivers must be up to date and started. You can find a list of compatible drivers at the following URL: <http://entsupport.symantec.com/rd/bews-drivers.htm>
- The primary server must have enough space for the catalogs of all the servers in the SAN.

For a fibre channel-connected installation, note the following additional requirements:

- A fibre channel host adapter and its device drivers must be installed and connected to the SAN.
- A hub or switch must be connected to all the fibre to SCSI bridges or to the fibre libraries on the SAN.

- All the robotic libraries must be connected to the SCSI bridges or to the fibre switches/hubs.
- The hub or switch must be powered up before the bridges or fibre channel libraries.
- All robotic libraries must be powered up before the bridges.
- The bridges must be powered up before Windows loads the fibre channel driver (usually during the boot phase).

Note: If the SAN Shared Storage Option is installed on a media server, Backup Exec disables all fibre channel-connected devices in Removable Storage. You cannot re-enable the devices in Removable Storage until the SAN Shared Storage Option and Symantec Device Drivers are uninstalled.

See [“About installing the SAN Shared Storage Option”](#) on page 1994.

About installing the SAN Shared Storage Option

You must install the Backup Exec SAN Shared Storage Option on the server that you want to designate as the primary server. You can then install the Backup Exec Shared Storage Option on other servers. The server containing the shared ADAMM database must be running before other media servers can be installed properly.

For best performance, install the shared ADAMM and catalog databases on the fastest server on the SAN that is not heavily loaded with non-Backup Exec tasks.

You can install the SAN Shared Storage Option while installing Backup Exec.

See [“Installing a custom installation of Backup Exec”](#) on page 119.

If you have already installed Backup Exec, you can install additional options.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

If installing the first server in the SAN select primary as your Shared Storage option type. You can install the Backup Exec device drivers as part of the Backup Exec installation.

If the SAN Shared Storage Option has already been installed to the primary server, and you are installing to a secondary server, select Secondary and enter the name for the Backup Exec database server.

If you installed Backup Exec on the primary database server, restart the database server. You must ensure that Backup Exec discovers all your devices before loading

Backup Exec on other servers. The server containing the shared ADAMM database must be running before other media servers can be installed properly.

The first time you run Backup Exec, use the Getting Started item in the Home view to configure environment settings. The Getting Started item provides a series of steps that you can follow to configure logon accounts, devices, and media sets. You should complete the applicable sections of this wizard as it guides you in preparing Backup Exec for operations. Setting the overwrite protection level is important since the media is shared throughout the SAN.

Note: When uninstalling Backup Exec, you must uninstall from the secondary servers before uninstalling from the primary server.

See [“Uninstalling Backup Exec”](#) on page 168.

See [“About scheduling and viewing jobs in the SAN Shared Storage Option”](#) on page 1997.

See [“Best practices for the SAN Shared Storage Option”](#) on page 2014.

See [“About enabling a SQL instance behind a firewall”](#) on page 414.

About devices in the SAN Shared Storage Option

At startup, Backup Exec recognizes all local storage devices as well as the storage devices on the SAN. If you do not see one or more of your attached storage devices or if the shared storage devices do not appear when Devices is selected from the navigation bar, use the **Tape Device Configuration Wizard**. This wizard guides you through installing the appropriate drivers for the storage hardware connected to your system.

See [“About configuring tape devices by using the Tape Device Configuration Wizard”](#) on page 453.

You can also view a list of the devices Backup Exec recognizes in the ADAMM.log file. The default path for the ADAMM.log file is \Program Files\Symantec\Backup Exec\Logs.

Storage devices are categorized as either robotic libraries or stand-alone drives. The Library Expansion Option is required to support robotic libraries with multiple tape drives.

See [“About the Library Expansion Option”](#) on page 453.

See [“Viewing license information”](#) on page 174.

If you need to add a new device to the SAN after you install Backup Exec, follow your storage network vendor’s instructions. After you add the new device, reboot

the primary server that contains the ADAMM database to verify that the new device is recognized. The new device may appear offline in the **Devices view** until the device discovery process completes. You should reboot the other media servers according to your storage network vendor's instructions. Some vendors do not support booting multiple servers concurrently or booting any server while active jobs are running.

Backup Exec's device management feature provides the following functionality for the secondary storage units on a SAN:

- **Device allocation.** Jobs must first reserve the shared secondary backup devices before they can be used. The job that gains a reservation on a drive keeps it reserved while the drive is in use. The drive is released after a job completes, which allows other jobs to compete for it.
- **Drive pools.** You can assign the drives to drive pools in which one or more drives are combined as a backup target. Jobs submitted to a particular drive pool run on the first available drive in that pool. You can also submit a job to a selected individual drive in the drive pool.

See [“How to use drive pools with the SAN Shared Storage Option”](#) on page 2004.

See [“About device operations with the SAN Shared Storage Option”](#) on page 2003.

About media rotation in the SAN Shared Storage Option

Media rotation jobs are treated the same as backup jobs. You can schedule a media rotation job to run on any device you have access to, such as a local device or a shared storage device. You cannot schedule a media rotation job to run on a device you do not have access to, such as a tape drive attached to the local SCSI bus of another server.

If a media rotation job is scheduled to begin, but all available devices are in use, the job will be placed into the queue.

To successfully use the Media Set Wizard in a shared storage environment, you must use one of the following strategies:

- Restrict use of the Media Set Wizard to a single media server.
- Use the same Overwrite Protection Periods and the same full backup day whenever you use the Media Set Wizard on all media servers in the shared storage environment.
- Edit the jobs and rename the media sets created by each Media Set Wizard used so the jobs and media sets are server-centric.

See [“About scheduling and viewing jobs in the SAN Shared Storage Option”](#) on page 1997.

How to catalog media in the SAN Shared Storage Option

The SAN Shared Storage Option uses a shared catalog database. A tape that has already been cataloged can be physically moved from one device to another and will not have to be recataloged.

If the server is not available over the network when a secondary server generates catalog information, the information is stored temporarily on the secondary server until automatic catalog synchronization occurs.

Because catalogs are shared, information can be restored using any server that has access to a device where the tape resides. If the tape resided in a shared device, or in a local device on the server where you want to perform the restore, simply start a restore job. Otherwise, you have to move the tape into a drive that is accessible.

See [“Creating a new catalog”](#) on page 248.

About sharing media in the SAN Shared Storage Option

The Backup Exec media servers can share media within the shared storage devices, but not simultaneously. For example, media server A can write a backup to a media, and when that job is finished, media server B can append another backup to the same media. Or, if overwrite protection is not enabled, media server B may overwrite that media.

Media sets are not server-centric. In the shared storage environment, all users have a view of all media and media sets. Each media set can contain media in the shared devices and media in any local devices attached to servers.

Note: The default media overwrite protection is not server-centric; this option is set in the shared ADAMM database and affects all media, including media in locally attached devices. For example, if media overwrite protection is set to None by one server, all media in the shared storage environment - including media in other servers' locally attached devices - are immediately available for overwriting.

Media stored in locally-attached secondary storage devices are not accessible by other media servers.

About scheduling and viewing jobs in the SAN Shared Storage Option

Creating backup and restore jobs with the SAN Shared Storage Option is identical to creating jobs in a non-shared storage environment. You can also create test run jobs, resource discovery jobs, and duplicate backup data jobs.

While the SAN Shared Storage Option does not provide a central view of the jobs scheduled on all servers on the SAN, you can always view your scheduled, active, and completed jobs on the media server to which you submitted the jobs by selecting Job Monitor from the navigation bar.

Note: If your job is awaiting a storage device, the Job Monitor window will not display a Device Name. Also, if the Job Status displays as Queued, the job is awaiting an available storage device.

With the SAN Shared Storage Option enabled, all of the media servers share access to the storage devices through the shared ADAMM database. The server that reserves the storage device first runs its job first. Therefore, a job scheduled by a server may not run exactly when scheduled if all the storage devices are being used by other servers.

When a server releases control of a device, there is a short delay before the server looks for additional jobs to process. This delay provides a window of opportunity for other media servers to reserve the shared storage device.

If a device fails during a nonrecurring job, that job will fail and will be rescheduled on hold. If a device fails during a recurring job, the job is rescheduled. The device is then released for the next job scheduled for that device. But, depending on why the device failed, the second job may become trapped. This might prevent other jobs from seeing the device, running to normal completion, or failing and being rescheduled to an "on hold" status. If you determine that a device is malfunctioning, you may want to retarget jobs to another drive or quickly replace the failed drive and resume the jobs that were placed on hold.

See [“Viewing the properties for completed jobs”](#) on page 572.

About sharing robotic libraries between Backup Exec for NetWare Servers and Backup Exec

Backup Exec for NetWare Servers and Backup Exec SAN Shared Storage Options inside the same fibre environment can share robotic libraries, which lowers hardware costs. With robotic library sharing, you first create partitions of the robotic libraries within Backup Exec. Then you can create additional partitions for the same robotic libraries for use within Backup Exec for NetWare Servers.

You can view NetWare servers from a NetWare console and Windows servers from a Windows console. If a Backup Exec for NetWare Servers job is targeted to a drive being used for a Backup Exec job, the drive appears as reserved.

See [“About robotic library sharing prerequisites”](#) on page 1999.

See [“Configuring partitions on Windows media servers for robotic library sharing”](#) on page 2000.

About robotic library sharing prerequisites

Before you can share libraries, you must have the following installed:

- Backup Exec on the Windows media servers.
- Backup Exec for NetWare Servers version 9.0 or later on the NetWare media servers.
- The Backup Exec SAN Shared Storage Option on each Windows media server you want to operate in the shared environment.
- Backup Exec for NetWare Servers SAN Shared Storage Option on each NetWare media server you want to operate in the shared environment.
- The Backup Exec Library Expansion Option or the Backup Exec for NetWare Servers Library Expansion Option.

Note: The drive licenses you purchase for your robotic library are not platform-specific for this implementation. For example, if you plan to share a ten-drive robotic library, you can purchase nine Backup Exec or Backup Exec for NetWare Servers drive licenses (the first drive in a robotic library does not require a Library Expansion Option license).

In order to successfully share robotic libraries, you need a working knowledge of both Backup Exec and of Backup Exec for NetWare Servers. You also should have complete access to hardware and the ability to restart media servers.

Before proceeding with the robotic library sharing configuration, plan how you want to use your robotic library in this shared environment. For example, if you have a robotic library that has 100 slots in it, you may want to partition your robotic library so Backup Exec uses 50 slots and Backup Exec for NetWare Servers uses 50 slots. Factors affecting how many slots you use for each operating system include media rotation schemes, the number of servers you are protecting, and the types of data stored on each server.

Label your media according to operating system or Backup Exec type. Color-coded or unique bar codes identifying whether the tapes are being used with Backup Exec or Backup Exec for NetWare Servers is recommended. This will help you identify the media when you need to restore data or rotate the media back into your media rotation schedule.

Configuring partitions on Windows media servers for robotic library sharing

Before you configure Windows media servers for robotic library sharing, you must ensure that there is currently no backup activity.

No fibre activity or backup jobs should be run until you configure all media servers for robotic library sharing.

To configure partitions on Windows media servers for robotic library sharing

- 1 At the Windows server where the SAN Shared Storage Option is installed, start the Backup Exec Administration Console .
- 2 On the navigation bar, click **Devices**.
The tree pane contains a list of any fibre-attached or locally-attached devices.
- 3 Select the robotic library you want to share.
- 4 Under **Robotic Library Tasks** in the task pane, select **Configure partitions**.
See “[Configure Partitions options](#)” on page 478.
- 5 Set up your partitions.
See “[About robotic library partitions](#)” on page 475.
- 6 From the Devices tree pane, select the drive in this partition that will not be used by your Backup Exec media servers.
- 7 Delete the drive to ensure no jobs run to the unused partition.
- 8 Repeat steps 6 and 7 for all drives in the unused partition.
- 9 Restart all other Windows servers and make sure you can see the shared robotic library on each server.

Configuring partitions on NetWare media servers for robotic library sharing

Before you configure NetWare media servers for robotic library sharing, you must ensure that there is currently no backup activity. You can use either the Backup Exec for NetWare Servers Administration Console or the Administration Console for NetWare to configure partitions on robotic libraries.

Table Y-1 Configuring partitions on NetWare media servers for robotic library sharing

Step	Action
Step 1	<p>Ensure that you are connected to the server and that the partition management feature is enabled.</p> <p>Refer to Using the SAN Shared Storage Option in the Symantec Backup Exec for NetWare Servers documentation for more details.</p>
Step 2	<p>Delete existing partitions.</p> <p>Do one of the following:</p> <ul style="list-style-type: none">■ See “Deleting robotic library partitions from the Backup Exec for NetWare Servers Administration Console” on page 2001.■ See “Deleting robotic library partitions from the Administration Console for NetWare” on page 2002.
Step 3	<p>Create partitions.</p> <p>Do one of the following:</p> <ul style="list-style-type: none">■ See “Creating robotic library partitions using the Backup Exec for NetWare Servers Administration Console” on page 2002.■ See “Creating robotic library partitions using the Administration Console for NetWare” on page 2003.
Step 5	<p>Restart Backup Exec for NetWare Servers on the group servers and make sure you can see the shared library on each server.</p>

Deleting robotic library partitions from the Backup Exec for NetWare Servers Administration Console

You can use the Backup Exec for NetWare Servers Administration Console to delete a partition in a robotic library.

Deleting a partition from the Backup Exec for NetWare Servers Administration Console

- 1 Click **Drives**.
- 2 Click **Partitions**.
- 3 Right-click the partition you want to delete, and then click **Delete**.
- 4 Click **OK**.

Deleting robotic library partitions from the Administration Console for NetWare

You can use the Administration Console for NetWare to delete a partition in a robotic library.

Deleting a partition from the Administration Console for NetWare

- 1 Select **Drives**.
- 2 Select **Partitions**.
- 3 Select the partition you want to delete, and press DEL.
- 4 Press Y.

Creating robotic library partitions using the Backup Exec for NetWare Servers Administration Console

You can use the Backup Exec for NetWare Servers Administration Console to create a robotic library partition in the shared library.

Creating robotic library partitions using the Backup Exec for NetWare Servers Administration Console

- 1 From the Administration Console, click **Drives**.
- 2 Right-click **Partitions**.
- 3 Click **New Partition**.
- 4 Enter options on the **New Partition** dialog box.
- 5 Click **OK**.
- 6 Enter options in the **New Partition General** dialog box.
- 7 Enter options in the **New Partition Configuration** dialog box.
- 8 Click **OK**.

Creating robotic library partitions using the Administration Console for NetWare

You can use the Administration Console for NetWare to create a robotic library partition in the shared library.

Creating robotic library partitions using the Administration Console for NetWare

- 1 From the Administration Console for NetWare, select **Drives**.
- 2 Select **Partitions**.
- 3 Press **INS**.
- 4 Select a drive to be included in this partition.
- 5 Ensure **Configuration** is selected, and press enter to view or edit options on the **Configuration** dialog box.
- 6 Press **F2** to return to the **General** dialog box.
- 7 Press **F2**.

About device operations with the SAN Shared Storage Option

Device operations allow you to manage the physical drives attached to the media servers and perform some operations on the media in the drives. The steps for deleting drives, pausing and resuming drives, completing the inventory drive process, cataloging media, labeling media, and erasing media are identical in the shared and non-shared storage environments.

See [“About scheduling and viewing jobs in the SAN Shared Storage Option”](#) on page 1997.

See [“About devices in the SAN Shared Storage Option”](#) on page 1995.

See [“About sharing media in the SAN Shared Storage Option”](#) on page 1997.

About renaming robotic libraries and drives in the SAN Shared Storage Option

You cannot rename a server, but you can rename robotic libraries and drives to make them more easily identifiable. You may want to use names that are more descriptive of your operations, or you may want to identify the device by the user or location, such as `DATA_CENTER_ROBOTIC LIBRARY`.

The names of all servers attached to the SAN appear when **Devices** is selected from the navigation bar. Press **F5** to manually refresh the screen and view new names.

The robotic libraries and drives can be renamed from any server sharing the ADAMM database, and the new names will appear on all servers on the SAN.

Note: The Device Management window on other Backup Exec administration consoles may need to be manually refreshed before the new names appear on them.

See [“Pausing storage devices”](#) on page 446.

How to use drive pools with the SAN Shared Storage Option

When Backup Exec is installed, the All Devices (<Server Name>) is created by default. In a non-shared storage environment, this default drive pool contains the server’s locally attached drives. In a shared environment, this default drive pool is created for each server using the SAN Shared Storage Option, and contains both locally attached and shared devices.

Symantec recommends creating a shared storage drive pool, which contains only shared devices.

See [“Creating device pools”](#) on page 516.

You can create other drive pools to meet your particular requirements. For example, you may want to create a drive pool for high-performance drives and create a second drive pool for lower-performance drives. High-priority jobs can then be sent to the high-performance drive pool for faster completion.

Drives can belong to more than one drive pool, and drive pools can contain different types of drives. In the shared storage environment, drive pools can contain both local and shared drives, but jobs will run only on those drives in the pool to which the server has access.

For example, suppose you create a drive pool that contains the local drives for both media server A and media server B. If a job is submitted at media server B to this drive pool, the job will run only on available drives attached to server B. If all of server B’s drives are in use, the job has to wait for a drive on server B to become available. If a job was submitted from server B to a drive pool that contained both its local and shared devices, the job would run on the first available drive.

The steps for creating and deleting drive pools, adding or deleting drives from a drive pool, and setting priorities for drives in a drive pool are the same in a shared storage environment as in a non-shared storage environment.

About viewing media in the SAN Shared Storage Option

If you select a drive or select Slots from the Devices view, information for the media contained in the drive or slot appears in the right pane. This information also appears when All Media is selected from the Media view.

See [“General properties for media”](#) on page 261.

See [“How to use drive pools with the SAN Shared Storage Option”](#) on page 2004.

See [“About device operations with the SAN Shared Storage Option”](#) on page 2003.

See [“About scheduling and viewing jobs in the SAN Shared Storage Option”](#) on page 1997.

How to monitor drives in the SAN Shared Storage Option

When you select Devices from the navigation bar, you can view all the physical drives attached to your server, as well as the logical groups they are associated with. All the logical groupings of the physical drives are displayed under Drive Pools as well as all devices locally attached (by SCSI, SATA, SAS, USB, or ATAPI) to all servers on the SAN and the secondary storage units accessed through the SAN.

Expanding the view for All Devices allows you to view all of the storage devices on the SAN. Robotic libraries are listed below each server that has access to that robotic library, even if the robotic library is not directly connected to the server.

You can also run the Device Summary report to monitor drives in the SAN Shared Storage Option.

You can view drive properties for shared devices.

See [“Viewing storage device properties”](#) on page 458.

With shared devices, the Write single block mode and Write SCSI pass-through mode options are selected by default. Selecting these options decreases the chances of dropping critical blocks of data and provides more detailed information when write errors occur. These options are required for FC-connected tape drives.

See [“Policy Jobs Summary Report”](#) on page 757.

About designating a new primary database server and setting up servers in the SAN Shared Storage Option

You can change the SAN Shared Storage Option configuration through Backup Exec Utility (Beutility.exe). This utility lets you designate a new primary database server.

You may wish to replace the primary database server for the following reasons:

- A newer, faster server has become available.
- The database server has stopped functioning.

If the current server is functioning, you should consider specifying the current primary server when installing SAN Shared Storage Option to the new system. This will allow you to test the fibre connections before designating a new database server. If the current database server is not functioning, installing the new system as the primary database server is recommended.

Refer to the *Backup Exec Utility* documentation for more information.

Tips for maintaining the Backup Exec database servers and the shared ADAMM database in the SAN Shared Storage Option

The ADAMM database and the Backup Exec database server are important components of the SAN Shared Storage Option. To protect against possible loss of the ADAMM and catalog databases, you should run frequent backup jobs of the entire Backup Exec directory tree on the primary server.

Note: Each secondary server also has its own local Backup Exec database instance that is independent of the other servers. You should ensure that you back up the Backup Exec database instance on the secondary server since it contains data about the server.

When scheduling backups of the database server's Backup Exec directory, base the frequency of these backup jobs on the rate at which backup sets are being created and the number of tapes affected by all of the media servers on the SAN. All backup sets and tapes affected since the last shared database/catalog server backup would have to be recataloged if all information on the database server was lost.

Create a special media set just for backing up the Backup Exec directory tree and the Windows operating system on the primary server. This will reduce the number of tapes that must be cataloged to find the files for restoring the ADAMM database and catalogs.

Caution: If you allow the backups of these files to go to a large media set, you may have to catalog every tape in that large media set in order to find the latest versions of the ADAMM database and catalogs to restore.

If the primary server is not operational, Backup Exec is unusable on all of the servers on the SAN. The Intelligent Disaster Recovery Option is strongly recommended for protecting each Backup Exec database server. Should this entire system be lost, you can use IDR to quickly recover this system.

If you deem the Backup Exec functions to have a high availability requirement, you should consider setting up one of the other media servers on the storage network as a standby primary server.

Creating a standby primary database server in the SAN Shared Storage Option

You should have a standby server configured and available in case your primary server fails. To avoid data loss if the primary database server fails, Symantec recommends that you save the bedb.bak file and the Catalogs directory to a separate server after the scheduled daily database maintenance.

Refer to the *Backup Exec Utility* documentation for more information about performing the following procedures.

Table Y-2 Creating a standby primary database server in the SAN Shared Storage option

Step	Action
Step 1	Use BEUtility to add all the SAN servers that will use the new primary SAN SSO server.
Step 2	Use BEUtility to create a media server group that contains all the servers you added in step 1. Note: Do not select the Create group from SAN SSO configuration option.
Step 3	Use BEUtility to promote a new SAN SSO server to primary in the media group you created.
Step 4	Use BEUtility to stop the Backup Exec services on all media servers in the media server group.

Table Y-2 Creating a standby primary database server in the SAN Shared Storage option (*continued*)

Step	Action
Step 5	<p>On the media server that you promoted to primary, navigate to \Program Files\Symantec\Backup Exec\Data directory and rename the bedb.bak file to indicate that this is the original file.</p> <p>For example, originalbedb.bak or bedborg.bak</p>
Step 6	<p>On the original primary server, navigate to \Program Files\Symantec\Backup Exec\Data directory. Copy the bedb.bak file to the same directory on the new primary server.</p> <p>If the original primary server is unavailable, locate the latest copy of the bedb.bak file and copy it to the new primary server.</p>
Step 7	<p>On the media server that you promoted to primary, navigate to \Program Files\Symantec\Backup Exec. Rename the Catalogs directory to indicate that this is the original.</p>
Step 8	<p>On the original primary server, navigate to \Program Files\Symantec\Backup Exec\. Copy the Catalogs directory to the same directory on the new primary server.</p> <p>If the original primary server is unavailable, locate the latest copy of the Catalogs directory and copy it to the new primary server.</p>
Step 9	<p>Use BEUtility to restore the database you copied in step 6. Be sure to select the Drop existing database and reload from backup option.</p>
Step 10	<p>Use BEUtility to start the Backup Exec services on all the media servers in the media server group.</p>

See [“About the the Intelligent Disaster Recovery Configuration Wizard”](#) on page 1810.

See [“Restoring data by setting job properties”](#) on page 605.

About starting and stopping Backup Exec Services on multiple servers in the SAN Shared Storage Option

Stopping Backup Exec services is the first step in system maintenance. After system maintenance, you can start the services again. You can start and stop services at the same time, called "bouncing", which refreshes the database, re-establishes connections, and forces the system back into synchronization. It is similar to rebooting all the servers.

See [“Starting and stopping Backup Exec services”](#) on page 168.

About reconfiguration of the SAN Shared Storage Option environment

You can change the primary server to which a secondary server is assigned. First, you must make the secondary server into a stand-alone server. Otherwise, all of the SAN SSO device configuration data migrates to the primary server database.

When the ADAMM service restarts, the standard device discovery process rediscovers the physical devices that are attached to the server.

When you configure a stand-alone server as a SAN SSO secondary server, the standard device discovery process rediscovers the physical devices that are attached to the server. You cannot recover original backup-to-disk specifications. You must recreate the backup-to-disk folder and enter the original path where the folder resides. You must then run an inventory job to discover the backup-to-disk media. Before you can restore data, catalog the media.

Reconfiguring management of a secondary server

Use the following process to reconfigure management of a secondary server.

Refer to the *Backup Exec Utility* documentation for more information about performing the following procedures.

Table Y-3 Reconfiguring management of a secondary server

Step	Action
Step 1	Use BEUtility to convert the secondary server to a stand-alone server.
Step 2	Use BEUtility to set the primary SAN SSO server.

Troubleshooting failed components in the SAN Shared Storage Option

Various problems can occur at any location in a SAN. In order for Backup Exec to work properly, a device has to be recognized in three locations; the bridge/router must recognize it as a SCSI device, the operating system must recognize it as a device, and Backup Exec must recognize it as a supported device. In some cases, there will be a problem with your hardware that will require you to contact your hardware vendor for technical support.

You may need to replace a component of your SAN, such as a bridge or hub. For specific steps for replacing your equipment, refer to your hardware vendor's documentation.

See "[Troubleshooting offline devices in the SAN Shared Storage Option](#)" on page 2010.

Troubleshooting offline devices in the SAN Shared Storage Option

If a device in your SAN has gone offline, follow these steps to determine the source of the problem.

Before you begin troubleshooting, verify that your devices are on the Backup Exec supported device list.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Also verify that all hardware drivers are up to date and started. If you find errors with your hardware, contact your hardware vendor for specific instructions.

Table Y-4 Troubleshooting offline devices in the SAN Shared Storage Option

Step	Action
Step 1	Use the Windows Device Manager to verify that the operating system recognizes the device. If the device is not recognized, you may need to troubleshoot the device. See " Finding hardware errors for the SAN Shared Storage Option " on page 2012.
Step 2	For robotic libraries, verify that robotic library support is installed.

Table Y-4 Troubleshooting offline devices in the SAN Shared Storage Option
(continued)

Step	Action
Step 3	<p>Check the system event log for the following errors, which indicate SAN communication errors: SCSI errors 9, 11, and 15, or timeout errors relating to storage. Check the application event log for multiple events 33152. These events indicate SAN communication errors.</p> <p>See "Finding hardware errors for the SAN Shared Storage Option" on page 2012.</p> <p>You may need to contact your hardware vendor.</p>
Step 4	<p>If the library is online, but some or all of the drives are offline, use Backup Exec to initialize the library.</p> <p>To initialize the library, do the following in the order listed:</p> <ul style="list-style-type: none"> ■ On the navigation bar, click Devices. ■ Select the library, and then under Device Tasks, click Initialize.
Step 5	<p>If initializing the library does not bring the devices online, check the library for an error display on the front panel, mechanical problems, or tapes inappropriately in the drives. Correct any errors that you find.</p>
Step 6	<p>If no errors exist on the library or if you corrected the errors and the devices are still offline, stop Backup Exec services and then restart them when all Backup Exec jobs are inactive in the SAN.</p>
Step 7	<p>If restarting the services does not bring the devices online, restart the operating system. Be sure that no Backup Exec jobs are running when you restart.</p>

Table Y-4 Troubleshooting offline devices in the SAN Shared Storage Option
(continued)

Step	Action
Step 8	<p>If restarting the operating system does not bring the devices online, reset the SAN to help identify problem tape devices. Recycling the SAN may also resolve fibre problems.</p> <p>See “Resetting the SAN in the SAN Shared Storage Option” on page 2013.</p>

Finding hardware errors for the SAN Shared Storage Option

Use the following steps to find common hardware errors that occur in a SAN environment. If you find errors with your hardware, contact your hardware vendor for specific instructions.

Table Y-5 Finding hardware errors for the SAN Shared Storage Option

Step	Action
Step 1	Verify that the proper device drivers were installed.
Step 2	Verify that the fibre cable is securely connected to the HBA and the switch.
Step 3	<p>Verify that the SCSI bridge is properly connected to the library and the switch.</p> <p>Apply normal SCSI troubleshooting techniques at the fibre to SCSI bridge. Use a bridge administration tool to verify that the bridge recognizes all of the devices. Also verify that the bridge firmware is up to date.</p>
Step 4	<p>Check for a failed hardware component between the server and the switch.</p> <p>Sometimes some of the servers in the SAN recognize the tape devices, but other servers do not. If none of the servers in the SAN recognize the tape devices, check for a failed hardware component between the switch and the tape devices.</p>

Table Y-5 Finding hardware errors for the SAN Shared Storage Option
(continued)

Step	Action
Step 5	Reset the SAN, which may identify problem hardware components and may resolve fibre problems.

Resetting the SAN in the SAN Shared Storage Option

Resetting the SAN involved powering down the components of the SAN, and then powering them on in a specific order.

Table Y-6 Resetting the SAN in the SAN Shared Storage Option

Step	Action
Step 1	Turn off all servers, robotic libraries, and SCSI to fibre bridges in the SAN. In rare cases, you may need to power down the switch also. If you need to power down the switch, you should power it on before any other components and wait for all checks to complete before powering on other components.
Step 2	Turn on the robotic library. See “Creating a job to initialize a robotic library” on page 484.
Step 3	Turn on the bridge.
Step 4	Verify that the switch recognizes the library.
Step 5	Turn on the primary SAN SSO server.
Step 6	Verify that the operating system recognizes the robotic library and drives.
Step 7	Turn on one of the secondary servers. Wait for the secondary server to boot before powering on the other secondary servers.

Bringing devices online after an unsafe device removal event in the SAN Shared Storage Option

If a device is in use by Backup Exec at the time of an unsafe device removal event, the device will go offline in Backup Exec.

Table Y-7 How to bring a device online after an unsafe device removal event

Step	Action
Step 1	Verify that no Backup Exec jobs are running in the SAN.
Step 2	Use Backup Exec to initialize the library if the library is online, but the drives are offline. See “Creating a job to initialize a robotic library” on page 484.
Step 3	Stop all Backup Exec services and then restart them if the library is offline or if the drives are offline after initialization. See “Starting and stopping Backup Exec services” on page 168. If the device is not online, you may need to troubleshoot the device. See “Finding hardware errors for the SAN Shared Storage Option” on page 2012.

Best practices for the SAN Shared Storage Option

Review the following recommendations for SAN SSO:

- Before installing Backup Exec, be sure that all hardware in the SAN is working and configured properly.
- Make sure that the primary server is the fastest server and has the fewest extraneous operations.
- Use a separate primary server for each SAN.
- Make sure that the HBA drivers, SCSI to fibre bridges, and library firmware have been updated to the hardware vendor’s most current release.
- Make sure that all HBA cards on the SAN are using the same and the most current firmware and driver levels.

- Keep the servers in the SAN in the same Microsoft Administration domain. Cross-domain environments can cause authentication problems during install and can block access to resources during backups.
- Change the display names of libraries and drives to names that reflect the servers or jobs for which you will use them.
- Run frequent backups of the entire Backup Exec directory tree on each Backup Exec database server in the SAN.
- Create a separate media set to use only for backups of the Backup Exec directory tree and the Windows operating system on the primary database server. This will reduce the number of tapes you need to catalog to find the files to restore the ADAMM database and catalogs.
- Use a switch administration tool to verify that each server is in a zone configuration with the tape devices.
- You must manually refresh Backup Exec administration consoles in a SAN SSO configuration if a member server updates the database. To manually refresh the user interface, right-click the proper component, and then press F5.

Symantec Backup Exec Storage Provisioning Option

This appendix includes the following topics:

- [About the Storage Provisioning Option](#)
- [Requirements for the Storage Provisioning Option](#)
- [Requirements for the Storage Provisioning Option in a CASO environment](#)
- [About installing the Storage Provisioning Option](#)
- [Viewing storage array components in Backup Exec](#)
- [About using the Storage Array Configuration Wizard](#)
- [Configuring a storage array by using the Storage Array Configuration Wizard](#)
- [Viewing properties for storage arrays](#)
- [Properties of physical disks on storage arrays](#)
- [About the All Virtual Disks device pool in the Storage Provisioning Option](#)
- [About virtual disks in the Storage Provisioning Option](#)
- [About hot spares in the Storage Provisioning Option](#)
- [Detecting a new storage array](#)
- [Renaming a virtual disk or storage array](#)
- [About identifying the physical disks of a virtual disk](#)
- [About predicting disk usage in the Storage Provisioning Option](#)

- [Configuring an alert for low disk space on storage arrays](#)
- [Default options for Storage Provisioning Alert](#)
- [Troubleshooting the Storage Provisioning Option](#)

About the Storage Provisioning Option

The Storage Provisioning Option lets you configure, manage, and monitor storage arrays that are attached to the media server.

Note: If you use a Dell DL Appliance, do not use this appendix. See the *Dell™ PowerVault™ DL Backup to Disk Appliance and the Symantec Backup Exec Storage Provisioning Option* documentation that Dell provides with the appliance.

Table Z-1 Features of the Storage Provisioning Option

Feature	Description
Discovery of new storage arrays, physical disks, and virtual disks	Backup Exec can discover new storage arrays, physical disks, and the virtual disks that you add to a storage array. If you create virtual disks by using storage array vendor tools or the Microsoft Storage Manager for SANs utility, Backup Exec also detects those virtual disks.
A wizard to help you configure a storage array for use with Backup Exec	Backup Exec provides the Storage Array Configuration wizard to help you configure virtual disks on a storage array. The virtual disks are added to the All Virtual Disks device pool. Backup Exec then uses the virtual disks in the device pool as destination devices for jobs.
Trend analysis of disk space usage	Backup Exec collects statistical information to predict the amount of disk space that is required on the storage arrays. Alerts are sent if the available disk space does not meet the predicted amount of disk space that is needed.
Alerts for low disk space	Backup Exec sends an alert when available disk space reaches each of three thresholds that you set for a virtual disk.

See [“Requirements for the Storage Provisioning Option”](#) on page 2019.

See [“Configuring a storage array by using the Storage Array Configuration Wizard”](#) on page 2022.

See [“About the All Virtual Disks device pool in the Storage Provisioning Option”](#) on page 2029.

Requirements for the Storage Provisioning Option

Do the following before you install the Backup Exec Storage Provisioning Option:

- Ensure that Virtual Disk Service (VDS) 1.1 is installed on the media server. VDS 1.1 is installed with Windows Server 2003 Service Pack 2 and Windows Server 2008. To install VDS 1.1 on Windows Server 2003 R2, install Service Pack 2 or the Microsoft Storage Manager for SANs management tool.
- Attach any storage arrays to the media server.
- Install the storage array vendor's VDS hardware provider on the media server.

See [“How to choose the location for CASO device and media data”](#) on page 1509.

See [“Requirements for the Storage Provisioning Option in a CASO environment”](#) on page 2019.

See [“About installing the Storage Provisioning Option”](#) on page 2020.

See [“About using the Storage Array Configuration Wizard”](#) on page 2021.

Requirements for the Storage Provisioning Option in a CASO environment

The following are required to run the Storage Provisioning Option in a Central Admin Server Option (CASO) environment:

- The Storage Provisioning Option must be installed on the media server to which the storage array is attached.
If the storage array is attached to a managed media server, install the Storage Provisioning Option on that managed media server. You do not need to install the Storage Provisioning Option on the central administration server if the storage array is not attached to it.
- The Central Admin Server Option must use a centralized database.
See [“About CASO catalog locations”](#) on page 1541.

See [“About using the Storage Array Configuration Wizard”](#) on page 2021.

See [“About installing the Storage Provisioning Option”](#) on page 2020.

About installing the Storage Provisioning Option

Install the Storage Provisioning Option on a local media server as a separate add-on component of Backup Exec.

You can install the Storage Provisioning Option when you upgrade from a previous version of Backup Exec. However, the default device pool from the previous version is kept. The **All Virtual Disks** device pool is not set as the default device pool.

See [“Installing additional Backup Exec options to the local media server”](#) on page 123.

See [“About the All Virtual Disks device pool in the Storage Provisioning Option”](#) on page 2029.

See [“Requirements for the Storage Provisioning Option in a CASO environment”](#) on page 2019.

Viewing storage array components in Backup Exec

After you install the Storage Provisioning Option, storage arrays appear in the **Devices** view. After you use the **Storage Array Configuration Wizard** to configure the storage array, virtual disks appear under the storage array to which they belong.

Physical disks do not appear in the **Devices** view under the storage arrays. You can view physical disks in the storage array properties, and in the right pane of the **Devices** view when you select a storage array.

To view storage array components in Backup Exec

- 1 On the navigation bar, click **Devices**.
- 2 Expand a media server that has an attached storage array.
- 3 Expand a storage array for which you want to view properties.
- 4 View the storage array components.

See [“About using the Storage Array Configuration Wizard”](#) on page 2021.

See [“Configuring a storage array by using the Storage Array Configuration Wizard”](#) on page 2022.

See [“Properties of physical disks on storage arrays”](#) on page 2026.

About using the Storage Array Configuration Wizard

The Storage Provisioning Option provides the **Storage Array Configuration Wizard** to help you configure virtual disks in a storage array.

With the **Storage Array Configuration Wizard**, you can:

- Specify a virtual disk size, and then let the storage array automatically choose the physical disks for the virtual disk's disk group.
With this option, the storage array attempts to spread the required disk space for the virtual disk equally across each physical disk it selects.
When you specify a virtual disk size, the **Storage Array Configuration Wizard** presents you with the total amount of disk space available to create a virtual disk. The total available disk space is derived from the physical disks that are present. You can use the entire space to create one virtual disk, or you can specify smaller amounts of space and create multiple virtual disks. Each virtual disk you create is based on the amount of disk space you enter. To create multiple virtual disks, run the wizard multiple times, until you have the virtual disks you want, or until you exhaust the disk space.
- Create a disk group by manually selecting the physical disks, from which you can then choose the number of virtual disks to create.
With this option, the **Storage Array Configuration Wizard** lets you customize the entire virtual disk creation process on the storage array.
- Configure hot spares.
You can also use this wizard to add or change hot spares for the disk groups that are already configured.

In addition, you can also use the **Storage Array Configuration Wizard** to configure a virtual disk as a deduplication storage folder. Deduplication reduces the disk storage and the network bandwidth that is used for backups by storing and sending only unique data.

Note: Deduplication may not be available for the storage array if a deduplication storage folder presently exists on the media server. Only one deduplication storage folder per media server is supported. You can delete the existing deduplication storage folder and then rerun the **Storage Array Configuration Wizard** to enable deduplication on the storage array.

When the wizard completes, it runs a utility job named Configure Storage Array. This utility job creates the virtual disks that you specified. Then, Backup Exec adds the virtual disks to a device pool named **All Virtual Disks**. You can submit jobs to the **All Virtual Disks** device pool, to the storage array, or to a specific virtual disk.

See [“About deduplication storage folders”](#) on page 1579.

See [“Requirements for the Deduplication Option”](#) on page 1572.

See [“Configuring a storage array by using the Storage Array Configuration Wizard”](#) on page 2022.

Configuring a storage array by using the Storage Array Configuration Wizard

Use the **Storage Array Configuration Wizard** to configure a storage array for use with the Backup Exec Storage Provisioning Option.

Backup Exec then submits a Configure Storage Array job to create the virtual disks.

Note: In a Central Admin Server Option (CASO) environment, run the **Storage Array Configuration Wizard** from the central administration server. You can run the **Storage Array Configuration Wizard** for any managed media server that has the Storage Provisioning Option installed. Managed media servers can share a single storage array but cannot share a virtual disk on a storage array.

See [“About using the Storage Array Configuration Wizard”](#) on page 2021.

See [“Viewing storage array components in Backup Exec”](#) on page 2020.

See [“About predicting disk usage in the Storage Provisioning Option”](#) on page 2046.

See [“Configuring an alert for low disk space on storage arrays”](#) on page 2047.

To configure a storage array by using the Storage Array Configuration Wizard

- 1 On the **Tools** menu, click **Configure Devices**.
- 2 Click **Configure Storage Array**.
- 3 Do one of the following:

In a non-CASO environment

On the **Welcome** panel, select the storage array that you want to configure, and then click **Next**.

In a CASO environment

Do the following in the order listed:

- On the **Welcome** panel, select the managed media server that you want to access the virtual disks.
- Select the storage array that you want to configure.
- Click **Next**.

4 Follow the on-screen directions.

Viewing properties for storage arrays

Properties provide detailed information, such as statistics and settings.

To view properties for storage arrays

- 1 On the navigation bar, click **Devices**.
- 2 Do one of the following:
 - Expand a media server that has an attached storage array.
 - Right-click the storage array for which you want to view properties, and then click **Properties**.
 - Select the item for which you want to view properties, and then in the task pane under **General Tasks**, click **Properties**.
See [“General properties for storage arrays”](#) on page 2023.

General properties for storage arrays

General properties provide information on the name, status, and health for a storage array.

See [“Viewing properties for storage arrays”](#) on page 2023.

Table Z-2 General properties for storage arrays

Item	Description
Name	<p>Displays the name that Backup Exec assigns to the storage array when you use the Storage Array Configuration Wizard. The name is ARRAYx, where x is a number that increments each time that you add a storage array.</p> <p>See “Renaming a virtual disk or storage array” on page 2044.</p>
Status	<p>Indicates the current status of the storage array.</p> <p>Statuses are as follows:</p> <ul style="list-style-type: none"> ■ Pause The storage array is temporarily stopped. You can pause a storage array to perform maintenance activities. Active jobs are not affected if they start before the storage array is paused. See “Pausing storage devices” on page 446. ■ Enable The storage array is available for use with Backup Exec. If the storage array is not enabled, it is available for use with other applications. ■ Online The storage array is available for use. ■ Offline Backup Exec cannot access the storage array. You can click this check box to try to bring the storage array online. ■ Failed Backup Exec cannot access the storage array.
Total capacity	<p>Displays the expected amount of total raw capacity of all of the physical disks in the storage array.</p> <p>If a storage array can read the total capacity, then it appears in this field. Otherwise, Backup Exec estimates the total capacity.</p>

Table Z-2 General properties for storage arrays (*continued*)

Item	Description
Used capacity	Displays the amount of raw capacity of all of the physical disks that are used in the storage array. Backup Exec calculates used capacity by subtracting available capacity from total capacity.
Configured capacity	Displays the expected amount of total raw capacity of all of the configured physical disks in the storage array.
Unconfigured capacity	Displays the expected amount of total raw capacity of all of the unconfigured physical disks in the storage array.
Hardware name	Displays the name that the storage array hardware or the vendor hardware provider assigns.
Hardware status	Indicates the hardware status. The values for hardware status are as follows: <ul style="list-style-type: none">■ OK The storage array is online.■ Offline The storage array and the virtual disks on the storage array are offline. Backup Exec cannot access them. Refer to the vendor documentation and management software that are supplied with the storage array. After the storage array is brought online, the virtual disks are automatically brought online.

Table Z-2 General properties for storage arrays (*continued*)

Item	Description
Hardware health	<p>Indicates the hardware health.</p> <p>Values for hardware health are as follows:</p> <ul style="list-style-type: none"> ■ OK The storage array is online. ■ Warning The storage array may fail or produce errors, but it is currently operational. The storage array and the virtual disks on the storage array are offline. Backup Exec cannot access them. ■ Critical The storage array has failed. The storage array and the virtual disks on the storage array are offline. Backup Exec cannot access them. <p>Refer to the vendor documentation and management software that are supplied with the storage array. After the storage array is brought online, the virtual disks are automatically brought online.</p>
Identification	<p>Displays the information that the vendor provides to assist you with identification of the attached storage array hardware. Refer to the vendor documentation that is supplied with the storage array for information.</p>

Properties of physical disks on storage arrays

You can view properties for the physical disks in a storage array.

See [“Viewing properties for storage arrays”](#) on page 2023.

Table Z-3 Properties of physical disks on storage arrays

Item	Description
Enclosure	<p>Identifies the enclosure that the physical disk is in.</p>

Table Z-3 Properties of physical disks on storage arrays (*continued*)

Item	Description
Slot	Identifies the slot that the physical disk occupies.
Capacity	Displays the total amount of available disk space on the physical disk in this slot.
State	Displays the state of the physical disk. States are as follows: <ul style="list-style-type: none"><li data-bbox="821 569 1245 651">■ Hot spare The physical disk is configured as a hot spare.<li data-bbox="821 661 1245 718">■ Configured The physical disk is configured for use.<li data-bbox="821 729 1245 840">■ Configurable The physical disk is available for configuration so that Backup Exec can use it.<li data-bbox="821 850 1245 939">■ Unconfigurable The physical disk cannot be configured because it is in a bad state, or it has failed.<li data-bbox="821 949 1245 1031">■ Allocated The physical disk is in the process of configuration.

Table Z-3 Properties of physical disks on storage arrays (*continued*)

Item	Description
Status	<p>Displays the hardware status.</p> <p>Values for hardware status are as follows:</p> <ul style="list-style-type: none"> ■ OK The physical disk is online. ■ Offline The physical disk is offline. The virtual disks that use this physical disk may also be offline. Backup Exec cannot access them. ■ Failed The physical disk has failed. The virtual disks that use this physical disk may also fail. Backup Exec cannot access the virtual disks. If hot spares are configured, the virtual disk is automatically rebuilt. If your storage array does not support an automatic rebuild capability, you must use vendor tools to perform a manual rebuild of the virtual disks. Refer to your vendor documentation for the storage array for more information. <p>To troubleshoot issues, refer to the vendor documentation and management software that are supplied with the storage array.</p>
Health	<p>Displays the hardware health.</p> <p>Values for hardware health are as follows:</p> <ul style="list-style-type: none"> ■ OK The physical disk is online. ■ Warning The physical disk may fail or produce errors, but it is currently operational. ■ Critical The physical disk may fail. You should replace the physical disk. <p>To troubleshoot issues, refer to the vendor documentation and management software that are supplied with the storage array.</p>

About the All Virtual Disks device pool in the Storage Provisioning Option

After you install the Storage Provisioning Option, Backup Exec adds the **All Virtual Disks** device pool to the list of device pools. The **All Virtual Disks** device pool contains all virtual disks from all storage arrays on all computers in the Backup Exec environment.

Note: Individual virtual disks do not appear in the **All Devices <computer_name>** device pool. However, you can select a specific virtual disk as a destination device for a job.

You cannot add nonvirtual disk devices to the **All Virtual Disks** device pool.

See [“About using the Storage Array Configuration Wizard”](#) on page 2021.

See [“About virtual disks in the Storage Provisioning Option”](#) on page 2029.

See [“About device pools”](#) on page 515.

About virtual disks in the Storage Provisioning Option

A virtual disk is a logical disk that you create on a storage array to provide virtual storage to the media server.

You can use any of the following to create a virtual disk:

- The Storage Array Configuration Wizard
- The management tools that the vendor of the storage array provides
- The Microsoft Storage Manager for SANs management tool

If you create a virtual disk with a tool other than the **Storage Array Configuration Wizard**, you must configure the virtual disk for use with Backup Exec. After you configure a virtual disk, Backup Exec uses it as a destination device for jobs. Backup Exec automatically adds configured virtual disks to the **All Virtual Disks** device pool.

See [“Configuring a virtual disk on a storage array”](#) on page 2034.

In the **Storage Array Configuration Wizard**, you specify the number of virtual disks to create from the physical disks that are in the storage array. The media server cannot access the physical disks. The media server can access only the virtual disks that you create.

Backup Exec uses a configured virtual disk in the same manner in which it uses a backup-to-disk folder.

See [“About backup-to-disk folders”](#) on page 496.

Backup Exec does not assign a drive letter to the virtual disk. You cannot browse for a virtual disk or access it from a command prompt. Since you cannot browse to the virtual disk, you cannot back it up with Backup Exec. Symantec recommends that you create a duplicate backup data job to move the data from the virtual disk to another device. For example, you can move the data to a tape or to another virtual disk on a separate storage array.

You can configure three low disk space thresholds for the virtual disks. As available disk space reaches each threshold, Backup Exec sends an alert. When the available disk space on the virtual disk reaches the third threshold, the alert warns you to create more disk space immediately.

You can configure these thresholds as defaults that apply to all new virtual disks, or as defaults that apply to a specific virtual disk.

See [“Editing default options for a virtual disk on a storage array”](#) on page 2030.

Note: You cannot share a virtual disk between two computers.

See [“Viewing storage array components in Backup Exec”](#) on page 2020.

See [“About predicting disk usage in the Storage Provisioning Option”](#) on page 2046.

See [“Editing the default options for all virtual disks on storage arrays”](#) on page 2032.

See [“Viewing properties for unconfigured virtual disks on a storage array”](#) on page 2035.

See [“About the All Virtual Disks device pool in the Storage Provisioning Option”](#) on page 2029.

Editing default options for a virtual disk on a storage array

You can set the default options that apply to individual virtual disks.

See [“About the All Virtual Disks device pool in the Storage Provisioning Option”](#) on page 2029.

See [“Editing general properties of virtual disks on storage arrays”](#) on page 2037.

To edit default options for a virtual disk on a storage array

- 1 On the navigation bar, click **Devices**.
- 2 Expand a media server that has an attached storage array.

- 3 Expand a storage array, and then select the virtual disk that you want to view.
- 4 In the task pane, under **General Tasks**, click **Properties**.
- 5 On the **General** tab, change the information as appropriate.
See “[General properties for virtual disks on storage arrays](#)” on page 2038.
- 6 On the **Advanced** tab, change the information as appropriate.
See “[Advanced properties for storage arrays](#)” on page 2031.
- 7 Click **OK**.

Advanced properties for storage arrays

Advanced properties for storage arrays provide information about the low disk space thresholds and buffered reads and writes.

See “[Editing default options for a virtual disk on a storage array](#)” on page 2030.

Table Z-4 Advanced properties for storage arrays

Item	Description
First threshold	<p>Displays the first low disk space threshold at which you want Backup Exec to send an alert. You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes.</p> <p>When the amount of used disk space reaches this threshold, Backup Exec sends an alert. The default threshold is 75%, which is a percentage of the total available disk space on this virtual disk.</p>
Second threshold	<p>Displays the second low disk space threshold at which you want Backup Exec to send an alert. You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes.</p> <p>When the amount of used disk space reaches this threshold, Backup Exec sends an alert. The default threshold is 85%, which is a percentage of the total available disk space on this virtual disk.</p>

Table Z-4 Advanced properties for storage arrays (*continued*)

Item	Description
Third threshold	<p>Displays the third low disk space threshold at which you want Backup Exec to send an alert. You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes.</p> <p>When the amount of used disk space reaches this threshold, Backup Exec sends an alert. The default threshold is 95%, which is a percentage of the total available disk space on this virtual disk.</p>
Auto detect settings	<p>Indicates if Backup Exec automatically detects the preferred settings for this device.</p> <p>Uncheck Auto detect settings if you want to set buffered reads or writes.</p> <p>The default setting is On.</p>
Buffered read	<p>Indicates the following:</p> <ul style="list-style-type: none"> ■ You do not want Backup Exec to automatically detect settings for this device. ■ You want this device to allow buffered reads, which is the reading of large blocks of data. <p>If you enable buffered read, increased performance may result.</p> <p>The default setting is Off.</p>
Buffered write	<p>Indicates the following:</p> <ul style="list-style-type: none"> ■ You do not want Backup Exec to automatically detect settings for this device. ■ You want this device to allow buffered writes, which is the writing of large blocks of data. <p>The default setting is On.</p>

Editing the default options for all virtual disks on storage arrays

You can set the defaults that apply to all new virtual disks on the storage arrays.

See “[Default options for all virtual disks on storage arrays](#)” on page 2033.

To edit the default options for all virtual disks on storage arrays

- 1 On the **Tools** menu, click **Options**.
- 2 In the **Properties** pane, under **Settings**, click **Virtual Disk**.
- 3 Set the following defaults as appropriate:
- 4 Click **OK**.

Default options for all virtual disks on storage arrays

Default options provide information about all of the virtual disks on storage arrays.

Table Z-5 Default options for all virtual disks on storage arrays

Item	Description
Maximum number of backup sets per file	<p>Displays the maximum number of backup sets to write to a file on a virtual disk. The maximum number can be from one to 8192. The default number is 100.</p> <p>If you specify fewer rather than more backup sets in a file, Backup Exec may be able to reclaim disk space faster. For example, you specify 100 backup sets per file. Backup Exec cannot reclaim any disk space until the overwrite protection period expires for all of the backup sets. If you specify one backup set per file, Backup Exec can reclaim disk space as soon as the overwrite protection period expires for that set.</p>
Maximum size for files	<p>Displays the maximum size for each file that this virtual disk contains. Select either MB or GB as the unit of size. The file size can be from one MB to 4096 GB. The default size is four GB.</p> <p>If you create small but numerous files, performance may slow since the computer must still process each file. However, if you create large files, file system limitations can cause memory allocation problems or network issues. These issues can be a problem if you store files across a network.</p>

Table Z-5 Default options for all virtual disks on storage arrays *(continued)*

Item	Description
Allow x concurrent jobs	Displays the number of concurrent operations that you want to allow to this virtual disk. This number can be from one to 16.
Threshold 1	Displays the low disk space threshold at which you want Backup Exec to send the first of three alerts. The default threshold is 75%, which is a percentage of the total available disk space on this virtual disk. You can change the default, and you can change the amount of space from a percentage to megabytes or gigabytes.
Threshold 2	Displays the low disk space threshold at which you want Backup Exec to send the second of three alerts. The default threshold is 85%, which is a percentage of the total available disk space on this virtual disk. You can change the default, and you can change the amount of space from a percentage to megabytes or gigabytes.
Threshold 3	Displays the low disk space threshold at which you want Backup Exec to send the last of three alerts. The default threshold is 95%, which is a percentage of the total available disk space on this virtual disk. You can change the default, and you can change the amount of space from a percentage to megabytes or gigabytes.

Configuring a virtual disk on a storage array

If you create a virtual disk with a tool other than Backup Exec, then you must configure the virtual disk to use it with Backup Exec. Backup Exec can only use configured virtual disks as destination devices for jobs. When you configure the virtual disk, Backup Exec submits a job named Configure Virtual Disk. When the job completes successfully, the virtual disk is configured and added to the **All Virtual Disks** device pool.

Note: Be careful when you select an unconfigured virtual disk. An unconfigured virtual disk may be in use as a Microsoft SQL Server database, an Exchange database, or a boot disk.

To configure a virtual disk on a storage array

- 1 On the navigation bar, click **Devices**.
- 2 Expand a media server that has an attached storage array.
- 3 Expand a storage array, and then select the unconfigured virtual disk.
- 4 In the task pane, under **Device**, click **Configure**.
- 5 When you are prompted, click **OK** to configure the virtual disk.

See [“Editing the default options for all virtual disks on storage arrays”](#) on page 2032.

Viewing properties for unconfigured virtual disks on a storage array

You can view the properties of an unconfigured virtual disk on a storage array.

Note: You must configure a virtual disk before Backup Exec can use it as a destination device for jobs.

See [“Configuring a virtual disk on a storage array”](#) on page 2034.

To view properties for unconfigured virtual disks on a storage array

- 1 On the navigation bar, click **Devices**.
- 2 Expand a media server that has an attached storage array.
- 3 Expand a storage array, and then select an unconfigured virtual disk.
- 4 In the task pane, under **General Tasks**, click **Properties**.
- 5 On the **Virtual Disk (Unconfigured) Properties** dialog box, view the properties information.

Properties for unconfigured virtual disks on storage arrays

Properties for unconfigured virtual disks provide information on the name, status, and health of the disks.

See [“Viewing properties for unconfigured virtual disks on a storage array”](#) on page 2035.

Table Z-6 Properties for unconfigured virtual disks on storage arrays

Item	Description
Name	<p>Displays the name of the unconfigured virtual disk.</p> <p>The default name is VIRTDISK x, where x is a number that increments each time that you create a virtual disk.</p> <p>See “Renaming a virtual disk or storage array” on page 2044.</p>
Hardware name	<p>Displays the name that you assign to a virtual disk if you use a vendor-specific tool to create the virtual disk.</p>
Hardware status	<p>Displays the hardware status.</p> <p>Values for the hardware status are as follows:</p> <ul style="list-style-type: none"> ■ OK The unconfigured virtual disk is online. ■ Offline The unconfigured virtual disk is offline. ■ Failed The unconfigured virtual disk has failed.
Hardware health	<p>Displays the hardware health.</p> <p>Values for the hardware health are as follows:</p> <ul style="list-style-type: none"> ■ OK The unconfigured virtual disk is online. ■ Warning The unconfigured virtual disk may fail or produce errors, but it is currently operational. ■ Critical The unconfigured virtual disk has failed. ■ Unspecified The unconfigured virtual disk is in the process of configuration.

Table Z-6 Properties for unconfigured virtual disks on storage arrays
(continued)

Item	Description
Disk classification	<p>Displays the type of disk group that the unconfigured virtual disk is on.</p> <p>Disk classifications are as follows:</p> <ul style="list-style-type: none"> ■ Simple (RAID 0) A single physical disk, no striping, or parity. No redundancy. ■ Span A set of multiple physical disks that are concatenated together. No striping or parity. No redundancy. ■ Stripe A set of multiple physical disk extents with data striped across the physical disks. No redundancy. ■ Mirror (RAID 1) A pair or multiple pairs of physical disks with the same data written to each physical disk of the pair. Provides for data redundancy. ■ Stripe with parity (RAID 5 or RAID 6) Three or more physical disks with data striped across the physical disks, with one disk's worth of space that is used for parity. Provides for data redundancy. ■ Unknown <p>Backup Exec creates only physical disk groups with a disk classification of Stripe with parity (RAID 5/RAID 6). If another disk classification appears, the disk group was created with a tool other than the Storage Provisioning Option</p>

Editing general properties of virtual disks on storage arrays

You can edit general properties for a virtual disk on a storage array.

To edit general properties of virtual disks on storage arrays

- 1 On the navigation bar, click **Devices**.
- 2 Expand a media server that has an attached storage array.

- 3 Expand a storage array, and then select the virtual disk that you want to view.
- 4 In the task pane, under **General Tasks**, click **Properties**.
- 5 On the **General** tab, edit the properties as appropriate.
See [“Viewing properties for unconfigured virtual disks on a storage array”](#) on page 2035.
- 6 Click **OK**.
See [“Editing default options for a virtual disk on a storage array”](#) on page 2030.

General properties for virtual disks on storage arrays

General properties provide information about virtual disks on storage arrays.

See [“Editing general properties of virtual disks on storage arrays”](#) on page 2037.

Table Z-7 General properties for virtual disks on storage arrays

Item	Description
Name	Displays the name that Backup Exec assigns to the virtual disk when you use the Storage Array Configuration Wizard . The name is VIRTDISKx, where x is a number that increments each time that you add a storage array. See “Renaming a virtual disk or storage array” on page 2044.

Table Z-7 General properties for virtual disks on storage arrays (*continued*)

Item	Description
Status	<p>Displays the current status of the virtual disk.</p> <p>Statuses for a virtual disk are as follows:</p> <ul style="list-style-type: none">■ Pause The virtual disk is temporarily stopped. See “Pausing storage devices” on page 446.■ Enable The virtual disk is available for use with Backup Exec. If the virtual disk is disabled, it is available for use with other applications. Backup Exec does not monitor the low disk space thresholds for a disabled virtual disk.■ Online The virtual disk is available for use.■ Offline Backup Exec cannot access the virtual disk. You can check Offline to try to bring the storage array online.
Used capacity	<p>Displays the amount of raw capacity of all of the physical disks in the storage array that were used. Backup Exec calculates used capacity by subtracting available capacity from total capacity.</p>
Hardware name	<p>Displays the name that the storage array hardware or the vendor hardware provider assigns.</p>

Table Z-7 General properties for virtual disks on storage arrays *(continued)*

Item	Description
Hardware status	<p>Displays the hardware status.</p> <p>Values for the hardware status are as follows:</p> <ul style="list-style-type: none"> ■ OK The virtual disk is online. ■ Offline The virtual disk is offline. Backup Exec cannot access it. To bring the virtual disk online, refer to the vendor documentation and management software that are supplied with the storage array. ■ Failed The virtual disk has failed. Backup Exec cannot access it. To troubleshoot the issue, refer to the vendor documentation and management software that are supplied with the storage array. After the issue is resolved, the virtual disk is automatically brought online.
Hardware health	<p>Displays the hardware health.</p> <p>Values for the hardware health are as follows:</p> <ul style="list-style-type: none"> ■ OK The virtual disk is online. ■ Warning The virtual disk may fail or produce errors, but it is currently operational. ■ Critical The virtual disk has failed. Backup Exec cannot access it. To troubleshoot the issue, refer to the vendor documentation and management software that are supplied with the storage array. ■ Unspecified The virtual disk is in the process of configuration.

Table Z-7 General properties for virtual disks on storage arrays (*continued*)

Item	Description
Disk classification	<p>Displays the type of disk group that the virtual disk is on.</p> <p>Backup Exec creates only physical disk groups with a disk classification of Stripe with parity (RAID 5). If another disk classification appears, the disk group was created with a tool other than the Storage Provisioning Option.</p>
Maximum file size	<p>Displays the maximum size for each file that is contained on this virtual disk. Select either MB or GB. The file size can be from one MB to 4096 GB. The default size is four GB.</p> <p>If you create small but numerous files, performance may slow since the computer must still process each file. However, if you create large files, file system limitations can cause memory allocation problems or network issues. These issues can be a problem if you store files across a network.</p>
Maximum number of backup sets per file	<p>Displays the maximum number of backup sets to write to each virtual disk file. The maximum number can range from one to 8192. The default number is 100.</p> <p>If you specify fewer rather than more backup sets in a file, Backup Exec may be able to reclaim disk space faster. For example, you specify 100 backup sets per file. Backup Exec cannot reclaim any disk space until the overwrite protection period expires for all of the backup sets. If you specify one backup set per file, Backup Exec can reclaim disk space as soon as the overwrite protection period expires for that set.</p>
Allow x concurrent jobs for this device	<p>Displays the number of concurrent operations that you want to allow to this virtual disk. This number can range from one to 16. The default number is one.</p>

About hot spares in the Storage Provisioning Option

If a storage array that has automatic rebuild capability loses virtual disk redundancy, it uses a physical disk as a hot spare to regain redundancy. If your storage array does not support an automatic rebuild capability, you must use vendor tools to manually rebuild the virtual disks. Refer to your storage vendor documentation for more information.

Use the **Storage Array Configuration Wizard** to specify the physical disks that you want to use as hot spares.

Before you specify a hot spare, refer to the following best practices:

- Specify at least one hot spare for each enclosure. Although you can specify only one hot spare for all of the enclosures, consider the risk if more than one physical disk fails.
- Specify the physical disks that are in slot 0 in the enclosures as hot spares. Then, you can quickly identify which disk is a hot spare.
- Specify a hot spare that is at least the same size as the physical disk that it replaces. If the hot spare is smaller than the physical disk, the storage array cannot rebuild the virtual disk.

For more recommendations, refer to your storage array vendor's documentation.

See “[Changing a hot spare by using the Storage Array Configuration Wizard](#)” on page 2043.

See “[About using the Storage Array Configuration Wizard](#)” on page 2021.

See “[Configuring a storage array by using the Storage Array Configuration Wizard](#)” on page 2022.

Adding a hot spare by using the Storage Array Configuration Wizard

You can use the **Storage Array Configuration Wizard** to add a hot spare in a storage array. When you complete this wizard, it submits a utility job named Configure Storage Array. When the job completes successfully, the hot spare has been added.

To add a hot spare by using the Storage Array Configuration Wizard

- 1 On the **Tools** menu, click **Wizards>Storage Array Configuration Wizard**.
- 2 On the **Welcome** panel, select the storage array that contains the hot spare that you want to add, and then click **Next**.
- 3 On the **How do you want to configure the storage array?** panel, click **I only want to configure hot spares**, and then click **Next**.

- 4 In the **Available Physical Disks** list, select the physical disk that you want to use as a hot spare.
- 5 Click the bottom right arrow to move the selected physical disk to the **Hot Spares** list.
- 6 Click **Next**.
- 7 On the **Summary** panel, ensure that the **Hot Spare Count** is correct, and then click **Finish**.

See [“About hot spares in the Storage Provisioning Option”](#) on page 2042.

Changing a hot spare by using the Storage Array Configuration Wizard

You can use the **Storage Array Configuration Wizard** to select a different physical disk to use as a hot spare in a storage array. When you complete this wizard, it submits a utility job named Configure Storage Array. When the job completes successfully, the hot spare has been changed.

To change a hot spare by using the Storage Array Configuration Wizard

- 1 On the **Tools** menu, click **Wizards>Storage Array Configuration Wizard**.
- 2 On the **Welcome** panel, select the storage array that contains the hot spare that you want to change, and then click **Next**.
- 3 On the **How do you want to configure the storage array?** panel, click **I only want to configure hot spares**, and then click **Next**.
- 4 Do one of the following:

To designate a hot spare as an available physical disk

Do the following in the order listed:

- In the **Hot Spares** list, select the hot spare that you want to return to the **Available Physical Disks** list.
- Click the bottom left arrow to move the selected hot spare to the **Available Physical Disks** list.

To designate an available physical disk as a hot spare

Do the following in the order listed:

- In the **Available Physical Disks** list, select one or more physical disks that you want to use as hot spares.
- Click the bottom right arrow icon to move the selected physical disks to the **Hot Spares** list.

- 5 Click **Next**.
- 6 On the **Summary** panel, ensure that the **Hot Spare Count** is correct, and then click **Finish**.

See [“About hot spares in the Storage Provisioning Option”](#) on page 2042.

See [“Viewing storage array components in Backup Exec”](#) on page 2020.

Detecting a new storage array

Backup Exec periodically searches for new storage arrays or new physical disks. If Backup Exec does not find a new storage array or physical disk that you added, then you should run a refresh operation. If a refresh operation does not discover the new devices, then restart the Backup Exec services.

After you restart the services, the new storage array appears in the **Devices** view.

You must install the Storage Provisioning Option before Backup Exec can detect a new storage array.

To detect a new storage array

- 1 On the navigation bar, click **Devices**.
- 2 Expand a media server where you added a new storage array or a physical disk.
- 3 On the menu bar, click **View > Refresh**.

The new storage array should appear in the **Devices** view.

- 4 If the refresh does not discover the storage array, restart the Backup Exec services.

See [“Starting and stopping Backup Exec services”](#) on page 168.

See [“Troubleshooting the Storage Provisioning Option”](#) on page 2048.

Renaming a virtual disk or storage array

You can rename a virtual disk or storage array. Names cannot exceed 128 characters. You cannot change the hardware name.

If you use a vendor tool to configure the storage array, the hardware name that you assign it in the vendor tool appears. To change the name of the storage array, you must use the vendor-supplied tool.

To rename a virtual disk or storage array

- 1 On the navigation bar, click **Devices**.
- 2 Expand a media server that has an attached storage array, and then select the device that you want to rename.
- 3 Expand a storage array, and then select the device that you want to rename.
- 4 Do one of the following:
 - Right-click the device that you want to rename, click **Properties**, and then on the **General** tab, select the **Name** field.
 - In the task pane, under **General Tasks**, click **Rename**.
- 5 Type the new name of the device.
- 6 Click **OK**.

About identifying the physical disks of a virtual disk

Many storage array enclosures incorporate a set of physical disks that use small status lights to indicate the operational status of physical disks. The Storage Provisioning Option uses these lights with its **blink** feature to help you quickly identify the physical disks that comprise a virtual disk. When you select the **blink** feature for a virtual disk, the status lights on the physical disks blink.

Note: Storage array support for the blink features depends on storage array hardware support for the feature. Not all storage array hardware supports blinking. See your storage array hardware documentation for more information.

You can use the blink feature in different ways. You can use it to assist with:

- Moving virtual disks from one storage array to another.

You can use the blink feature when you want to move a virtual disk from one enclosure to another. If you have many enclosures, you can use the blink feature to identify the physical disks that comprise the virtual disk. Without it, determining the physical disks that comprise a virtual disk can be difficult.
- Identifying problematic physical disks.

When the Storage Provisioning Option generates an alert for a physical disk issue, you can use the blink feature to assist you in finding problematic physical disks.

When you use the blink feature, the following applies:

- The blink feature works on one virtual disk at a time.

You cannot use it to simultaneously identify the physical disks in multiple virtual disks.

See [“Identifying the physical disks of a virtual disk”](#) on page 2046.

Identifying the physical disks of a virtual disk

Use the following steps to identify the physical disks of a virtual disk.

See [“About identifying the physical disks of a virtual disk”](#) on page 2045.

To identify the physical disks of a virtual disk

- 1 On the navigation bar, click **Devices**
- 2 Expand a media server that has an attached storage array.
- 3 Expand a storage array.
- 4 Select a virtual disk.
- 5 In the task pane, under **Devices**, click **Blink**
- 6 To turn off the blink feature, in the task pane, under **Devices**, click **Unblink**.

About predicting disk usage in the Storage Provisioning Option

After you install the Storage Provisioning Option, Backup Exec can predict the usage of disk space on the storage arrays. You can configure Backup Exec to send an alert when it predicts that the amount of available disk space for all storage arrays is low. The alert provides information about whether the current disk space resources are sufficient, and can help you plan when to increase disk space.

Backup Exec gathers sample data for statistical analysis. For example, you can use the defaults to let Backup Exec gather data for a one-week period of 24-hour days. To ensure statistical accuracy, by default Backup Exec keeps 35 sample groups of data. Backup Exec examines the job history data for each sample group and uses the new data to recalculate the disk usage trends.

Backup Exec uses the sample data to estimate the rate at which future jobs will use space in the disk array. Backup Exec computes the statistical average for previous usage and determines any upward trends or downward trends. Backup Exec also calculates the amount of used disk space that becomes available as the overwrite protection periods expire for previous backup sets.

Backup Exec combines these estimates with the amount of available disk space on storage arrays. Backup Exec can then predict how much time remains before the disk space on all storage arrays is depleted.

The accuracy of the prediction is reduced if any of the following conditions occur:

- The amount of history data is insufficient for a statistically valid estimate.
- The history data shows substantial variance and lack of repeatability.

Backup Exec sends an alert to notify you when these conditions occur.

See [“Configuring an alert for low disk space on storage arrays”](#) on page 2047.

See [“Default options for Storage Provisioning Alert”](#) on page 2047.

Configuring an alert for low disk space on storage arrays

Backup Exec gathers disk usage information for all attached storage arrays. Through statistical analysis, Backup Exec then estimates how much time remains before the disk space on all storage arrays is depleted. You can also specify how many days before low disk space occurs that you want Backup Exec to send an alert.

See [“About predicting disk usage in the Storage Provisioning Option”](#) on page 2046.

To configure an alert for low disk space on storage arrays

- 1 On the **Tools** menu, click **Options**.
- 2 In the task pane, under **Settings**, click **Storage Provisioning Alert**.
- 3 Change the defaults as appropriate.
See [“Default options for Storage Provisioning Alert”](#) on page 2047.
- 4 Click **OK**.

Default options for Storage Provisioning Alert

Default options provide information for how Backup Exec gathers disk usage information.

See [“Configuring an alert for low disk space on storage arrays”](#) on page 2047.

Table Z-8 Default options for Storage Provisioning Alert

Item	Description
Send low disk space alert x days before predicted low disk space condition for all storage arrays	Displays when Backup Exec sends an alert before low disk space occurs. The default number is 30 days.

Table Z-8 Default options for Storage Provisioning Alert (*continued*)

Item	Description
Sample groups	<p>Displays the number of sample groups for Backup Exec to average to ensure a valid analysis.</p> <p>The default number is 35 sample groups.</p>
Samples per group	<p>Displays the number of samples per group. Each sample is a period of time during which Backup Exec gathers data. For example, if you specify the default of seven samples per group, Backup Exec gathers data for seven periods of time. The period of time is the sample interval that you specify. Backup Exec averages the samples in the group for statistical analysis.</p> <p>The default number is seven samples per group.</p>
Sample interval	<p>Displays the number of hours during which Backup Exec gathers data for a sample. For example, if you specify the default sample interval of 24 hours, Backup Exec gathers data for 24 hours.</p> <p>The default sample interval is 24 hours.</p>
Suppress hardware-related information alerts (Dell Backup-to-disk devices)	

Troubleshooting the Storage Provisioning Option

If problems occur with the Storage Provisioning Option or with storage array hardware, ensure the following:

- The operating system is supported. If the media server runs Windows Server 2003 R2, ensure that Service Pack 2 or the Microsoft Storage Manager for SANs management tool is installed.
- The vendor storage array and the vendor hardware provider are supported. You can find a list of compatible devices at the following URL:
<http://entsupport.symantec.com/umi/v-269-2>
- The storage array has power and is turned on.

- All lights and indicators on the storage array appear normal.
- The storage array is properly zoned if it is on a SAN.
- The cables are plugged into the correct ports.
- The Microsoft DiskRAID command line tool or the Microsoft Storage Manager for SANs management tool can detect and exercise the storage array hardware.
- The Disk Manager can detect the unmasked virtual disks.
- The refresh operation has run to detect new virtual disks.

If you installed the Storage Provisioning Option as an evaluation license, ensure that the license is still within the evaluation time period. When the evaluation time period expires, the option functions in a very limited mode.

Symantec Online Storage for Backup Exec

This appendix includes the following topics:

- [About Symantec Online Storage for Backup Exec](#)
- [Best practices for using Symantec Online Storage for Backup Exec](#)
- [Setting up Symantec Online Storage for Backup Exec](#)
- [About Symantec Online Storage folders](#)
- [About creating duplicate backup jobs for Symantec Online Storage for Backup Exec](#)
- [About managing Symantec Online Storage for Backup Exec jobs](#)
- [Erasing Symantec Online Storage for Backup Exec files](#)
- [Deleting Symantec Online Storage folders](#)
- [About restoring Symantec Online Storage for Backup Exec jobs](#)

About Symantec Online Storage for Backup Exec

Symantec Online Storage for Backup Exec provides online backup and restoration services as part of the Symantec Protection Network. The Symantec Protection Network offers Symantec technologies as online services. Its integration with Backup Exec means that you do not have to learn a new application to benefit from the security of online backups.

Symantec Online Storage for Backup Exec lets you back up your most critical data in Backup Exec and then send duplicate copies of the backups off-site. Your data

is securely stored on Symantec's servers where it is safe from hardware failure, malware, and natural disasters. Using Symantec Online Storage for Backup Exec can be an important part of your backup strategy.

Backing up your Backup Exec catalogs to the Symantec Protection Network protects your data even if you lose your entire Backup Exec media server. You can install Backup Exec and the Symantec Online Storage for Backup Exec Protection Agent on any supported computer to restore your data from the online catalogs.

Best practices for using Symantec Online Storage for Backup Exec

You can use Symantec Online Storage for Backup Exec to back up any data that you would normally back up with Backup Exec. Symantec recommends that you use Symantec Online Storage for Backup Exec for the small jobs that contain your most critical data however. Symantec Online Storage for Backup Exec is not intended to replace your local backup process. It can, however, help to provide security from natural disasters and hardware failure for your business' most important data.

Bandwidth capabilities may limit your ability to back up large Symantec Online Storage for Backup Exec jobs. You may find that it takes longer to run the same jobs over the Internet than it does to run them locally. Since you pay for the storage space that you use, backing up only your critical jobs is also the most cost-effective use of this service.

The most efficient way to use Symantec Online Storage is to run the same duplicate backup job periodically. You can create a policy and run the job according to a schedule. After the first time you run a job, Symantec Online Storage for Backup Exec examines the backup data in subsequent occurrences of that job. Any data that is unchanged from the previous occurrence is skipped. Subsequent backups include only the files that have changed since the last occurrence. This process reduces the amount of time and bandwidth that is required to run recurring backup jobs.

You may want to back up the following types of critical information with Symantec Online Storage for Backup Exec:

- Backup Exec catalogs
- Customer relationship management databases
- Employee or payroll information

Setting up Symantec Online Storage for Backup Exec

Before you can run a duplicate online backup job, you must set up Symantec Online Storage for Backup Exec. You sign up for the service and download the Symantec Online Storage for Backup Exec Protection Agent from the Symantec Protection Network Web site. Once you have completed these steps, you can create a Symantec Online Storage folder and run duplicate online backups.

You must sign up for Symantec Online Storage and download the Symantec Online Storage for Backup Exec Protection Agent before you create Symantec Online Storage folders.

Table AA-1 Setting up Symantec Online Storage for Backup Exec

Step	Description
Step 1	Sign up for Symantec Online Storage for Backup Exec on the Symantec Protection Network Web site. See “About signing up for Symantec Online Storage for Backup Exec” on page 2053.
Step 2	Download the Symantec Online Storage for Backup Exec Protection Agent. See “About downloading the Symantec Online Storage for Backup Exec Protection Agent” on page 2054.
Step 3	Create a Symantec Online Storage folder. See “Creating a Symantec Online Storage folder” on page 2054.

About signing up for Symantec Online Storage for Backup Exec

You sign up for Symantec Online Storage for Backup Exec on the Symantec Protection Network Web site. You must select a service plan that fits your needs.

To sign up for the Symantec Online Storage for Backup Exec, go to the following Web site:

<https://signup.spn.com>

For more information, refer to the Symantec Protection Network Web site's online Help.

See [“Setting up Symantec Online Storage for Backup Exec”](#) on page 2053.

About downloading the Symantec Online Storage for Backup Exec Protection Agent

Before you can use Symantec Online Storage for Backup Exec you must download the Symantec Online Storage for Backup Exec Protection Agent. The Symantec Online Storage for Backup Exec Protection Agent lets you create and configure Symantec Online Storage folders. Symantec Online Storage folders are the online storage devices that can be used as backup destinations like any other device in Backup Exec.

To download the Symantec Online Storage for Backup Exec Protection Agent go to the following Web site:

<http://www.spn.com>

Log on to your account and then follow the instructions that display for Symantec Online Storage for Backup Exec. For more information, refer to the Symantec Protection Network Web site's online Help.

See “[Setting up Symantec Online Storage for Backup Exec](#)” on page 2053.

About Symantec Online Storage folders

Symantec Online Storage folders are the backup destinations for online, duplicate backup jobs. You can create and configure multiple Symantec Online Storage folders for different online duplicate backup jobs. However, you can only run one Symantec Online Storage job at a time. After you download the Symantec Online Storage for Backup Exec Protection Agent and create a Symantec Online Storage folder, you can use it as a device. You can view Symantec Online Storage folders on the **Devices** tab.

See “[Setting up Symantec Online Storage for Backup Exec](#)” on page 2053.

See “[Creating a Symantec Online Storage folder](#)” on page 2054.

See “[Pausing a Symantec Online Storage folder](#)” on page 2056.

See “[Resuming a Symantec Online Storage folder](#)” on page 2056.

See “[Sharing an existing Symantec Online Storage folder](#)” on page 2057.

Creating a Symantec Online Storage folder

Symantec Online Storage folders are the backup destinations for online, duplicate backup jobs.

See “[About Symantec Online Storage folders](#)” on page 2054.

Note: You cannot create Symantec Online Storage folders until you sign up for Symantec Online Storage and download the Symantec Online Storage for Backup Exec Protection Agent.

See [“Setting up Symantec Online Storage for Backup Exec”](#) on page 2053.

Symantec Online Storage can never be part of a device pool, including the All Device pool. You must always specify the individual Symantec Online Storage folder to which you want to target a duplicate backup job. This feature helps ensure that you don’t accidentally send backup jobs to a Symantec Online Storage folder.

To create a Symantec Online Storage folder

- 1 On the navigation bar, click **Devices**.
- 2 In the task pane, under **Device Tasks**, click **Configure devices assistant**.
- 3 Click **Symantec Online Storage**.
If you have not set up Symantec Online Storage for Backup Exec, the **Configure Devices Assistant** displays **Symantec Protection Network** instead of **Symantec Online Storage**. You must sign up for the service and download the Symantec Online Storage for Backup Exec Protection Agent.
- 4 Type a name for the new folder.
- 5 Ensure that **Enable** is selected to make the folder available for online duplicate backup jobs.
- 6 Click **OK**.

Symantec Online Storage folder properties

You can create and configure multiple Symantec Online Storage folders for different online duplicate backup jobs.

See [“Creating a Symantec Online Storage folder”](#) on page 2054.

Table AA-2 Symantec Online Storage folder options

Item	Description
Name	Designates the name of the Symantec Online Storage folder. When you choose a name for a Symantec Online Storage folder you cannot change it. Symantec Online Storage folder names must not exceed 128 characters.
Pause	Pauses or unpauses the folder.

Table AA-2 Symantec Online Storage folder options (*continued*)

Item	Description
Enable	Enables the folder for use by Backup Exec. Clear this check box to disable the folder.

Pausing a Symantec Online Storage folder

You can pause a Symantec Online Storage folder. When you pause a Symantec Online Storage folder, duplicate backup jobs do not run on it. If a duplicate backup job is already running on a Symantec Online Storage folder when you pause it, the job completes. Any subsequent duplicate backup jobs do not run until the folder is resumed.

See [“Resuming a Symantec Online Storage folder”](#) on page 2056.

To pause a Symantec Online Storage folder

- 1 On the navigation bar, click **Devices**.
- 2 Expand the icon for the computer where the Symantec Online Storage folder is located.
- 3 Select the Symantec Online Storage folder you want to pause.
- 4 Under **General Tasks** in the task pane, select **Pause**.

Resuming a Symantec Online Storage folder

You can pause a Symantec Online Storage folder. When you pause a Symantec Online Storage folder, duplicate backup jobs do not run on it. To run a duplicate backup job to the Symantec Online Storage folder, you must resume it.

See [“Pausing a Symantec Online Storage folder”](#) on page 2056.

To resume a Symantec Online Storage folder

- 1 On the navigation bar, click **Devices**.
- 2 Expand the icon for the computer where the Symantec Online Storage folder is located.
- 3 Select the paused Symantec Online Storage folder that you want to resume.
- 4 Under **General Tasks** in the task pane, select **Pause**.

Sharing an existing Symantec Online Storage folder

If you have the Central Admin Server Option (CASO) or the SAN Shared Storage Option installed, you can share Symantec Online Storage folders between computers. Shared Symantec Online Storage folders are listed in the **Devices** view under each computer that can access them. Symantec Online Storage folder names are unique. You cannot have more than one folder with the same name for an account.

To share an existing Symantec Online Storage folder

- 1 On the computer on which you want to add the folder for sharing, on the navigation bar, click **Devices**.
- 2 In the task pane, under **Device Tasks**, click **Configure devices assistant**.
- 3 Click **Symantec Online Storage**.
- 4 Click **Add shared Symantec Online Storage**.
- 5 Type the name of the shared Symantec Online Storage folder that you want to add to this computer.
- 6 Click **OK**.

About creating duplicate backup jobs for Symantec Online Storage for Backup Exec

Backup jobs for Symantec Online Storage for Backup Exec must be created as duplicate jobs. You can duplicate existing backup sets or you can duplicate backup sets immediately following the scheduled job in which they are created.

See [“Duplicating backed up data”](#) on page 369.

See [“Creating duplicate backup jobs for Symantec Online Storage for Backup Exec”](#) on page 2058.

You can also create a duplicate backup set for Symantec Online Storage for Backup Exec as part of a template.

See [“About duplicate backup set templates”](#) on page 548.

See [“Adding a duplicate backup template to a policy”](#) on page 550.

After the first time you run a duplicate backup job, Symantec Online Storage for Backup Exec examines the backup data in subsequent occurrences of that job. Any data that is unchanged from the previous occurrence is skipped. Subsequent backups include only the files that have changed since the last occurrence. This process reduces the amount of time and bandwidth that is required to run recurring backup jobs.

You should consider the best practices when you create duplicate backup jobs for Symantec Online Storage for Backup Exec.

See “[Best practices for using Symantec Online Storage for Backup Exec](#)” on page 2052.

Creating duplicate backup jobs for Symantec Online Storage for Backup Exec

Backup jobs for Symantec Online Storage for Backup Exec must be created as duplicate jobs. You can duplicate existing backup sets or you can duplicate backup sets immediately following the scheduled job in which they are created.

See “[Duplicating backed up data](#)” on page 369.

To create duplicate backup jobs for Symantec Online Storage for Backup Exec

- 1 On the navigation bar, click **Job Setup**.
- 2 In the task pane, under **Backup Tasks**, click **New job to duplicate backup sets**.
- 3 Do one of the following:

To copy existing backup sets to another destination

Do the following in the order listed:

- Click **Duplicate existing backup sets**, and then click **OK**.
- Select the backup sets that you want to copy. For Oracle or DB2 jobs that were created with multiple data streams, under the instance name, select the date on which the backup set was created.

To duplicate the backup sets that are created when a scheduled backup job runs

Do the following in the order listed:

- Click **Duplicate backup sets following a job**, and then click **OK**.
- Select the scheduled backup job to be used as the source.

- 4 In the **Properties** pane, under **Destination**, click **Device and Media**.

5 Complete the following options.

Device	Select the Symantec Online Storage folder to which you want to copy the duplicate backup job.
Overwrite media	Select this option. You cannot append Symantec Online Storage for Backup Exec backup jobs to existing backup sets.

Complete any other options as necessary.

See [“Device and media options for backup jobs and templates”](#) on page 339.

6 In the **Properties** pane, under **Settings**, click **General**.

See [“General options for new duplicate backup set jobs”](#) on page 379.

7 In the **Properties** pane, under **Settings**, click **Advanced**.

8 In **Compression type**, click **None**.

Symantec Online Storage for Backup Exec duplicate backup jobs do not support hardware compression. If the original backup job used software compression, the Symantec Online Storage for Backup Exec job is compressed also.

Complete any other options as necessary.

See [“Advanced options for new duplicate backup set jobs”](#) on page 379.

9 In the **Properties** pane, under **Settings**, click **Network and Security**.

10 In **Encryption type**, click **Software**.

Symantec Online Storage for Backup Exec duplicate backup jobs must be encrypted. Symantec Online Storage does not support hardware encryption.

If the original backup job was encrypted, the encryption key for the original job is applied to the Symantec Online Storage for Backup Exec job.

Complete any other options as necessary.

See [“Network and Security backup options”](#) on page 407.

11 If you want Backup Exec to notify someone when the backup job completes, in the **Properties** pane, under **Settings**, click **Notification**.

See “[Notification options for jobs](#)” on page 682.

12 Do one of the following:

To duplicate data from a scheduled backup job Click **Run Now**.

To duplicate data from an existing backup set Click **Run Now**, or under **Frequency**, click **Schedule**.

About managing Symantec Online Storage for Backup Exec jobs

You cannot append Symantec Online Storage for Backup Exec backup jobs to existing backup sets. All Symantec Online Storage for Backup Exec duplicate backup jobs are either new jobs or they overwrite existing jobs with expired overwrite protection periods.

See “[About media overwrite protection](#)” on page 223.

To view your Symantec Online Storage for Backup Exec account information, log on to the Symantec Protection Network Web site. You can view how much data you have used, among other things. If you configured Backup Exec to verify your backup jobs, you can also view the results on the Web site.

You can access the Symantec Protection Network Web site at the following URL:

<https://www.spn.com>

Erasing Symantec Online Storage for Backup Exec files

You can erase Symantec Online Storage for Backup Exec files if you no longer need them. When you erase Symantec Online Storage for Backup Exec files, Backup Exec removes the data from both the Symantec Online Storage folder and the disk. It also removes the file references from the catalog. However, the file remains for use with future backup jobs.

Caution: You cannot restore the data that you erase. Before you erase files, be sure that you no longer need them.

To erase a Symantec Online Storage for Backup Exec file

- 1 On the navigation bar, click **Devices**.
- 2 Expand the icon for the computer where the Symantec Online Storage folder is located.
- 3 Select the Symantec Online Storage folder that contains the file that you want to erase.
- 4 In the **Results** pane, select the file that you want to erase.
- 5 In the task pane under **Media Tasks**, select **Erase media, quick**.
- 6 Click **Yes**, or if more than one file was selected, click **Yes to All**.
- 7 Complete the appropriate options.
See [“General options for utility jobs”](#) on page 482.
- 8 If you want a person or group to be notified when the job completes, in the **Properties** pane, under **Settings**, click **Notification**.
See [“Notification options for jobs”](#) on page 682.
- 9 Do one of the following:

To run the job now

Click **Run Now**.

To set scheduling options

Under **Frequency**, click **Schedule**.

See [“Scheduling jobs”](#) on page 356.

Deleting Symantec Online Storage folders

You can delete a Symantec Online Storage folder if you no longer want to store the folder or the data in it on the Symantec Protection Network.

You should erase any backup files that are contained within the Symantec Online Storage folder before you delete the folder. If you delete the folder without erasing the backup files, the files remain stored on the Symantec Protection Network, however you cannot view them inside Backup Exec.

See [“Erasing Symantec Online Storage for Backup Exec files”](#) on page 2060.

Note: If you deleted a Symantec Online Storage folder without erasing the files within it first, you can recreate the folder using the folder's original name. Then run an inventory job on the folder. The backup files within the Symantec Online Storage folder are visible after the inventory operation is complete. You can then follow the steps in this procedure to erase the files and delete the folder, if necessary.

To delete Symantec Online Storage folders

- 1 Erase any Symantec Online Storage backup files that reside in the folder you want to delete.
- 2 On the navigation bar, click **Devices**.
- 3 Expand the icon for the computer where the Symantec Online Storage folder is located.
- 4 Select the Symantec Online Storage folder that you want to delete.
- 5 In the task pane, under **General Tasks**, click **Delete**.
- 6 Click **Yes**.

About restoring Symantec Online Storage for Backup Exec jobs

You may find that it takes longer to restore jobs over the Internet using Symantec Online Storage for Backup Exec than it does to restore them locally. Symantec recommends that if you have to restore data you should first try to restore from the source of the duplicate backup before you use Symantec Online Storage for Backup Exec.

See “[Advanced options for restore jobs](#)” on page 613.

Accessibility and Backup Exec

This appendix includes the following topics:

- [About accessibility and Backup Exec](#)
- [About keyboard shortcuts in Backup Exec](#)
- [List box navigation in Backup Exec](#)
- [Tabbed dialog box navigation in Backup Exec](#)
- [About setting accessibility options](#)

About accessibility and Backup Exec

Symantec products meet federal accessibility requirements for software as defined in Section 508 of the Rehabilitation Act:

<http://www.access-board.gov/508.htm>

Symantec products are compatible with operating system accessibility settings as well as a variety of assistive technologies. All manuals also are provided as accessible PDF files, and the online help is provided as HTML displayed in a compliant viewer.

Keyboard shortcuts are available for all graphical user interface operations and menu items. Backup Exec uses standard operating system navigation keys and keyboard shortcuts. For its unique functions, Backup Exec uses its own keyboard shortcuts, which are documented.

See “[About keyboard shortcuts in Backup Exec](#)” on page 2064.

Items in the task pane that do not have keyboard shortcuts can be accessed by using the operating system's "mouse keys", which allow you to control the mouse through the numerical keyboard.

To see a table of the standard Microsoft navigation keys and keyboard shortcuts, select your version of Microsoft Windows from the table at:

<http://www.microsoft.com/enable/products/keyboard.aspx>

About keyboard shortcuts in Backup Exec

All menu items can be selected by using accelerator or mnemonic keyboard shortcuts. An accelerator is a key combination that provides shortcut access to a user interface function. A mnemonic (sometimes referred to as a "hot key") is a single-key equivalent (used in combination with the ALT key) for selecting user interface components such as menu items. The mnemonic "hot key" letter is underlined in the user interface.

Select secondary menu items by opening the main menu and using the UP or DOWN ARROW key until the desired item is highlighted. Press the RIGHT ARROW key to open a submenu, and ENTER to select your choice.

Keyboard shortcuts are not case-sensitive. Mnemonic keystrokes may be pressed either sequentially or simultaneously. All menu items have mnemonics, but not all menu items have accelerators.

Routine functions such as opening, saving, and printing files can be performed using the standard Microsoft keyboard shortcuts. Other menu items are unique to Backup Exec.

See "[Keyboard shortcuts unique to Backup Exec](#)" on page 2064.

See "[Keyboard shortcuts unique to Backup Exec](#)" on page 2064.

See "[Keyboard shortcuts unique to Backup Exec Desktop and Laptop Option Administration Console](#)" on page 2067.

See "[Keyboard shortcuts unique to Backup Exec Desktop and Laptop Option Desktop Agent](#)" on page 2069.

Keyboard shortcuts unique to Backup Exec

The following table lists the keyboard shortcuts unique to Backup Exec.

See "[About keyboard shortcuts in Backup Exec](#)" on page 2064.

Table AB-1 Keyboard shortcuts unique to Backup Exec

Backup Exec Accelerator	Backup Exec Mnemonic	Result
ALT	F	The File menu expands. From the File menu, you can create new jobs, devices and media, print selected items, view properties, or exit Backup Exec.
ALT	E	The Edit menu expands. From the Edit menu, you can rename, delete, copy, and select items. In addition, you can work with selection lists and search catalogs.
ALT	V	The View menu expands. From the View menu, you can change the information that displays on the screen. The options on the View menu change according to which item is selected on the navigation bar.
ALT	N	The Network menu expands. Use the Network menu to work with Backup Exec logon accounts, connect to media servers on the network, or to reconnect to a local media server.

Table AB-1 Keyboard shortcuts unique to Backup Exec (*continued*)

Backup Exec Accelerator	Backup Exec Mnemonic	Result
ALT	T	The Tools menu expands. The Tools menu provides many important options for working with Backup Exec, including starting and stopping services, using device and media operations, using Wizards, and setting default options.
ALT	W	The Window menu expands. Use the Window menu to move to a new window or view.
ALT	H	The Help menu expands. Use the Help menu to access Backup Exec documentation and various Symantec Web sites.

Keyboard shortcuts unique to Backup Exec Utility

The following table lists the shortcut keys in Backup Exec Utility.

See [“About keyboard shortcuts in Backup Exec”](#) on page 2064.

Table AB-2 Keyboard shortcuts unique to Backup Exec Utility

Backup Exec Accelerator	Backup Exec Mnemonic	Result
ALT	F	The File menu expands. From the File menu, you can create new media servers and media server groups, view properties, or exit Backup Exec Utility.
ALT	E	The Edit menu expands. From the Edit menu, you can rename, delete, and select items.
ALT	V	The View menu expands. From the View menu, you can change the information that displays on the screen.
ALT	H	The Help menu expands. Use the Help menu to access Backup Exec documentation and various Symantec Web sites.

Keyboard shortcuts unique to Backup Exec Desktop and Laptop Option Administration Console

The following table lists the shortcut keys in the Backup Exec Desktop and Laptop Option Administration Console.

See [“About keyboard shortcuts in Backup Exec”](#) on page 2064.

Table AB-3 Keyboard shortcuts unique to Backup Exec Desktop and Laptop Option Administration Console

Backup Exec Accelerator	Backup Exec Mnemonic	Result
ALT	F	The File menu expands. From the File menu, you can create new Profiles and Storage Locations, and add users.
ALT	E	The Edit menu expands. From the Edit menu, you can restore files, search for files to restore, manage alerts, and delete items.
ALT	V	The View menu expands. From the View menu, you can change the information that displays on the screen.
ALT	N	The Network menu expands. Use the Network menu to work with administrator accounts, connect to media servers on the network, or to reconnect to a local media server.

Table AB-3 Keyboard shortcuts unique to Backup Exec Desktop and Laptop Option Administration Console (*continued*)

Backup Exec Accelerator	Backup Exec Mnemonic	Result
ALT	T	The Tools menu expands. Use the Tools menu to set global excludes, access all DLO wizards, and manage service credentials.
ALT	W	The Window menu expands. Use the Window menu to move to a new window or view.
ALT	H	The Help menu expands. Use the Help menu to access Backup Exec documentation and various Symantec Web sites.

Keyboard shortcuts unique to Backup Exec Desktop and Laptop Option Desktop Agent

The following table lists the shortcut keys in the Backup Exec Desktop and Laptop Option Desktop Agent.

See [“About keyboard shortcuts in Backup Exec”](#) on page 2064.

Table AB-4 Keyboard shortcuts unique to Backup Exec Desktop and Laptop Option Desktop Agent

Backup Exec Accelerator	Backup Exec Mnemonic	Result
ALT	F	The File menu expands. From the File menu, you can minimize or exit the Desktop Agent.
ALT	V	The View menu expands. From the View menu, you can change the information that displays on the screen.
ALT	K	The Tasks menu expands. Use the Tasks menu to run a job or refresh the view.
ALT	T	The Tools menu expands. Use the Tools menu to reset dialog boxes and accounts.
ALT	H	The Help menu expands. Use the Help menu to access the online help for the Desktop Agent.

General keyboard navigation within the Backup Exec user interface

You can navigate and use Backup Exec with only the keyboard. In the user interface, the current active tree or table has a dark blue highlight, and the current active tab, radio button, or checkbox is enclosed within a rectangle formed by dotted lines. These areas are said to have focus and will respond to commands.

All Symantec user interfaces use the following keyboard navigation standards:

- The TAB key moves the focus to the next active area, field, or control, following a preset sequence. SHIFT+TAB moves the focus in the reverse direction through the sequence.
- CTRL+TAB exits any Console area that you internally navigate with the TAB key.
- UP and DOWN ARROW keys move focus up and down the items of a list.
- The ALT key in combination with the underlined mnemonic letter for a field or command button shifts the focus to that field or button.
- Either ENTER or the SPACEBAR activates your selection. For example, after pressing the TAB key to select Next in a wizard panel, press the SPACEBAR to display the next screen.
- SHIFT+F10 provides access to context menus.

Keyboard navigation within dialog boxes in Backup Exec

Dialog boxes contain groups of controls necessary to set options or settings for programs.

The following list contains some general rules about dialog box navigation:

- The TAB key moves focus between controls within the dialog box along a preset sequence.
- Controls displaying a mnemonic (an underlined letter) can be selected regardless of focus by typing ALT and the underlined letter.
- A dark border indicates the default command button. Press ENTER at any time to choose the button with a dark border.
- ESC chooses the Cancel button if one exists.
- SPACEBAR chooses a control you select with the TAB key.
- SPACEBAR changes the state of a checkbox that has focus. Typing a mnemonic (if one is available) will move the focus to the checkbox and change its state.
- Arrow keys move focus within radio buttons, list boxes, sliders, groups of option controls, or groups of page tabs.
- Items that cannot be changed are not visited by the TAB key sequence. Options that are unavailable are grayed-out and can neither be selected nor given focus.

While the controls described here are typically found in dialog boxes, they also can occur in other contexts. The same navigation standards will apply.

List box navigation in Backup Exec

List boxes display a column of available choices.

There are different kinds of list boxes with the following additional navigation conventions:

- Drop-down list boxes by default show only the selected item. A small button to the right of the control shows a downward-pointing arrow. Select the arrow to display more items from the list box. If there are more choices than can fit in the preset list box area, a slider appears along the side of the list box. Show or hide the list using ALT+DOWN ARROW, ALT+UP ARROW, or F4. The TAB key selects an item.
- Extended selection list boxes support selecting single items, blocks of items, or combinations of the two. After selecting an item, hold down CTRL+navigation keys to select or clear additional items or blocks of items.

Tabbed dialog box navigation in Backup Exec

Some dialog boxes use tabbed pages to subcategorize groups of many options. Each tabbed page contains different groups of controls. Use TAB to move the focus between tabbed pages within a dialog box. Typing the mnemonic for a tab also moves the focus to the tabbed page and displays its page of controls.

The following table lists keyboard navigation rules within tabbed dialog boxes:

Table AB-5 Keyboard navigation within tabbed dialog boxes

Keyboard input	Result
CTRL+PAGE DOWN or CTRL+TAB	Switches to the next tab and displays the page.
CTRL+ PAGE UP	Switches to the previous tab and displays the page.
RIGHT ARROW or LEFT ARROW	When the focus is on a tab selector, chooses the next or previous tab in the current row and displays the page.

About setting accessibility options

Symantec software responds to operating system accessibility settings.

Symantec products are compatible with Microsoft's accessibility utilities. In Windows operating systems, accessibility options involving keyboard

responsiveness, display contrast, alert sounds, and mouse operation can be set through the Control Panel.

Accessibility features are primarily for the English version. Localized versions of this product include support for keyboard (mouseless) navigation using accelerator keys and mnemonic keys.

For more information on setting accessibility options, see the Microsoft documentation.

Glossary

ADAMM (Advanced Device and Media Management)	A Backup Exec database that automates the tracking of media and storage devices and ensures that backups are written to the appropriate media.
administration console	The user interface that allows you to run Backup Exec operations. The user interface can be run from the media server or a remote computer.
agent	A component that allows workstations or other computers, for example, Microsoft SQL Server, to interact with the Backup Exec media server.
alert category	A group of one or more events that occur in Backup Exec and that can generate an alert. Examples of alert categories include Job Success, Install Warning, and Database Maintenance Failure.
alert source	A source that can generate an alert. Alert sources include jobs, media, devices, and systems.
alert type	The classification of an alert that lets you determine the severity of the alert. Alert types include Error, Warning, Information, and Attention Required.
alert	An event that usually requires some form of user interaction or acknowledgment.
allocated media	The media that are associated with a media set and that have current append and overwrite protection periods.
append period	The length of time that data can be added to the media. The append period starts when the first backup job is written to the media.
archive	A logical group of archived items that the Backup Exec Archiving Option creates. Archives are contained in vault store partitions. Each archived file system share has its own archive. Each archived Exchange mailbox has its own archive.
audit log	A running history of all actions that are performed in Backup Exec. An entry into the log is created each time an action that is configured to display in the audit log occurs.
Backup Exec service account	A user account that is configured for the Backup Exec system services. It contains a user name and password and provides the rights to log on as a service and act as a Backup Exec administrator.
backup method	An option that you select when you run a backup job to specify how Backup Exec sets each file's backup status. For example, depending on the method that you

select, Backup Exec may reset the archive bit or use modified time to determine if a file needs to be backed up.

backup set	The data that is selected from a single resource, such as a Microsoft Exchange dataset, and placed together on media when a backup job is run. Files selected from multiple resources create multiple backup sets.
backup strategy	The procedures that you implement for backing up your network. A good backup strategy requires minimal time to get a computer running in the event of a disaster.
backup-to-disk folder	A storage device that you create that enables you to back up data to a folder on a hard disk.
baseline	The first backup job to run in a synthetic backup policy. The baseline backup runs one time only and backs up all of the files on the selected resources. A full backup is assembled, or synthesized, from a baseline backup and the subsequent incremental backups that are also contained in a policy.
catalog	A database that tracks the contents of media created during a backup or archive operation. You can only restore information from fully cataloged media.
central administration server	A Backup Exec media server on which the Central Admin Server Option (CASO) is installed. In a CASO environment, the central administration server provides centralized administration and delegated job processing and load balancing functionality for Backup Exec media servers in your storage environment.
centralized catalog	A catalog location in the Central Admin Server Option. All of the files in the catalog are kept on the central administration server.
centralized restore	A process in which you can run and manage all restore operations from a central administration server. Centralized restore is only available with the Central Admin Server Option.
common encryption key	A type of encryption key that anyone can use to back up data using encryption and to restore encrypted data.
custom error-handling rule	An error-handling rule that you can define for a specific error code in an error category. When a job fails with the error code that is associated with the custom error-handling rule, the retry options and the final job disposition are applied to the job.
custom filter	A filter that you can define to display only the information that you specify in the Job Monitor.
device pool	A group of devices that can be used for Backup Exec operations. Jobs assigned to the device pool are run on the first available device.
device	A robotic library drive, stand-alone drive, backup-to-disk folder, virtual drive, removable storage drive, online storage folder, or other type of data storage that is supported by Backup Exec.

Differential - Back up changed files since last full	A backup method that includes all files that have been changed (based on archive bit) since the last full or incremental backup. This method does not affect any media rotation scheme because the archive bit is not reset.
Differential - Using modified time	A backup method that includes all files since the last full backup using the files' last modified date and time stamp.
distributed catalog	A catalog location in the Central Admin Server Option. Image files in the catalog are distributed to the central administration server from every managed media server. These distributed files are small because they do not contain the entire catalog. They contain only information about the backup set. The history files, which contain detailed information about the backup set, remain on the managed media server.
Duplicate Backup Sets template	A template that allows you to use a multi-stage backup strategy for backing up data to disk and then copying it to tape.
error-handling rule	A default or custom rule that sets retry options and the final job disposition for failed or canceled jobs. Retry options let you specify how often to retry a job if it fails and the time to wait between retry attempts. The final job disposition lets you place the job on hold until you can fix the error.
event	An action that occurs during a Backup Exec operation; for example, job cancellation.
Full - Back Up Files - Back up and delete the files	A backup method that backs up the selected data as a copy backup, verifies the media, and then deletes the data from the volume.
Full - Back Up Files - Copy the files	A backup method that includes all selected data. It does not affect any media rotation scheme because the archive bit is not reset.
Full - Back Up Files - Using archive bit	A backup method that backs up all of the files selected for backup and resets the archive bit to indicate that the files have been backed up.
Full - Back Up Files - Using modified time	A backup method that includes all of the files selected for backup and allows the use of incrementals and differentials using the modified date and time stamp.
granular restore	A restore of individual items from a backup for which the Granular Recovery Technology option was enabled.
GRT (Granular Recovery Technology)	A backup option that is available with some Backup Exec Agents. Granular Recovery Technology lets you restore individual items from database backups. A separate backup of the individual items is not required for you to recover one item.
Home view	A central place in Backup Exec from which you can access the features you use frequently. You can customize the Home view by adding or deleting items with Backup Exec data and links to features.

imported media	The media that are created by a product other than this installation of Backup Exec, but are in storage devices in the Backup Exec environment.
Incremental - Back up changed files since last full or incremental - Using archive bit (reset archive bit)	A backup method that backs up only the files that have changed (based on the archive bit) since the last full or incremental backup. It resets the archive bit to indicate that the files have been backed up
Incremental - Back up changed files since last full or incremental - Using modified time	A backup method that backs up all files that have changed since the last full or incremental backup using the files' last modified date and time stamp.
job delegation	A process by which jobs are distributed by a central administration server to available storage devices on managed media server. Job delegation is only available with the Central Admin Server Option.
job history	A report of what happened during the processing of the job (statistics, errors, and so on).
job log	A log that contains the results of a job. It is created when the job runs. You can review the job log for job errors and job details.
job	An operation that has been scheduled for processing by the media server. For example, if you make selections and submit a backup based on those selections, you have created a backup job. Jobs contain source or destination information, settings, and a schedule. Types of jobs include backup, restore, media rotation, resource discovery, report, test run, and utility jobs.
load balancing	<p>A feature in Backup Exec that automatically distributes jobs among any available storage devices.</p> <p>Also a feature of the Backup Exec Central Admin Server Option in which jobs are automatically distributed from a central administration server to multiple managed media servers for processing among the various storage devices.</p>
logon account	An account that stores the credentials of a Windows user account and that enables Backup Exec to manage user names and passwords. It can be used to browse resources or process jobs.
mailbox group	A group of user mailboxes to which you want to assign the same archive rules, retention categories, and vault stores in the Backup Exec Archiving Option. In Enterprise Vault, this is called a provisioning group.
managed media server	A media server that is managed by a central administration server. Managed media servers are responsible for the actual processing of backup and restore jobs in a Central Admin Server Option environment. Managed media servers are only available with the Backup Exec Central Admin Server Option.

media ID	A unique internal label that Backup Exec assigns to each media used in Backup Exec. The ID keeps statistics for each media. The media ID cannot be erased or changed.
media label	A label used to identify media. Backup Exec can assign the label automatically or you can specify a label prefix and number to be assigned for a type of media. If the media was first used in a library with a bar code reader, the media label will already have a bar code label.
media overwrite protection level	A global setting in Backup Exec that lets you specify whether to overwrite scratch, imported, or allocated media regardless of the media's overwrite protection period.
media rotation	A strategy that determines when media can be reused, or rotated back into use, by Backup Exec. Common examples of a media rotation strategy are Son, Father/Son, and Grandfather/Father/Son.
media server pool	A feature of the Backup Exec Central Admin Server Option that lets you group managed media servers in a pool to which you can restrict backup jobs.
media server	The computer on which Backup Exec is installed and where the Backup Exec services are running.
media set	A set of rules that apply to media that are associated with a media set. These rules specify append periods, overwrite protection periods, and vaulting periods.
media vault	A user-defined logical representation of the actual physical location of media, such as a special media room, a scratch bin, or an offsite location.
offhost backup	A feature of the Backup Exec Advanced Disk-based Backup Option that enables the backup operation to be processed on a Backup Exec media server instead of on the remote computer, or host computer. Moving the backup from the remote computer to a media server enables better backup performance and frees the remote computer as well.
offline media location	A node on the Media view that lists media that are onsite but are not in drives, slots, or media vaults. Media are automatically moved to the offline media location if you use Backup Exec to remove media from a device or slot.
online media location	A node on the Media view that lists media that reside in a storage device, robotic library slot, or backup-to-disk folder.
overwrite protection period	The length of time that data is retained on a specific media before being overwritten (unless the media is erased, formatted, moved to scratch media, or if the media overwrite protection level is set to None). The overwrite protection period is measured from the last time data was appended to the media.
policy	A method for managing backup jobs and strategies. Policies contain settings for jobs.

preferred server configuration	A collection of one or more servers and sites that you select as preferred backup sources. Preferred server configurations take priority as backup sources in instances where data is replicated between multiple servers.
primary database server	The server on which the shared Advanced Device and Media Management (ADAMM) database and the shared catalog database reside when the Backup Exec SAN Shared Storage Option is installed.
recyclable media	Media that is assigned to a media set but has expired data overwrite protection periods.
remote administrator	The Backup Exec user interface (Administration Console) that is run on remote computers.
remote agent	A Backup Exec system service that runs on Microsoft Windows computers or NetWare remote servers and workstations and allows remote backup and restore of those computers and provides increased backup throughput.
replicated catalog	A catalog location in the Central Admin Server Option. All of the files in the catalog are replicated from the managed media server to the central administration server.
resource discovery	A Backup Exec operation that allows detection of new backup resources within a Windows domain.
resource	Data files and databases, such as Windows shares and Microsoft SQL databases, that can be selected for backup.
restricted encryption key	A type of encryption key that anyone can use to back up data using encryption. Only the key owner or a user with knowledge of the pass phrase can restore data that was encrypted with a restricted encryption key.
retention category	A setting in the Backup Exec Archiving Option that lets you specify the period of time for which you want to keep items in the archives. You can name a retention category to make it easier to search for and retrieve archived items.
retired media	Media that has been taken out of service, usually because of an excessive number of errors. Media that is retired is available for restore jobs but not for backup jobs. Media must be retired before it can be deleted. If you want to use media that has been deleted, Backup Exec will recognize it as imported media. It must be cataloged before you can restore from it.
scratch media	Media that are not associated with a media set and that can be overwritten. Scratch media includes new or blank media, erased media, and media moved from another group.
selection list	The data selected to be backed up or restored. Selection lists can be saved and used for multiple jobs.
simulated tape library	A tape library that emulates an Advanced Intelligent Tape (AIT) media type and has the AIT media type label. A simulated tape library is created by the Tape Library Simulator.

Symantec Online Storage folder	A storage device that you create to back up data to the Symantec Protection Network.
Symantec Online Storage for Backup Exec	An optional Backup Exec component that provides online backup and restoration services as part of the Symantec Protection Network.
Symantec Protection Network	Symantec's software-as-a-service provider. The Symantec Protection Network offers Symantec technologies as online services.
synthetic backup	A feature of the Advanced Disk-based Backup Option that enables a full backup to be assembled, or synthesized, from a baseline and subsequent incremental backups.
Tape Library Simulator	A utility that lets you create a virtual device on a hard disk or on any mounted volume on a computer on which the Backup Exec Remote Media Agent for Linux Servers is installed. The virtual device that is created is called a simulated tape library.
template rule	A method of setting up relationships between templates in a policy.
template	A required element of a policy that defines how and when Backup Exec processes a job. Templates specify the device, settings, and schedule options to be used for the job. Each policy must contain at least one template.
true image restore	A feature of the Advanced Disk-based Backup Option that enables Backup Exec to restore the contents of directories to what they were at the time of any full or incremental backup. Restore selections are made from a view of the directories as they existed at the time of the particular backup. Files that were deleted before the time of the backup are not restored. In true image restore, only the correct versions of files are restored from the appropriate full or incremental backups that contain them. Previous versions are not unnecessarily restored and then overwritten.
UMI (Unique Message Identifier)	A unique code that is associated with an error reported in the job log, or on some alerts. These codes contain hyperlinks that you can click to go to the Symantec Technical Support Web site. You can access technical notes and troubleshooting tips that are related to a specific error.
vault store partition	The physical location on a disk where archived items that the Backup Exec Archiving Option creates are stored. Backup Exec creates one vault store partition in each vault store by default. As the data in a vault store grows, you can create more vault store partitions to provide additional capacity.
vault store	A disk-based container for the archived data that the Backup Exec Archiving Option archives from one server.
virtual disk	A logical disk that you configure on a storage array to provide storage to the media server.

Working Set - Back up files - Changed today A backup method that backs up all files that were created or modified today.

Working Set - Back up files - Last accessed in (x) days A backup methods that backs up data that has been accessed in a specified number of days. If you select this backup method, you can then indicate the number of days in the Files accessed in x days field.

Index

A

- accelerator
 - defined 2064
- accessibility
 - dialog boxes 2071
 - keyboard navigation 2070
 - keyboard shortcuts 2064
 - overview 2063
 - settings 2072
- active alerts
 - defined 645
 - responding to 654
 - viewing 645
- Active Alerts by Media Server Report 730
- Active Alerts report 729
- Active Directory
 - backing up in Exchange 1105
 - for Automated User Assignments in DLO 1684
 - for connection policies 1778
- Active Directory Domains
 - adding a domain 288
 - deleting a domain 289
- Active Directory Recovery Agent
 - about 882
 - about restoring individual objects 888
 - Granular Recovery Technology (GRT)
 - overview 883
 - installing 881
 - passwords 889
 - recreating purged objects 892
 - requirements 880
 - tombstones 888
 - Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from Active Directory backups
 - option 884, 887
- Active File Exclusion 1105
 - Exchange data 1073
- active jobs
 - canceling 562
 - placing scheduled occurrences on hold 563
- active jobs (*continued*)
 - removing hold 563
 - viewing properties 558, 561
- add Remote Media Agent 1972
- add users in DLO
 - CSV file to 1699
- administration console
 - Desktop Agent Console 1763
 - Desktop and Laptop Option (DLO) 1638
 - overview 91
 - role in backup process 68
- Advanced Device and Media Management (ADAMM)
 - database overview 1992
 - device management overview 441
- Advanced Disk-based Backup Option
 - baseline
 - setting 899
 - best practices for offhost backup 923
 - host computer
 - defined 919
 - offhost backup 928
 - List Snapshot Providers option 925
 - offhost backup for Exchange Server with Granular Recovery Technology 929
 - offhost backup overview 919
 - offhost backup requirements 921
 - offhost backup snapshot provider
 - choosing 927
 - synthetic backup
 - creating 906
 - option to enable 904
 - policy. *See* example
 - template requirements 901
 - synthetic backup template rules 910
 - example 906
 - transportable snapshots
 - defined 920
 - true image restore
 - enabling 916
 - option to enable 904
 - overview 912
 - requirements 915

- Advanced Disk-based Backup Option *(continued)*
 - VSW FlashSnap option
 - using with offhost backup 922
- Advanced File Selection options 298
- Advanced Open File Option
 - backup jobs 948
 - cache file location 946
 - changing default settings for Symantec VSP 945
 - databases
 - backing up 939
 - default options 949
 - setting 943
 - encrypted files
 - backing up 941
 - installing 941
 - installing in an Active Directory network 140
 - installing to a remote computer in the backup selections list 140
 - installing to remote computers 131, 134
 - installing using a command script 148
 - installing using the command line 942
 - job log 951
 - overview 937
 - requirements 941
 - Snap Start
 - overview 944
 - Snap Start for VSW volumes 945
 - snapshot providers 940
 - Symantec Volume Snapshot Provider 949
 - Symantec VSP cache file size 947
 - uninstalling using a command script 149
- Advanced options for backup 348
- Agent for Microsoft Hyper-V
 - backing up 1175
 - backup options 1175
 - backup selections 1173
 - default options 1186
 - enabling Granular Recovery Technology (GRT) 1175
 - highly available virtual machines 1188
 - installation overview 1170
 - overview 1169
 - requirements 1171
 - restoring data to virtual server 1182
 - restoring virtual machine to different host 1184
 - selecting resources for restore 1180
- Agent for Microsoft SharePoint
 - about 1190
 - about restoring SharePoint 2003 resources 1237
 - Agent for Microsoft SharePoint *(continued)*
 - about restoring SharePoint 2007 resources 1221
 - about restoring SharePoint 2010 resources 1203
 - about restoring SharePoint Foundation 2010 resources 1203
 - about restoring SharePoint Services 3.0 resources 1221
 - about SharePoint Portal Server 2003 and Windows SharePoint Services 2.0 1235
 - adding a server farm 1191, 1219
 - automatic protection for SharePoint 1194
 - backing up a Microsoft Office SharePoint 2007 Server 1219
 - backing up a Microsoft Office SharePoint 2010 Server 1202
 - backing up a Windows SharePoint Foundation 2010 farm 1202
 - backing up a Windows SharePoint Services 3.0 farm 1219
 - backing up individual SharePoint 2007 web applications 1220
 - backing up individual SharePoint 2010 web applications 1203
 - backing up SharePoint Portal Server 2003 resources 1236
 - backup options 1246
 - changing the default name of a farm 1195
 - default options 1197
 - deleting a farm 1195
 - disabling or enabling communication between Web servers and Backup Exec 1196
 - installing 1191
 - overview 1190
 - redirecting a restore for SharePoint Portal Server 2003 1241
 - redirecting a restore for SharePoint Server 2007 1227
 - redirecting a restore for SharePoint Server 2010 1208
 - redirecting individual SharePoint 2003 objects to a file path 1243
 - redirecting individual SharePoint 2007 objects to a file path 1228
 - redirecting individual SharePoint 2010 objects to a file path 1210
 - redirecting restore jobs for SharePoint 2003 document library data 1242
 - redirecting the restore of a SharePoint 2003 database to an alternate SQL instance 1244

Agent for Microsoft SharePoint *(continued)*

- redirecting the restore of a SharePoint 2007 database to an alternate SQL instance 1230
 - redirecting the restore of a SharePoint 2010 database to an alternate SQL instance 1213
 - redirecting the restore of SharePoint 2007 web applications 1229
 - redirecting the restore of SharePoint 2010 web applications 1211
 - redirection options 1250
 - requirements 1190
 - restore options 1247
 - restoring a Microsoft Office SharePoint Server 2007 Shared Services Provider 1224
 - restoring a SharePoint Server 2007 Web application 1225
 - restoring a SharePoint Server 2010 Web application 1207
 - restoring individual SharePoint 2003 objects 1238
 - restoring individual SharePoint 2007 objects 1223
 - restoring individual SharePoint 2010 objects 1205
 - restoring SharePoint 2003 document libraries 1240
 - restoring SharePoint 2003 documents from document library backups 1241
 - restoring SharePoint 2003 resources 1238
 - restoring SharePoint 2007 resources 1222
 - restoring SharePoint 2010 resources 1204
 - restoring SharePoint Foundation 2010 resources 1204
 - restoring SharePoint Services 3.0 resources 1222
 - selecting SharePoint Portal Server 2003 resources for backup 1236
 - server farms options 1193
 - setting default options for SharePoint Server 2003 and later 1196
 - system requirements 1190
 - using with SharePoint Server 2007 and Windows SharePoint Services 3.0 1218
 - using with SharePoint Server 2010 and Windows SharePoint Foundation 2010 1200
- Agent for VMware**
- adding VMware vCenter and ESX servers 1387
 - backing up resources 1390
 - backup defaults 1406

Agent for VMware *(continued)*

- backup methods 1388
 - components 1386
 - delete existing virtual machines 1402
 - deleting VMware vCenter and ESX servers 1388
 - full system recovery 1400
 - Granular Recovery Technology (GRT)
 - overview 1396
 - requirements 1386
 - Granular Recovery Technology (GRT), setting for backup 1394
 - installing 1387
 - overview 1386
 - redirecting restores 1404
 - requirements 1386
 - restore defaults 1406
 - restore overview 1400
 - restoring resources 1401
 - selecting individual files and folders for restore 1400
 - selecting network for redirected restore 1406
 - selecting storage location for redirected restore 1405
 - selecting transport method for VMDK file 1393, 1407
 - transport mode priority 1403
 - turn on virtual machine after restore 1402
 - VSS Provider 1398
- Agents**
- Backup Exec
 - list of agents with descriptions 80
 - trial version 166
 - upgrading on remote computers 180
 - alert history
 - defined 645
 - viewing 645
 - Alert History report 730
 - alert notification
 - printers 675, 677
 - alert notification in DLO
 - printers 1731
 - alerts
 - alert types 644
 - assigning recipients 680
 - defined 644
 - defined for DLO 1721
 - deleting in DLO 1726
 - displaying in DLO 1724
 - filters 645, 648

alerts *(continued)*

- grooming in DLO 1720
- handling with SAP Agent 1366
- managing in DLO 1725
- monitoring in DLO 1720
- properties 651
- removing recipients 681
- responding to 654
- viewing 648
- viewing job log 653
- views 645

all media

- defined 220

All Virtual Disks device pool

- description 2029

allocated media

- overwriting 232

append period

- defined 223, 231
- setting for media set 231

append to media option 341, 375

Application Event Log report 732

application settings

- Backup Exec Retrieve 874
- catalogs 601

archive bit

- defined 273

Archive Job Success Rate 768

archive logging

- Lotus Domino 1074
- recovery of Lotus Domino 1090

Archive Selections by Archive Rules and Retention

- Categories 768

archive settings

- overview 1455

archives

- deleting 1455
- deleting items with expired retention periods 1454
- editing properties 1454
- overview 1453

Archiving Option

- allowing archiving from tapes 1441
- archives, overview 1453
- archiving from encrypted data 1440, 1495
- archiving from recent backups 1495
- archiving from tape devices 1495
- arranging archiving rules for file system selections 1441

Archiving Option *(continued)*

- arranging mailbox groups 1443
- arranging mailbox groups for provisioning 1463
- assigning a vault store 1439
- assigning settings to file system selections 1437
- assigning system mailbox 1443
- audit log entries 1432
- backing up components 1477, 1481
- backing up from remote media server 1491
- best practices 1433
- configuring archiving rules for file system selections 1443
- configuring options for Exchange mailboxes 1443
- creating a job 1434
- creating a vault store 1446
- creating a vault store partition 1451
- data not archived 1431
- deleting a vault store 1450
- deleting archives 1455
- deleting data after archiving 1477
- deleting items after archiving 1446
- deleting items after vault store backup 1446
- deleting items from archives 1474
- deleting items with expired retention periods 1454, 1493, 1495
- Disabled job status 1429
- disabling backup mode 1480
- displaying Backup Exec Retrieve links to end users 1432
- editing archive properties 1454
- editing default retention category 1495
- editing retention categories 1458
- editing vault store partition properties 1452
- editing vault store properties 1447
- enabling single instance storage 1495
- finding recent data to archive 1440
- granting permissions on Exchange Server 1419
- how it works 1429
- including and excluding file system selections 1460
- installing 1428
- installing Enterprise Vault 1429
- item deletion mode 1448
- items not supported 1415
- managing mailbox groups 1463, 1495
- overview 1414
- overview of archive settings 1455
- overview of components 1477

Archiving Option *(continued)*

- overview of mailbox groups 1461
 - overview of retention categories 1457
 - redirecting restores for Directory database 1486
 - redirecting restores of all components 1485
 - redirecting restores of components 1490
 - redirecting restores of Exchange items 1473
 - redirecting restores with Backup Exec Utility 1485-1486
 - reinstalling 1428
 - reports 1500
 - requirements 1415
 - restoring components 1483
 - restoring Exchange data from archives 1471
 - restoring file system data from archives 1472
 - restoring from remote media server 1491
 - restoring items from archives 1468
 - retaining directory structure during restore 1470
 - running Backup Exec Utility 1491
 - running consistency checks on databases 1480
 - running Enterprise Vault services 1422
 - searching archives for data 1464
 - searching archives for data to restore 1470
 - selecting administrative shares 1436
 - selecting items to restore from archives 1470
 - selection file system shares and folders 1436
 - setting backup job defaults 1492
 - setting job defaults 1495
 - setting restore options for databases 1484
 - setting rules for archiving mailbox groups 1462
 - skipping and overwriting items during restore 1470
 - specifying archive settings 1456
 - specifying retention period 1459
 - synchronizing permissions and settings 1493, 1495
 - troubleshooting 1499
 - uninstalling 1428
 - updating SQL Server name 1485
 - using Backup Exec Retrieve 1432
 - vault store partitions, overview 1451
 - vault stores, overview 1445
 - viewing the Enterprise Vault event log 1500
 - viewing vault store partition status 1452
 - viewing vault store status 1448
- ARCserve media**
- about restoring data from 624
 - restoring data from 625

- ASR files in IDR 1811
- audit log
 - about 205
 - Archiving Option entries 1432
 - configuring 206
 - for media operations 241
 - removing entries 208
 - saving to a file 208
 - viewing 206
- Audit Log report 732
- auto-inventory media after import job completes 490
- AUTOEXEC.NCF file
 - Remote Agent for NetWare Servers 1934
- Automated User Assignment
 - creating 1683
 - defined 1604, 1756
 - deleting 1687
 - modifying 1686
 - priority
 - changing 1686
 - properties
 - viewing 1686
- Automated User Assignment in DLO 1683
 - using Active Directory 1684
- automatic exclusion of files during volume level backups 1073
- automatic exclusion of SQL data during volume level backups 1280
- automatic updates
 - about scheduling 171
 - scheduling 172
- availability windows
 - setting 307
 - setting default 307
- availability windows
 - about 306

B

- back up and delete the files method
 - freeing disk space 367
 - using for a backup 368
- backing up
 - Microsoft clusters
 - local disks 837
 - shared disks 838
 - VERITAS clusters
 - database files 849
 - local disks 848
 - shared disks 848

BACKINT

- handling alerts 1366
- overview 1365
- using with CCMS console 1371

backup

- how to prepare 270
- overview 329
- using Remote Media Agent 1978

Backup Exec

- accessibility 2063
 - dialog boxes 2071
 - keyboard shortcuts 2064
- application protection agents 82
- client protection agents 86
- Desktop and Laptop ports 413
- installing
 - command line installation 153
 - silent mode installation 153
 - uninstalling 169
 - upgrading 180
 - using Repair option 167
- listening ports 412
- media server components 80
- media server storage options 87
- options 80
- overview
 - how it works 68
 - new agent and option features 75
 - new features 70
- ports 411
- repairing 167
- server protection agents 81
- system requirements 114
- upgrading
 - overview 180
 - using with Symantec Endpoint Protection 408
- virtual machine agents 85

Backup Exec 2010

- described 64

Backup Exec Archiving Site

- backing up 1477

Backup Exec diagnostic application

- diagnostic file
 - generating 805
 - generating using command line 806
- options 805
- overview 804

Backup Exec Environment Check 104**Backup Exec License Assessment Tool 176****Backup Exec media server in SAN 1992****Backup Exec Migrator**

- about 1041
- about retrieving Enterprise Vault data 1058
- about staged migrations 1046
- about the Backup Exec restore browse view 1057
- Backup Exec media server
 - working with 1051

best practices 1061**communicating with Enterprise Vault 1054****configuring 1049****data migration process 1046****Enterprise Vault retention periods 1048****events****about 1046****how it works 1042****log file location 1048****logs****about 1047****migrated files****about deleting 1048****Migrator for Enterprise Vault options 1053****requirements 1041****retrieving Enterprise Vault data 1058****troubleshooting 1062****Backup Exec Retrieve****application settings 874****description 1632****displaying links to end users 1432****end users, requirements for using 868****requirements for installing 866****retrieving files 1790****troubleshooting 876****using with Archiving Option 1432****using with DLO 1632****Backup Exec services****Backup Exec Services Manager dialog box 168****stopping and starting 168****Backup Exec Utility****redirecting restores for Archiving Option****Directory database 1486****redirecting restores of Archiving Option 1485****running for Archiving Option 1491****backup job****Advanced options 348****Agent for Microsoft Hyper-V 1175****choices for creating 329****copying to another server 554****creating manually 332**

- backup job *(continued)*
 - creating using the wizard 331
 - deduplication 1595
 - Device and Media options 339
 - pre/post commands 352, 399
 - required user rights 331
 - scheduling 356
 - selecting devices and data 280
 - Selections options 336
 - Symantec Online Storage for Backup Exec 2058
 - tasks to do before 330
- Backup Job Properties dialog box
 - SQL Agent 1273
- Backup Job Success Rate 733
- backup methods
 - selecting 343
 - setting default 391
 - using modified time 279
 - VMware resources 1388
- backup network
 - configuring 404
 - for a backup job 407
 - overview 402
 - setting up 404
- backup options
 - Configure desktop and laptop 1638
- Backup Recommendations 733
- backup selections
 - adding in DLO 1660
 - changing order 338
 - macros in DLO 1667
 - using fully qualified computer names 282
- backup selections list
 - Computer Name 282
 - Domains 287
 - Favorite Resources 284
 - User-defined Selections 290
- Backup Set Details by Resource report 734
- Backup Sets by Media Set report 735
- Backup Size by Resource report 736
- backup strategies
 - amount of data to be backed up 271
 - choosing resources to back up 272
 - defined 270
 - frequency of backups 271
 - how to choose 270
 - increase throughput with Remote Agent for Windows Systems 1946
 - length of data retention 272
- backup strategies *(continued)*
 - multiple resources per job 273
 - one job per resource 273
 - protecting against viruses 272
- backup types
 - about 274
 - Back up and delete the files 275
 - copy 275
 - daily 277
 - differential 276
 - full 275
 - incremental 276
 - working set 277
- Backup Wizard
 - configuring to launch from the Backup
 - button 332
 - launching 331
 - preventing from launching from the Backup
 - button 332
- backup-to-disk file
 - defined 496
 - deleting 509
 - erasing 510
 - recreating a deleted file 509
 - renaming 508
- backup-to-disk folder
 - advanced properties 501
 - auto detect settings 502
 - Backup-to-Disk Wizard 498
 - buffered reads 502
 - changing the path 506
 - concurrent jobs 496, 505
 - creating 499
 - defined 496
 - deleting 507
 - editing application settings 499, 505
 - low disk space threshold 501
 - maximum number of backup sets 505
 - maximum size 504
 - allocate at 504
 - overview 496
 - recommendations for using with Granular Recovery Technology 511
 - recreating 507
 - requirements 497
 - sharing 506
 - using in IDR 496

- bandwidth settings
 - DLO
 - for users in 1643
 - bar code labels
 - and media ID 242
 - default 244
 - mixed media libraries 245
 - overview 244
 - robotic library support 244
 - bar code rules
 - deleting 246
 - editing 245
 - enabling 472
 - setting up 245
 - baseline
 - setting for synthetic backup 899
 - bediag.fax file 1944
 - bediag.nlm utility
 - saving configuration information 1943
 - beoper group
 - creating 1880
 - Remote Agent for Linux or UNIX Servers,
 - about 1880
 - besernum.xml file
 - importing license keys with 121, 124
 - BESTART command
 - to start Remote Agent for NetWare Servers 1934
 - BESTOP command
 - to stop Remote Agent for NetWare Servers 1934
 - biparam.ini
 - options 1369
 - overview 1363
 - specifying Backup Exec parameters 1368
 - BKUPEXCDLO MSDE database instance 1738
 - blackout window
 - setting in DLO 1705
 - blink feature
 - about 2045
 - how to identify the physical disks 2046
 - block size
 - setting for devices 460
 - boot managers
 - restoring in IDR 1828
 - bootable media for IDR
 - CD image
 - creating 1817
 - comparing types 1810
 - tape image
 - creating 1819
 - bootable media for IDR (*continued*)
 - types of media 1809
 - BRRESTORE
 - restoring data with 1372
 - BRTOOLS
 - using with SAP Agent 1371
 - buffer count
 - setting for devices 462
 - buffer size
 - setting for devices 462
 - byte count
 - incorrect 796
- ## C
- cache file location
 - for AOFO 946
 - calendar
 - viewing the job workload from 588
 - CASO
 - about configuring 1529
 - alerts
 - configuring 1538
 - alias for managed media server 1521
 - alias for SQL Express 1520
 - Backup Exec Utility
 - running 1527
 - catalog location
 - changing 1542
 - displayed 1559
 - overview 1541
 - central administration server 1520
 - installing 1512
 - pausing storage devices from 1562
 - setting for a managed media server 1526
 - centralized catalog
 - overview 1542
 - centralized restore
 - multiple storage devices 1553
 - overview 1552
 - communication status
 - none 1536
 - communications
 - disabling 1562
 - device and media data 1509
 - distributed catalog
 - overview 1541
 - duplicate backup data job requirements 1551
 - features in 1507
 - installing across a firewall 1518

CASO (continued)

- job delegation 1504, 1544
- job history options
 - setting 1536
- managed media server
 - device and media data. *See* choosing the location of
 - installing 1513
 - installing for SAN Shared Storage Option 1515
 - job history options 1536
 - job log options 1536
 - pausing 1561
 - stalled 1536
 - status messages 1558
 - viewing properties 1564
- media server
 - changing to a managed media server 1526
- media server pool
 - advantages of 1545
 - apply settings to all servers in a pool 1550
 - deleting 1548
 - filter data for 1545
 - overview 1545
 - removing a managed media server 1548
 - renaming 1548
- monitor jobs on local managed media server 1532
- network interface cards
 - using any available 1540
- network traffic
 - reducing 1530
- notification
 - configuring 1540
- overview 1504
- port numbers for SQL instance 1520
- recovered jobs 1535–1536
- recovering failed jobs 1538, 1560
- replicated catalog
 - overview 1542
- requirements 1508
- selection list
 - restrict backup of 1546
- setting default for managed media servers 1531
- stopping and starting Backup Exec services 1563
- synthetic backup job requirements 1551
- time differences between servers 1533
- uninstalling Backup Exec from central administration server 1528

CASO (continued)

- uninstalling Backup Exec from managed media server 1528
- upgrading 1521
- catalog
 - changing application settings 601
 - levels 601, 603
 - media in drive 248
 - media with encrypted backup sets 423
 - restore jobs 600
 - searching 631
- catalog database
 - in SAN Shared Storage Option 1997
- catalog operation errors
 - DLT tape drive hangs 794
- catalogs
 - application settings 601
- CCMS console
 - using with SAP Agent 1371
- cell phone
 - notification 662
- centralized catalog
 - in CASO 1542
- centralized restore
 - best practices 1555
- CHECKCATALOG utility 1261
- CHECKDB utility 1261
- CHECKFILEGROUP utility 1261
- checkpoint restart on Microsoft cluster failover
 - enabling or disabling 824
 - overview 822
- circular logging
 - Exchange Agent
 - reviewing in 1106
 - Lotus Domino 1074
 - recovery of Lotus Domino server 1090
- cleaning slot
 - defining 472
- cleaning slots
 - defining for robotic libraries 471
- Cleaning tab for device properties 465
- client-side deduplication
 - adding Remote Agent for Deduplication 1592
 - how to configure 1591
 - overview 1590
- cluster
 - cluster shared volumes 839
- Cluster Failover error-handling rule 591, 595

- cluster shared volumes 839
 - quorum names 839
- clusters
 - backing up SAP database 1377
 - Desktop and Laptop Option 1738
 - disaster recovery
 - entire cluster manually 853
 - nodes using IDR 851
 - using IDR to prepare 851
 - installation of Backup Exec on a VERITAS cluster server 844
 - Microsoft 825, 838
 - adding or removing a failover node 825
 - all drives pool 821
 - backing up shared disks 838
 - BEUtility 825
 - changing the order in which nodes fail over 824
 - configurations 827–829, 832
 - disaster recovery 854–856
 - disaster recovery of Backup Exec on a cluster using IDR 852
 - failover restart 816
 - installation 818, 820
 - local disks 837
 - overview 836
 - restoring 840
 - uninstalling Backup Exec 820
 - restoring
 - specifying a new drive letter for the Microsoft cluster quorum disk 841
 - SAP Agent 1377
 - sizes for FAT partitions 796
 - troubleshooting 857
 - using with Backup Exec 814
 - VERITAS
 - backing up 847–849
 - disaster recovery 850, 856
 - overview 846
- combination SAP database server/media server recovery 1382
- command line installation of Backup Exec 153
- common encryption keys 417
- completed jobs
 - job log overview 577
- components of DLO 1604
- compression
 - delta file transfer 1666
 - delta file transfer in DLO 1671
 - compression (*continued*)
 - DLO backup selections
 - setting 1666
 - enable hardware compression option 460
 - hardware 460
 - setting backup defaults 395
 - using with encryption 415
 - configuration settings
 - copying to another server 199
 - Configuration Settings report 737
 - Configuration tab
 - for robotic library 471
 - configuration wizard for DLO 1641
 - Configure Alerts dialog box (DLO) 1724
 - Configure desktop and laptop backups option 1638
 - Configure Devices Assistant
 - about 443
 - configuring storage devices 444
 - configuring
 - devices 460
 - holidays 366
 - configuring Backup Exec Retrieve 871
 - connection based policies
 - configuring in DLO 1778
 - using Active Directory 1778
 - consistency check options
 - Exchange Agent 1137
 - SQL Agent 1260
 - continuing Active Directory backup if consistency check fails 884–885, 888
 - continuing Exchange backup if consistency check fails 1137
 - Continuous backup option for Exchange data 1135
 - continuous protection
 - Exchange data
 - best practices 1117
 - configuring for 1115
 - stopping 1120
 - troubleshooting 1123
 - for Exchange data
 - overview 1112
 - job statuses in Exchange 1118
 - requirements for Exchange 1113
 - review disk space 1119
 - viewing the console 1121
 - control connection with Remote Media Agent 1966
 - copy jobs
 - selection lists
 - and policies 554

- current jobs
 - custom filters 583
- custom installation 119
- custom reports
 - copying 718
 - creating 699
 - deleting 719
 - editing 718
 - graph options 706
 - grouping fields 702
 - overview 699
 - previewing 712
 - sorting fields 704
- D**
- daily backups
 - defined 277
- Daily Device Utilization report 738
- damaged media
 - removing 259
- DAOS
 - .nlo files 1066
 - about the Lotus Domino Agent and DAOS 1066
 - DAOS-enabled databases 1066
- data connection to remote computers 1966
- data source
 - adding 871
 - deleting 874
 - editing 873
 - options 872
- database 1607
- database files
 - backing up in a Microsoft cluster 838
 - backing up in a VERITAS cluster 849
- database instance
 - BKUPEXCDLO MSDE 1738
- database maintenance
 - configuring 209
 - overview 209
- database server
 - defined 1992
 - in Microsoft clusters 825
- Database snapshots
 - SQL 1274
- Date Modified tab 632
- DB2 Agent
 - archive logging methods. using with Backup Exec 974
 - DB2 Agent (*continued*)
 - archive logs template name
 - configuring for DB2 instance 962
 - authentication
 - configuring on DB2 instance 960
 - backing up 964
 - credentials
 - updating for instance 961
 - database access
 - adding on media server 956
 - deleting on media server 957
 - editing on media server 957
 - db2.conf configuration file 973
 - db2.conf file
 - creating 975
 - overview 975
 - DBA-initiated jobs
 - about 973
 - job template name for 962
 - defaults for backup and restore 958
 - example script
 - for command line processor 973
 - features 953
 - installation and configuration 954
 - multiple data streams
 - specifying 966
 - overview 953
 - redirected restore 970
 - troubleshooting 978
 - user exit db2uext2.exe
 - installed 973
 - user exit method
 - configuring for 955
 - vendor library db2sqluv.dll
 - installed 973
 - DBA-initiated job settings
 - about configuring 423
 - for SAP 1369
 - SAP Agent 1370
 - DBA-initiated jobs
 - creating a template 424
 - deleting a template 435
 - editing 434
 - Debug Monitor 810
 - Deduplicaiton Option
 - preparing for disaster recovery 1598
 - Deduplication device summary report 739
 - Deduplication Option
 - about backing up 1595

Deduplication Option *(continued)*

- about copying deduplicated data to tapes 1597
- about restoring 1597
- adding deduplication storage folder 1580
- adding OpenStorage device 1574
- adding Remote Agent for Deduplication 1592
- changing logon account password for
 - deduplication storage folder 1584, 1588
- client-side deduplication configuration 1591
- client-side deduplication overview 1590
- copying data between OpenStorage devices or deduplication storage folders 1595
- deduplication methods for agents 1570
- deduplication storage folder overview 1579
- deduplication storage folder properties 1584
- disaster recovery of deduplication storage folder 1597
- disaster recovery of OpenStorage devices 1599
- installing 1573
- OpenStorage device overview 1574
- OpenStorage device properties 1578
- overview 1568
- performing disaster recovery of deduplication storage folder 1598
- Remote Agent for Deduplication properties 1594
- requirements 1572
- setting up optimized duplication 1596
- sharing devices 1589
- with encryption 1597
- deduplication storage folder
 - adding 1580
 - changing logon account password 1584, 1588
 - overview 1579
- deduplication storage folders
 - about disaster recovery 1597
 - performing disaster recovery 1598
 - preparing for disaster recovery 1598
 - requirements 1572
 - viewing properties 1584
- default options
 - backup 392
 - backup and restore for Agent for Microsoft Hyper-V 1186
 - backup and restore for Agent for VMware 1406
 - backup and restore for Exchange Agent 1123
 - IDR
 - setting 1812
 - NDMP backup and restore 1867
 - network and security 405

default options *(continued)*

- overview 69
- pre/post commands 400
- restore 638
- setting 194
- setting for backup jobs 391
- SQL Agent
 - backup and restore 1265
- default preferred configuration settings for devices 462
- default settings
 - changing for DLO 1625
- defaults
 - device and media job 196
- deleting
 - Automated User Assignments in DLO 1687
 - device pools 519
 - devices from pools 519
 - DLO desktop computer 1704
 - media 260
 - revisions in DLO 1667
 - Storage Location in DLO 1682
 - user entry from DLO 1700
 - vault 253
- delta file transfer 1666, 1671
- deploying Silverlight in your organization 869
- desktop
 - defined for DLO 1756
- Desktop Agent
 - defined 1756
 - see also Desktop and Laptop Option (DLO) 1610
- Desktop Agent (DLO)
 - advanced view option 1768
 - backing up data 1765
 - backup selections
 - modifying 1767
 - overview 1765
 - console 1763
 - customizing installation 1612
 - Desktop User Data Folder
 - moving 1777
 - filter options
 - History view 1791
 - History view 1790
 - install set default location 1758
 - installing 1610
 - log files
 - overview 1790

Desktop Agent (DLO) *(continued)*

- menu bar
 - described 1765
 - overview 1755
 - Reset accounts option 1762
 - reset dialogs option 1762
 - Restore dialog box 1787
 - restoring files 1786
 - scheduling backup jobs 1773
 - standard view option 1767
 - Status view 1782
 - synchronization
 - create new sets 1779
 - delete synchronized folder 1781
 - Synchronized Selections view 1778
 - tasks bar
 - described 1765
 - views menu
 - described 1765
- Desktop Agent Users
- managing 1696
- desktop and laptop backups option
- Configure 1638
- Desktop and Laptop Option (DLO) 1604
- access
 - disabling/enabling 1700
 - adding user 1698
 - administration console 1638
 - administrator accounts 1618
 - administrators
 - creating 1618
 - alert history 1724
 - alert notification
 - printers 1732
 - alerts
 - categories 1721
 - Configure Alerts dialog box 1724
 - deleting 1726
 - displaying 1724
 - managing 1725
 - monitoring 1720
 - Automated User Assignment
 - creating 1684
 - defined 1604
 - deleting 1687
 - modifying 1686
 - priority 1686
 - properties 1686

Desktop and Laptop Option (DLO) *(continued)*

- automated user assignment
 - defined 1756
- backup selection
 - adding 1658
 - deleting 1670
 - modifying 1670
- BEUtility.exe utility
 - using 1738
- clustering 1738
- command line interface
 - emergencyrestore 1752
 - enableuser 1743
 - keytest 1745
 - listprofile 1746
 - listsl 1747
 - listuser 1748
 - logfile 1748
 - setrecoverypwd 1752
 - update 1749
 - assignSL 1742
 - changeserver 1744
 - remote server options 1741
 - syntax 1740
- Computer History pane 1716
- Configuration Wizard 1641
- configuring 1640
- deleting entry from DLO database 1704
- desktop
 - defined 1756
- desktop user data folder
 - defined 1756
- encryption
 - setting for backup selection 1666
- filter options
 - History view 1718
- History view 1715
- import multiple users in CSV file 1699
- include/exclude 1662
- installing 1610
- Job History pane 1716
- Move priority down option 1686
- Move priority up option 1686
- MSDE database instance
 - maintaining 1738
- network user data folder
 - defined 1756
- overview 1604

- Desktop and Laptop Option (DLO) *(continued)*
 - profile
 - creating 1641
 - defined 1756
 - properties
 - changing user 1699
 - removing user 1700
 - reset dialogs and accounts 1762
 - restoring 1707
 - revisions
 - defined 1663
 - deleting automatically 1667
 - setting number to keep in DLO 1664
 - Search history log file option 1720
 - see also Desktop Agent 1610
 - storage limits for user data 1644
 - Storage Location
 - defined 1604
 - deleting 1682
 - moving users 1701
 - synchronization
 - defined 1756
 - user bandwidth settings 1643
 - User Data Folder 1676
 - User Properties dialog box 1699
 - View history log file option 1717
 - viewing users 1703
- Desktop data
 - backing up with DLO 1765
- desktop user data folder
 - defined 1756
- destination media server
 - adding in CASO environment 201
 - adding in non-CASO environment 200
 - adding multiple 200
- device
 - allocation in a shared storage environment 1996
 - selecting for backup job 339
 - selecting for duplicate backup job 373
- device and media data
 - location of in CASO 1509
- Device and Media options for backup job 339
- device management 441
- device operations
 - cataloging media 248
 - ejecting media 487
 - Enable Hardware Compression Option 460
 - formatting media 485
 - labeling media 486
- device operations *(continued)*
 - overview of utility jobs 480
 - retensioning a tape 484
 - robotic library 453
 - using with SAN Shared Storage Option 2003
 - Virtual Tape Library Unlimited Drive
 - Option 452
- device pools
 - adding devices 517
 - All Virtual Disks 2029
 - creating 516
 - deleting 519
 - overview 515
 - prioritizing devices 518
 - properties 520
 - removing devices 519
 - using the default device pool 516
- device properties
 - Configuration tab 460
 - General tab 458
 - media types 466
 - SCSI Information tab 463
- Device Summary report 740, 2005
- Device Usage by Policy report 741
- devices
 - about 441
 - adding deduplication storage folder 1580
 - adding OpenStorage device 1574
 - block size 460
 - buffer count 462
 - buffer size 462
 - configuring 460
 - default settings 462
 - high water count 462
 - Hot-swappable Device Wizard 453
 - iSCSI-attached
 - adding 453
 - OpenStorage overview 1574
 - pausing 446
 - pausing a media server 446
 - renaming 447
 - resuming 447
 - resuming a media server 446
 - SCSI information 463
 - setting default for jobs 196
 - sharing deduplication devices 1589
 - specifying media types for 457
 - statistics on usage 463
 - statistics since cleaning 465

- devices *(continued)*
 - Symantec Device Driver Installation Wizard 455
 - USB tape devices
 - reconnecting 453
- diagnostic file
 - command line switches 806
 - remote media server 808
- dialog box (DLO)
 - Move User 1701
- differential backups
 - advantages and disadvantages 278
 - defined 276
- Direct Access Recovery 1863, 1867
- DirectCopy to tape
 - copying data 382
 - overview 381
- directories
 - about including and excluding for NDMP 1854
 - excluding from EMC backup 1860
 - excluding from NetApp backup 1859
 - including in NDMP backup 1856
- Directory database
 - backing up for Archiving Option 1477
- disable backup-to-disk folders for Backup Exec 503
- disable device for Backup Exec 459
- Disabled job status 1429
- disabling backup mode in Archiving Option 1480
- disaster preparation
 - Disaster Preparation Plan (DPP) 776
 - emergency repair disk 779
 - Exchange Server 1165
 - hardware protection 776
 - Lotus Domino Agent 1086
 - off-site storage 777
 - overview 775
- disaster recovery
 - clusters
 - Backup Exec on a Microsoft cluster using IDR 852
 - entire cluster manually 853
 - nodes using IDR 851
 - using IDR to prepare 851
 - data protected by Backup Exec agents 780
 - deduplication storage folder 1598
 - different types of computers
 - overview 780
 - Exchange Server 1166
 - local Windows 2000 computers
 - (non-authoritative) 780
 - disaster recovery *(continued)*
 - Lotus Domino Agent 1087
 - manual recovery of Windows system 780
 - Microsoft clusters
 - Backup Exec 856
 - data files 854
 - shared disks 855
 - OpenStorage devices 1599
 - overview 780
 - performing using SAP Agent 1380
 - remote Windows 2000 computers
 - (non-authoritative) 785-786
 - VERITAS clusters
 - overview 850
 - shared disks 856
 - disaster recovery alternate data path
 - in IDR 1813
 - disaster recovery data path
 - in IDR 1812
 - disaster recovery file (*.dr file) in IDR
 - defined 1810
 - setting locations for 1811
 - Disaster Recovery Wizard
 - requirements 1830
 - running 1831
 - disk space usage trends in the Storage Provisioning Option
 - configuring 2047
 - description 2046
 - distributed catalog
 - in CASO 1541
 - Distributed File System (DFS), backing up 294
 - DLO Administration Console
 - restoring from 1707
 - DLO Administration Server
 - connecting to 1639
 - DLT tape
 - drive hangs when cataloging 794
 - Domain Controller
 - using redirected restore to install from media 636
 - domains
 - host and target
 - defined 110
 - drive pools
 - creating in a Microsoft cluster 821
 - SAN Shared Storage Option 1996
 - drivers
 - download latest 792

- duplicate backup sets template
 - adding to a policy 550
 - overview 548
- duplicating backup data
 - about 369
 - creating duplicate backup jobs 369
- duplication between OpenStorage devices or deduplication storage folders 1595
- dynamic inclusion
 - for Hyper-V 1174
 - for SharePoint 1194

E

- e-mail
 - configuring MAPI notification 664
 - configuring notification in DLO 1728
 - configuring SMTP notification 662
 - configuring VIM notification 665
- editions of Backup Exec
 - listed and described 63
- eject media
 - after job completes 341, 375
 - from a drive 487
- emergency repair disk
 - creating 779
- enabling
 - backup-to-disk folders for Backup Exec
 - option 503
 - device for Backup Exec option 459
- encrypted files
 - about cataloging media 423
 - backing up with AOFO 941
- encrypted SQL database restore 1295
- encryption
 - about 415
 - hardware 416
 - in DLO backup selections 1666
 - restoring encrypted SQL databases 1295
 - SAP data 1366
 - software 415
 - types 415
 - with deduplication 1597
- encryption keys
 - 128-bit AES 415
 - 256-bit AES 415
 - about deleting 421
 - common 417
 - creating 420
 - deleting 422

- encryption keys *(continued)*
 - encryption types 415
 - managing 418
 - overview 416
 - pass phrases 417
 - replacing 421
 - restoring encrypted data 422
 - restricted 417
 - setting a default 404
 - using with compression 415
- Enterprise Vault
 - running services for 1422
 - viewing event log 1500
- Enterprise Vault Agent
 - about restoring 1012
 - About restoring individual files and folders 1029
 - Audit database
 - restoring 1021
 - automatic redirection of Enterprise Vault
 - components 1014
 - available backup methods 989
 - Backup Exec media server
 - log file location 1048
 - logs 1047
 - Backup Exec Migrator
 - about 1041
 - about deleting migrated files 1048
 - about events 1046
 - about logs 1047
 - about retrieving Enterprise Vault data 1058
 - about staged migrations 1046
 - about the Backup Exec restore browse
 - view 1057
 - best practices 1061
 - communicating with Enterprise Vault 1054
 - configuring 1049
 - data migration process 1046
 - Enterprise Vault retention periods 1048
 - how it works 1042
 - log file location 1048
 - Migrator for Enterprise Vault options 1053
 - requirements 1041
 - retrieving Enterprise Vault data 1058
 - troubleshooting 1062
 - VxBSA logs 1047
 - working with a Backup Exec media
 - server 1051
 - best practices for the 1041

Enterprise Vault Agent *(continued)*

- closed partition
 - backing up 995
- collections
 - configuring 1050
 - vault store partition properties 1051
- compliance accelerator configuration database
 - backing up 1004
 - restoring 1024
- compliance accelerator customer database
 - backing up 1004
 - restoring 1025
- Directory database
 - backing up 998
 - restoring 1015
 - restoring to a different SQL Server 1039
- discovery accelerator configuration database
 - backing up 1005
 - restoring 1026
- discovery accelerator custodian database
 - backing up 1006
 - restoring 1027
- discovery accelerator customer database
 - restoring 1028
- discovery accelrator customer database
 - backing up 1005
- Enterprise Vault 7.x server
 - backing up 1009
 - restoring 1033
- Enterprise Vault 8. x audit database
 - backing up 1001
- Enterprise Vault 8.x Fingerprint database
 - backing up 1003
- Enterprise Vault 8.x FSA Reporting database
 - backing up 1002
- Enterprise Vault 8.x site
 - about backing up 1009
- Enterprise Vault server
 - about backing up 1009
- Enterprise Vault site
 - backing up 1010
- Fingerprint database
 - restoring 1023
- FSA Reporting database
 - restoring 1022
- index locations
 - backing up 1011
- installing 988

Enterprise Vault Agent *(continued)*

- migration
 - vault store partition properties 1056
- Monitoring database
 - backing up 999
 - restoring 1016
- non-operational state 1012
- open partition
 - backing up 994
 - restoring individual files 1030
- Partition Recovery Utility
 - about 1059
 - finding an archive ID 1060
 - log file location 1048
 - logs 1047
 - requirements 1059
 - running 1060
 - troubleshooting 1062
- partitions
 - restoring 1017
- ready partition
 - backing up 997
- ready-to-use state 1012
- redirecting a restore job 1036
- redirection options 1037
- requirements 987
- restore options 1034
- restoring closed partitions 1017
- Restoring folders from an Enterprise Vault
 - index 1032
- Restoring individual files from an open
 - partition 1030
- restoring open partitions 1017
- restoring ready partitions 1017
- selecting a backup method 988
- setting a default backup method 993
- vault store
 - backing up 1007
- vault store database
 - backing up 1000
 - restoring 1019
- Environment Check
 - running for Backup Exec 104–105
- error codes
 - Unique Message Identifier
 - viewing 577, 658
- error-handling rules
 - Cluster Failover rule 591, 595
 - configuring 591

error-handling rules *(continued)*

- custom rules
 - defined 591
- default rules
 - defined 591
- overview 590
- Recovered Jobs custom rule 591
- ESX server, adding 1387
- ESX server, deleting 1388
- EV Agent
 - Automatically backed up components 983
- Event Recipients report 743
- example policies
 - re-creating 528
 - using 526
- Exchange Agent
 - Active Directory
 - backing up 1105
 - automatic exclusion of files during volume level
 - backups 1105
 - backing up
 - Exchange 2003/2007 overview 1129
 - Exchange Server 2010 overview 1130
 - individual mailboxes 1143
 - recommended selections 1105
 - backup methods 1124, 1134
 - backup options 1134
 - backup selections
 - adding a forest 1130
 - managing a forest 1131
 - best practices 1100
 - change password when recreating
 - mailboxes 1128
 - circular logging
 - reviewing 1106
 - Continuous backup option 1135
 - continuous protection
 - best practices 1118
 - configuring 1117
 - overview 1112
 - requirements 1113
 - restoring latest full transaction 1151
 - restoring point in time 1152
 - restoring the information store 1151
 - review disk space 1119
 - stopping backup jobs 1120
 - troubleshooting 1123
 - viewing the console 1121
 - creating backup job 1132

Exchange Agent *(continued)*

- databases
 - configuring 1146
 - dismounting for restore 1146
- default backup and restore options 1123
- disaster recovery 1166
- Exchange 2003 with VSS
 - backing up 1108
- Exchange 2007 snapshot backup method 1109
- Exchange high availability server option 1138
- Exchange Server 2010 forest
 - management options 1132
 - options 1131
- Exchange Web Services
 - overview 1107
- exclude specific folders 1141
- excluding files during volume level
 - backups 1105
- Granular Recovery Technology (GRT)
 - overview 1106
 - requirements for 1097
- Granular Recovery Technology (GRT) option
 - setting for backup 1136
- Guide Me wizard for backup 1136
- installation 1099
- Internet Information Service (IIS) metabase
 - backing up 1104
- legacy mailbox backup options 1125, 1140
- legacy mailbox or public folders
 - enabling access 1162
- mailbox access requirements 1101
- mount database after restore option 1159
- offhost backup
 - configuring 1111
 - with Granular Recovery Technology (GRT) 929, 1106
- overview 1094
- protecting Exchange using VSS 1108
- recovery point option 1135
- recovery points 1122
- recreating mailboxes and user accounts 1156
- redirecting data 1159
- redirecting mailboxes 1160
- redirecting storage groups and databases 1160
- Redirection dialog box 1163
- redirection options 1163
- requirements 1095
- resource discovery feature
 - using with 1096

Exchange Agent (*continued*)

- restore from continuous protection
 - backups 1150
 - Restore Job Properties dialog box 1155
 - restore of individual items
 - requirements for 1097
 - restore options 1155
 - restore over existing messages and folders when
 - restoring individual items 1156
 - restore requirements 1145
 - restoring
 - commit after restore completes
 - option 1159
 - DS/IS consistency adjuster after
 - restore 1162
 - Exchange data 1154
 - temporary location for log and patch
 - files 1158
 - restoring an Exchange 2007 database to a
 - recovery storage group 1148
 - restoring data from snapshot backups 1149
 - restoring data to server 1144
 - restoring Exchange 2003 and 2007 with
 - Recovery Storage Group 1147
 - restoring individual public folder messages from
 - tape 1153
 - restoring mailboxes and public folders
 - overview 1152
 - restoring mailboxes and user accounts 1156
 - services accounts
 - overview 1095
 - setting defaults 1123
 - snapshot backup
 - configuring 1110
 - snapshot technology
 - and 1108
 - storage groups
 - backing up 1129
 - strategies for backing up 1102
 - system state
 - backing up 1104
 - troubleshooting snapshot and offhost jobs 1109
 - volume level backups
 - automatic exclusion of files 1105
- Exchange Mailbox Archiving Option
- overview 1414
- Exchange Mailbox Group Archive Settings 769
- Exchange Redirection page 1163

Exchange Server Agent

- excluding files during volume level
 - backups 1073
- Exchange Web Services
 - using with the Exchange Agent 1107
- executing a command
 - after backup 353
 - after restore 353
 - before backup 353
 - before restore 353
- export media template
 - about 536
 - adding to a policy 537
- exporting media 490

F

- failback
 - defined 816
- Failed Archive Jobs 770
- Failed Backup Jobs report 744
- failover
 - adding or removing a failover node 825
 - changing the order in which nodes fail over 824
 - defined 814
 - restart 816
- farms
 - adding 1191, 1219
 - changing the default farm name 1195
 - deleting 1195
- FAT
 - cluster size 796
 - partition 796
- father/son media rotation strategy 266
- Favorite Resources
 - about 284
 - adding a Windows system 285
 - deleting a Windows system 286
- file access
 - secured in Backup Exec Retrieve 862
- file history
 - enabling for NDMP 1867
- file permissions
 - restoring 619
- File Replication Service (FRS), backing up 294
- File System Archive Settings 770
- File System Archiving Option
 - overview 1414
- file to add users in DLO
 - CSV 1699

filegroups

restoring

- nonprimary SQL 2000 1294
- primary SQL 2000 1294
- SQL Agent 1298

files

- about including and excluding for NDMP 1854
- erasing from Symantec Online Storage for Backup Exec 2060
- excluding from EMC backup 1860
- excluding from NetApp backup 1859

filters

custom

- current jobs 583
- job history 585
- jobs 583

- for alerts 645, 648
- for jobs 582

History view filters

- setting in Desktop Agent 1791
- setting in DLO 1718

finding media in a location or vault 254

fingerprint database

- backing up for Archiving Option 1477
- for vault stores 1445

firewall

- browsing systems through 414
- enabling a SQL instance behind 414
- using Backup Exec with 409

formatting media 485

full backups

- advantages and disadvantages 278
- defined 275

fully qualified computer name 291

G

Gather Utility 808

- collecting log file 808

general job defaults

- setting 196

general options for restore jobs 612

global excludes

- adding 1687–1688, 1692
- deleting 1687–1688
- email 1690
- encryption 1693
- macros 1696

grandfather media rotation strategy 267

Granular Recovery Technology (GRT)

- about restoring individual items 321
- Agent for Microsoft Servers 1177
- Agent for Microsoft SharePoint 1196, 1203, 1205, 1220, 1223, 1238
- enabling for Microsoft Hyper-V 1175
- Exchange data 1106
 - offhost backup 1106
- reclaiming disk space for 513
- recommendations for using backup-to-disk folders with 511
- recommended devices for 324
- requirements 325
- temporary staging location
 - setting as a default 641
 - setting for a job 617
- using Exchange Web Services 1107
- VMware resources 1394

grooming

- files in DLO 1664

groups

- configuring recipients 677, 1732

H

hardware

- creating profile 778
- enable hardware compression option 460
- protection in case of disaster 776
- troubleshooting 789

hardware compression

- enabling 457

high water count

- setting for devices 462

highly available virtual machines

- about backing up and restoring 1188

hold jobs 568

holding jobs that back up selection lists 303

Home view

- about 95
- configuring 95
- Detail items 99
- editing items 96
- Help and Technical Support items 96
- restoring the default configuration 95
- Summary items 98

host domain

- defined 110

hot key

- defined 2064

- hot spare
 - best practices 2042
 - changing or adding 2043
 - description 2042
 - specifying 2021–2022
- I**
- IBM computers
 - recovering with IDR 1830
- IDR Configuration Wizard 1810
- IMG subfolders
 - described 496
- imported media
 - labeled by Backup Exec 244
 - overwriting 232
- importing media 489
- importing templates 538
- include/exclude
 - DLO backup selections 1662
 - files for backup 355
- incremental backups
 - advantages and disadvantages 279
 - defined 276
- index locations
 - backing up for Archiving Option 1477
- initialize job for robotic library 484
- install from media 637
- installation
 - about typical and custom 116
 - custom 119
 - NDMP Option 1849
 - Remote Media Agent for Linux Servers 1968
 - typical 116
 - using installation program on installation media 119
- installation log 166
 - Remote Agent for Linux or UNIX Servers 1877
- installation overview 1170
- installation parameter file
 - creating 165
 - defined 164
 - using 165
- installed updates
 - viewing 173
- installing
 - additional Backup Exec options on a Microsoft cluster 820
 - Backup Exec in a Microsoft cluster 818
 - Backup Exec in a VERITAS cluster 844
 - installing (*continued*)
 - besernum.xml file
 - import license keys 121
 - Desktop Agent 1610
 - Desktop Agent options 1612
 - Desktop and Laptop Option 1610
 - Domain Controllers from media 636
 - Environment Check
 - running pre-install 104–105
 - import license keys 121
 - locally
 - additional options 124
 - methods 102
 - Microsoft SQL Server 2005 Desktop Engine (MSDE 2005) 111
 - Remote Administrator 150
 - SharePoint Agent 1191
 - silent install of DLO 1612
 - to an existing Microsoft SQL Server 2005 instance 112
 - trial version 121
 - using Repair option 167
 - using Terminal Services 119
 - Windows Management Instrumentation
 - performance counter 687
 - Windows Management Instrumentation SNMP provider 688
 - installing Backup Exec Retrieve 869
 - instance
 - BKUPEXCDLO MSDE database 1738
- Intelligent Disaster Recovery (IDR)
 - ASR files 1811
 - automated restore 1832
 - backup-to-disk folder
 - using in 496
 - boot managers 1828
 - catalog entries
 - added to *.dr file 1806
 - clusters
 - recovering Backup Exec on a Microsoft Cluster 852
 - recovering nodes 851
 - disaster recovery file (*.dr file)
 - defined 1810
 - encrypted backup sets 1831
 - hard drive partition
 - altering sizes using IDR 1841
 - IBM computers 1830
 - installing 1807

Intelligent Disaster Recovery (IDR) *(continued)*

- Microsoft Exchange Server
 - recovering 1844
 - Microsoft SQL Server
 - recovering 1844
 - Options - Set Application Defaults dialog
 - box 1812
 - OS/2 boot manager
 - restoring 1828
 - overview 1806
 - Recovery Wizard
 - running 1831
 - requirements 1807
 - restoring from a locally-attached media
 - device 1835
 - restoring from a remote media server 1839
 - restoring from remote backup-to-disk
 - folders 1837
 - SharePoint Portal Server
 - recovering 1844
 - System Commander boot manager
 - restoring 1828
 - utility partitions
 - backing up 1815
 - Windows Automated System Recovery (ASR)
 - files 1811
- Internet Information Services (IIS) metabase
- backing up 1104
- inventory
- all drives when Backup Exec starts 197
 - robotic libraries when Backup Exec starts 483
- IPv4 404
- IPv6 404
- iSCSI-attached devices
- adding 453
- ## J
- job delegation
 - in CASO 1504
 - Job Distribution by Device report 745
 - job history 1715
 - custom filters 585
 - deleting report 698
 - History view filters in DLO
 - setting in DLO 1718
 - saving report 696
 - viewing 572
 - job history (DLO)
 - viewing in Desktop Agent 1790

- job log
 - configuring default options 580
 - setting options in CASO 1536
 - status overview 577
- job monitor
 - views 557
- job priority, overview 195
- job progress indicators
 - displaying 197
- job queue
 - hold 569
- job status
 - and setting thresholds for 596
- jobs
 - about creating from policies 544
 - about restoring from Symantec Online Storage
 - for Backup Exec 2062
 - about scheduling 356
 - calendar
 - managing jobs from 589
 - calendar view of workload 588
 - changing priority for scheduled 570
 - configuring default Lotus Domino options 1069
 - configuring default schedule 366
 - configuring error-handling rules 591
 - configuring schedule 356
 - creating from policies 544
 - creating from selection lists 545
 - deleting jobs created from policies 546
 - deleting scheduled 571
 - editing the next occurrence of policy-based
 - jobs 546
 - filtering 582
 - filters 582
 - hold queue 569
 - holiday scheduling 366
 - Lotus Domino backup properties 1076
 - managing custom filters 583
 - removing hold
 - active jobs 563
 - scheduled jobs 570
 - renaming jobs created from policies 547
 - restarting during a time interval 365
 - run report 694
 - running scheduled job 568
 - running test for scheduled job 570
 - scheduling 356
 - sending notification when complete 682
 - setting general defaults 196

jobs *(continued)*
 viewing and scheduling in a SAN 1997
 viewing completed 572

Jobs Summary report 745

junction points
 backing up 350

K

keyboard navigation
 dialog boxes 2071
 shortcuts 2063
 standards 2070
 keyboard shortcuts 2064

L

labeling media
 creating default labels 239
 imported media label 244
 in drive 486
 renaming 243
 using bar code labels 244
 last known good menu 777
 legacy mailbox backup methods
 in Exchange 1140
 Library Expansion Option
 overview 453
 SCSI addresses for hardware 468
 setting up hardware 468
 license information
 finding in your environment 176
 license keys 124
 Backup Exec
 adding 175
 adding and removing 120, 124
 viewing 174
 finding in your environment 176
 Remote Agent for Windows Systems 1946
 list boxes
 navigation 2072
 LiveUpdate
 about 170
 about scheduling automatic updates 171
 running manually 173
 scheduling automatic updates 172
 local media server
 breaking connection with 151
 location
 media in Backup Exec 250

lock open files for backup 352, 398

lock robotic library panel 493

logon accounts
 changing default 193
 changing for a resource being backed up 337
 changing for a resource being restored 337
 changing the password 191
 creating 187
 default
 defined 185
 deleting 193
 editing 190
 overview 184
 replacing 192
 restricted 186
 SQL resources 1256
 system logon account 189
 testing 337

logon information
 copying to another server 204

Lotus Domino Agent
 APIs 1071
 archive logging 1074
 backup options 1076
 circular logging 1074
 configuring default options 1069
 database backup overview 1071
 database backup requirements 1065
 disaster preparation 1086
 disaster recovery
 archive logging 1090
 circular logging 1090
 of server 1087
 Microsoft Cluster Server 1065
 restoring 1079
 overview 1064
 redirecting restore 1083
 requirements 1064
 restore overview 1078
 selecting for restore 1079
 selecting restore options 1082
 supported configurations 1073
 viewing databases 1068

Lotus Notes
 backing up nsf files with DLO 1770

low disk space thresholds
 backup-to-disk folder option 501
 editing for a virtual disk 2030
 editing global defaults for virtual disks 2032

M

Machines Backed Up report 746

macros

global excludes 1696

mailbox access requirements for Exchange 1101

mailbox groups

arranging for provisioning 1463

arranging the order of 1443

creating for an archive job 1443

managing 1463, 1495

overview 1461

setting rules for archiving 1462

mailboxes

exclude specific folders 1141

redirecting restores 1160

restore overview 1152

Maintenance contract information

about 177

Renewal Assistant 177

running the 178

serial number and license key formats 177

viewing details 179

maintenance server (DLO)

delegation

maintenance server (DLO) 1673

majority node in a cluster 817

managed media server

copying jobs to 1551

defaults

setting 1531

installing 1513

network connection speed to central

administration server 1532

network interface card

using any available 1540

pools 1545

upgrading 1521

Managed Media Servers report 747

manually update server list

for NetWare 1942

MAPI

configuring recipients 668, 1729

e-mail notification method 664

e-mail notification method in DLO 1728

master database (SQL)

backup 1260

restore 1290

MaxDB databases

protecting with SAP Agent

backing up 1378

overview 1362

restoring 1380

media

about inventorying 448

adding to the offline location or user-defined

media vault 255

all media

defined 220

append backup to 341

associating with a media set or vault 229

categories 219

creating default labels 239

damaged 259

deleting 260

deleting vault 253

displaying media ID 261

drag and drop

to move media 258

erasing 449

finding in a location or vault 254

general properties 261

inventorying in a device 448

moving to a location or vault 254

overwrite for backup 341

overwrite for duplicate backup 374

overwrite options 233

overwriting allocated or imported 232

properties 261

retired

defined 221

scanning bar code labels 255

scheduling a job to move media 255

scratch

defined 221

setting default for jobs 196

setting default options 237

statistical properties 263

testing integrity of 382

Vault Wizard 256

with excessive errors 259

Media Audit report 748

media capacity

testing before backup runs 388

media catalogs 1992

Media Errors report 749

- media ID
 - defined 242
- media label
 - bar code rule
 - in mixed media libraries 245
 - bar codes 244
 - deleting bar code rule 246
 - editing bar code rule 245
 - imported 244
 - overview 242
 - renaming 243
- media location
 - updating 257
- media operations
 - associating media with media sets 229
 - audit log for 241
 - deleting media 260
- media overwrite protection
 - overview 226
- media overwrite protection level
 - defined 232
- Media Required for Recovery report 749
- media rotation 521
 - strategies
 - father/son 266
 - grandfather 267
 - son 265
- media server
 - connecting to in DLO 1639
- media set
 - creating 227
 - creating by using a wizard 227
 - default 226
 - defined 220
 - deleting 228
 - overview 226
 - overwrite and append properties 230
 - renaming 228
 - selecting for backup job 341
 - selecting for duplicate backup job 374
 - shared storage environments 1997
 - vault rule properties 252
- Media Set report 750
- Media Set Wizard 227
- media types
 - specifying for devices 457
- media vault
 - defined 251
 - deleting 253
- media vault (*continued*)
 - finding media in a vault 254
 - moving media 254
 - renaming 254
 - user-defined 251
- Media Vault Contents report 751
- media view 600
- menu bar
 - overview 92
- menus
 - described
 - Help 2070
 - Window 2066
- messages
 - error 794
- Microsoft Cluster Server
 - using with Backup Exec 816
- Microsoft clusters
 - database files 838
- Microsoft SQL Server 2005 Desktop Engine (MSDE)
 - installing 111
- Microsoft Terminal Services
 - and installing Backup Exec 119
- Microsoft Virtual Hard Disk files
 - about managing 293
 - about redirecting restore jobs to VHD files 636
- Microsoft Volume Shadow Copy Service (VSS)
 - and AOFO 950
- Missed Availability report 752
- mixed media library bar code labeling 245
- mnemonic
 - defined 2064
- mobile phone notification 662
- Monitor jobs on local managed media server 1532
- mounted local drives
 - backing up 350
 - backing up files and directories 350
- Move Media to Vault report 752
- Move priority options in DLO 1686
- Move User dialog box in DLO 1701
- moving media
 - using drag and drop 259
- MSCS
 - using with Backup Exec 816
- MSDE 1607
 - 2005 components
 - installed with Backup Exec 111
 - database instance
 - BKUPEXCDLO 1738

MSDE (*continued*)

- database instance (*continued*)

- DLO 1738

- multi-stage backup strategy 548

N

- Name & Resource tab 632

- named transaction

- include in restore

- SQL 2000 1293

- restore up to

- SQL 2000 1293, 1297

- navigation

- list boxes 2072

- tabbed pages 2072

- navigation bar

- overview 93

- NDMP Option

- adding an NDMP server 1849

- backing up resources 1851

- duplicate backed up data 1862

- excluding EMC directories and files 1860

- excluding NetApp directories and files 1859

- how to use patterns 1857

- including specific directories 1856

- installing 1849

- overview 1847

- requirements 1848

- restoring data 1863

- sharing devices on an NDMP server 1850

- viewing server properties 1872

- NDMP option

- redirecting restored data 1867

- Net Send

- configuring recipients 674, 1731

- network

- for a backup job 407

- overview of backup networks 402

- setting up a backup network 404

- Network Attached Storage (NAS)

- protecting 1847

- network traffic

- reducing in CASO 1530

- network user data folder

- creating 1697

- defined 1756

- nodes

- configurations in a Microsoft cluster 827

- defined 814

- nodes (*continued*)

- disaster recovery using IDR 851

- Microsoft

- adding or removing a failover node 825

- changing the order in which nodes fail

- over 824

- notification

- recipients 1728

- notifications

- assigning recipients to alerts 680

- configuring in DLO 1728

- configuring MAPI e-mail 664

- configuring pager 666

- configuring SMTP e-mail 662

- configuring SNMP 683

- configuring VIM e-mail 665

- defined 645

- modifying recipient properties 678

- scheduling recipients 677

- sending for completed jobs 682

- sending when selection list used in jobs 682

- notifications (DLO)

- modifying recipient properties 1733

- Novell OES

- about restoring 1898

- backing up components 1896

- requirements for back up 1895

- supported components 1896

- nsf files

- backing up with DLO 1770

- NTFS

- cluster size 796

- partition 796

- numbering

- specify for media label 240

O

- off-site storage of backups 777

- offhost backup

- best practices 923

- for Exchange Server

- with Granular Recovery Technology (GRT) 928

- host computer

- defined 919

- List Snapshot Providers option 925

- overview 919

- requirements 921

- single volume snap 928

- offhost backup (*continued*)
 - snapshot provider
 - choosing 927
 - transportable snapshots
 - defined 920
 - VFSW FlashSnap option
 - using with 922
- offline
 - when backup-to-disk folders display as 503
 - when devices display as 459
- offline media location
 - adding media to 255
 - defined 251
- online
 - bring device 508
- online media location 250
- open files
 - backing up with AOFO 938
 - unable to back up 795
- OpenStorage devices
 - adding 1574
 - disaster recovery 1599
 - overview 1574
 - requirements 1572
 - viewing properties 1578
- Operations Overview report 753
- optimized duplication 1595
 - setting up 1596
- optimizing remote backups 351
- options
 - additional Backup Exec options described 80
 - default for job log 580
 - Set Application Defaults
 - Preferences 197
- Oracle Agent
 - authentication credentials 1329
 - deleting 1333
 - editing 1332
 - setting 1330
 - authentication credentials options 1331
 - authentication for Oracle operations 1329
 - back up with 1335
 - backing up resources 1337
 - backup options 1339
 - channel time-out
 - change default for 1355
 - configure 1316
 - configuring 1318
- Oracle Agent (*continued*)
 - database time-out
 - change default 1354
 - DBA-initiated backup 1340
 - DBA-initiated job settings
 - create template for 423
 - DBA-initiated jobs
 - job template name for 1326
 - DBA-initiated restore 1343
 - default options 1320, 1334
 - defaults for backup and restore operations 1334
 - device and media options 1340
 - features 1315
 - install 1316
 - legacy GRFS Oracle Agent database backups
 - restore from 1348
 - Linux servers
 - configuring an Oracle instance 1324
 - deleting an Oracle instance 1327
 - editing an Oracle instance 1326
 - enabling database access 1327
 - viewing an Oracle instance 1326
 - multiple data streams
 - specify 1340
 - Oracle Net Service name 1321
 - port
 - configure for DB2 and Oracle
 - operations 1329
 - publish Oracle databases on Linux 1325
 - Real Application Cluster (RAC) 1324–1325, 1337
 - recovery catalog 1321, 1325
 - redirected restore 1347
 - redirection options 1347
 - Remote Agent Utility options 1322
 - restore 1341
 - restore options 1344
 - restoring data 1343
 - troubleshooting 1354
 - update credentials for instances 1320, 1325, 1333
 - upgrading 1317
 - Windows computers
 - configuring an Oracle instance 1319
 - deleting an Oracle instance 1323
 - editing an Oracle instance 1322
 - enabling database access 1323
 - viewing an Oracle instance 1321
- Outlook PST files
 - backing up with DLO 1769

- Overnight Archive Summary 771
- overwrite default media label 239
- overwrite media option
 - specifying for backup job 341
 - specifying for duplicate backup job 374
- overwrite protection
 - disabling 238
- overwrite protection levels
 - full 237
 - partial 237
- overwrite protection period
 - defined 223, 231
 - setting for media set 231

P

- pager
 - configuring recipients 671, 1730
 - notification method 666, 1728
- partial overwrite protection 237
- partition
 - creating for robotic library 475
 - FAT 796
 - NTFS 796
 - redefining for robotic library 478
- Partition Recovery Utility
 - troubleshooting 1062
- Partition Recovery Utility
 - about 1059
 - finding an archive ID 1060
 - log file location 1048
 - logs
 - about 1047
 - requirements 1059
 - running 1060
- pass phrases 417
- password
 - changing for logon account 191
- password database
 - Remote Agent for NetWare Servers 1935
- patterns in NDMP excludes 1857
- performance
 - increase during backups of remote Windows computers 1946
- physical check
 - SQL 2000 1260, 1278
- physical disk
 - capacity 2026
 - creating a physical disk group 2021–2022
 - hardware health 2026

- physical disk *(continued)*
 - hardware status 2026
 - viewing properties 2023, 2026
- PHYSICAL_ONLY utility 1261
- placing scheduled job on hold if test run fails 389
- point in time log restore option
 - SQL Agent 1292, 1296
- policy
 - about creating jobs 544
 - adding a backup template 530
 - adding a duplicate backup sets template 550
 - adding an export media template 537
 - changing template rules 542
 - copying to another server 554
 - creating jobs 544
 - creating manually 523
 - creating synthetic backup using the Policy Wizard 905
 - creating using the Policy Wizard 523
 - deleting 526
 - deleting a template 539
 - deleting jobs created from policies 546
 - deleting template rules 543
 - duplicate backup sets template overview 548
 - editing 525
 - editing a template 539
 - enabling true image restore 916
 - importing templates 538
 - overview 521
 - re-creating example policies 528
 - renaming jobs created from policies 547
 - setting template schedules 532
 - template rules 542
 - using an example policy 526
 - using templates 529
 - viewing 545
- Policy Jobs by Resource Summary report 756
- Policy Jobs Summary report 757
- Policy Properties report 758
- Policy Protected Resources report 759
- port number, changing for Remote Media Agent for Linux Servers 1973
- portal support 490
- ports used by Backup Exec
 - default 411
 - Desktop and Laptop 413
 - listening 412
- post-job command
 - for backup jobs 399

- post-job command (*continued*)
 - for restore jobs 399
 - setting defaults 400
 - setting for backup job 353
 - setting for restore job 618
 - pre-job command
 - for backup jobs 399
 - for restore jobs 399
 - setting defaults 400
 - setting for backup job 353
 - setting for restore job 618
 - preferred server configurations
 - about 435
 - creating 436
 - deleting 438
 - designating a default 438
 - editing settings 438
 - removing as default 439
 - prefix
 - creating for media label 239
 - preserve tree option
 - for backup job 346
 - for restores 613
 - primary server defined 1992
 - printer
 - configuring recipients 675, 1731
 - priority
 - Automated User Assignment
 - changing in DLO 1686
 - changing for scheduled job 570
 - default for selection lists 307
 - for selection lists 306–307
 - options in DLO
 - Move 1686
 - setting for devices in pools 518
 - setting for restore job 612
 - Problem Files report 759
 - profile
 - defined 1756
 - profile (DLO)
 - copying 1657
 - creating 1641
 - defined 1604
 - prompt before overwriting allocated or imported
 - media 238
 - properties
 - active job 558, 561
 - alerts 651
 - editing job 556
 - properties (*continued*)
 - Lotus Domino 1076
 - media
 - general 261
 - statistical 263
 - report 721
 - user
 - changing in DLO 1699
 - viewing for NDMP servers 1872
 - PST files
 - backing up with DLO 1769
 - publish
 - default interval 1953
 - disable on remote computer 1953
 - to media servers
 - using Remote Agent for Windows Systems 1951
 - Publish Linux, UNIX, and Macintosh computers to
 - media server
 - about 1882
 - Publish Linux, UNIX, and Macintosh computers to
 - media servers
 - how to 1883
- ## Q
- QuickStart Edition of Backup Exec
 - described 68
- ## R
- ralus.cfg
 - about, for the Remote Agent for Linux or UNIX Servers 1881
 - configuration options 1885
 - editing configuration options in 1884
 - for the Remote Agent for Macintosh Systems 1917
 - reassigning how slots appear 472
 - Recently Written Media report 760
 - recipients
 - assigning alert categories 680–681
 - configuring groups 677
 - configuring MAPI e-mail 668
 - configuring Net Send 674
 - configuring pager 671, 675, 1731
 - configuring SMTP 667
 - configuring VIM e-mail 670
 - defined 666

- recipients (DLO)
 - configuring groups 1732
 - configuring MAPI e-mail 1729
 - configuring Net Send 1731
 - configuring pager 1730
 - configuring SMTP 1729
 - configuring SNMP Trap 1731
 - configuring VIM e-mail 1730
 - defined 1728
- Recovered Jobs custom error-handling rule 591
- recovering jobs
 - threshold for 596
- recovery password
 - setting in DLO 1616
- recovery points
 - default interval in Exchange 1122
 - in Exchange
 - overview 1122
 - setting in Exchange 1135
- recovery requirements in IDR 1828
- Recovery Storage Group 1147
- redirected restore
 - Exchange data 1159
 - Microsoft Virtual Hard Disk files 636
 - Microsoft virtual machine 1184
 - SAP data 1373
 - using to install Domain Controllers from Media 636
 - VMware virtual machines 1404
- redirecting scheduled jobs 520
- Remote Administrator
 - installing using the command line 162
 - running 151
- Remote Agent for Deduplication 1592
 - viewing properties 1594
- Remote Agent for Linux or UNIX Servers
 - about backing up 1891
 - about exclusions from backup 1884
 - about publishing to media servers 1882
 - backing up Novell OES components 1896
 - backup job options 1892
 - beoper group, defined 1880
 - configuration options in the ralus.cfg file 1885
 - configuring the ralus.cfg file 1881
 - creating the beoper group 1880
 - default options 1899
 - Editing configuration options in the ralus.cfg file 1884
 - editing default options 1899
- Remote Agent for Linux or UNIX Servers *(continued)*
 - installing 1877
 - manual install and uninstall 1925
 - Novell OES, requirements for backup 1895
 - publishing to media servers 1883
 - push-installing 1877
 - requirements 1876
 - restore job options 1898
 - restoring 1897
 - runtime scripts 1905
 - saving installation log 1877
 - setting backup job properties 1892
 - starting the Remote Agent daemon 1907
 - stopping the Remote Agent daemon 1907
 - troubleshooting 1908
 - uninstalling 1903
 - uninstalling manually 1904
 - using SSH during push-install 1877
- Remote Agent for Macintosh
 - about publishing to media servers 1882
- Remote Agent for Macintosh Systems
 - about the ralus.cfg file 1917
 - backup job options 1892
 - configuration options in the ralus.cfg file 1885
 - default options 1899, 1919
 - editing backup job options 1918
 - Editing configuration options in the ralus.cfg file 1884
 - editing default options 1899, 1919
 - installing 1914
 - manually starting 1924
 - manually stopping 1924
 - publishing to media servers 1883
 - requirements 1912
 - restore job options 1898
 - restore options 1918
 - restoring 1897, 1919
 - setting backup job properties 1892
 - supported backup methods 1917
 - troubleshooting 1926
 - uninstalling 1923
- Remote Agent for NetWare Servers
 - AUTOEXEC.NCF file 1934
 - backing up
 - BEDIAG.NLM utility 1943
 - create BEDIAG.FAX 1944
 - decompressed files 1942
 - password database 1935
 - rights for backup 1935

Remote Agent for NetWare Servers *(continued)*

- backing up *(continued)*
 - single server backup strategies 1936
 - strategies for multiple administrators 1937
 - strategies for single administrator 1937
- backing up NetWare servers 1938
- backup options 1939
- creating Advrtms.dat file 1933
- default options 1940–1941
 - setting 1941
- installing 1931
- overview 1929
- publishing NetWare servers 1933
- restoring 1939
 - overview 1939
- system requirements 1930–1931
- TCP dynamic port ranges
 - specifying 1943

Remote Agent for Windows Systems

- hardware requirements 1946
- installing 139
- installing in an Active Directory network 140
- installing on a Microsoft cluster 818
- installing on a VERITAS cluster server 844
- installing to a remote computer in the backup
 - selections list 140
- installing using a command script 148
- installing using the command line 146
- license keys 1946
- publish to media servers 1951
- Remote Agent Utility 1948
- stopping and starting 1947
- uninstalling using a command script 149
- uninstalling using the command line 148

Remote Agent Utility

- activity status
 - viewing 1950
- command line applet 1958
 - switches 1959
 - using 1958
- database access
 - configuring 1955
 - options 1956
- DB2 archive logs job template name
 - adding 962
- DB2 DBA-initiated job template name
 - adding 962
- DB2 instance
 - configure for database access 1324

Remote Agent Utility *(continued)*

- DB2 instances
 - configuring for database access 960
- default publishing interval 1953
- Event Viewer
 - open 1948
- job template name for DBA-initiated jobs 1325
- Linux
 - configure Oracle instance on 1328
- port
 - configure for DB2 and Oracle
 - operations 1329
- publish to media servers 1951, 1953
- publishing
 - adding media servers 1952
 - editing media server information 1954
 - removing media servers 1955
- publishing options 1953
- Real Application Cluster (RAC)
 - publish to media server 1324
- refresh interval 1950
 - setting 1951
- Registry Editor
 - open 1948
- Services
 - open 1948
- start the utility at log on 1950
- starting 1949
- starting automatically 1951
- status options 1950
- update credentials for Linux instances 1325
- view status 1949
- Windows
 - configure Oracle instance on 1324

Remote Media Agent for Linux Servers

- adding to Backup Exec database 1972
- backing up data 1978
- beoper group 1968
- changing port number 1973
- creating a simulated tape library 1980
- deleting simulated tape library 1983
- determining server status 1976
- how it works 1966
- ICMP ping 1973
- installing 1968
- managing simulated tape libraries from the
 - command line 1984
- overview 1966
- requirements 1967

Remote Media Agent for Linux Servers *(continued)*

- restoring data 1978
- Tape Library Simulator Utility 1979
- troubleshooting 1987
- uninstalling 1985
- using with SAN Shared Storage Option 1973
- viewing properties 1975
- viewing properties of simulated tape libraries 1981
- remote storage
 - backing up 351
- removable backup-to-disk folder requirements 498
- renaming
 - media labels 243
 - vault 254
- Renewal Assistant
 - about Backup Exec maintenance contract information 177
 - Maintenance contract information 177
 - running the 178
 - viewing maintenance contract information 179
- Renewal Assitant
 - serial number and license key formats 177
- replicated catalog
 - in CASO 1542
- reports
 - active alerts 729
 - Active Alerts by Media Server 730
 - alert history 730
 - application event log 732
 - archive job success rate 768
 - archive selections by archive rules and retention categories 768
 - audit log 732
 - available in Backup Exec 722
 - Backup Job Success Rate report 733
 - Backup Recommendations report 733
 - Backup Set Details by Resource 734
 - backup sets by media set 735
 - Backup Size by Resource 736
 - configuration settings 737
 - Daily Device Utilization 738
 - Deduplication 739
 - deleting in job history 698
 - device summary 740
 - Device Usage by Policy 741
 - Event Recipient 743
 - exchange mailbox group archive settings 769

reports *(continued)*

- failed archive jobs 770
- Failed Backup Jobs 744
- file system archive settings report 770
- Job Distribution by Device 745
- Jobs Summary 745
- Machines Backed Up 746
- Managed Media Servers 747
- media audit 748
- media errors 749
- Media Required for Recovery 749
- media set 750
- media vault contents 751
- Missed Availability 752
- Move Media to Vault 752
- operations overview 753
- overnight archive summary 771
- overview 690
- Policy Jobs by Resource Summary 756
- Policy Jobs Summary 757
- Policy Properties 758
- Policy Protected Resources 759
- Problem Files 759
- Recently Written Media 760
- Resource Backup Policy Performance 761
- Resource Risk Assessment 761
- Restore Set Details by Resource 762
- Retrieve Media from Vault 763
- running 691
- running job 694
- saving 696
- Scheduled Server Workload 765
- scheduling report jobs 698
- Scratch Media Availability 766
- setting notification recipients 698
- test run results 767
- vault store usage details 772
- vault store usage Summary 773
- viewing 691
- viewing properties 721
- reports (DLO) 1734
 - running 1736
 - viewing 1734
- requirements
 - Agent for Microsoft Hyper-V 1171
 - Backup Exec 114
 - Central Admin Server Option 1508
 - Exchange Agent 1095
 - Lotus Domino Agent 1064

- requirements (*continued*)
 - NDMP Option 1848
 - Remote Media Agent for Linux Servers 1967
 - SAN Shared Storage Option 1993
- requirements for end users
 - Backup Exec Retrieve 868
- requirements for installing
 - Backup Exec Retrieve 866
- reset accounts option
 - in DLO Desktop Agent 1762
- Reset Cleaning Statistics 465
- reset dialogs option
 - in DLO Desktop Agent 1762
- resource
 - credentials
 - changing for restore job 630
 - order of 338
- Resource Backup Policy Performance report 761
- resource discovery
 - about 316
 - creating job 316
 - used with Exchange Agent 1096
- Resource Risk Assessment report 761
- responding to active alerts 654
- Restore dialog box
 - DLO Desktop Agent 1787
- restore job
 - Advanced File Selection 610
 - advanced options 613
 - canceling 641
 - copying to another server 554
 - creating for the Remote Media Agent for Linux Servers 1978
 - creating through dialog boxes 605
 - creating using the wizard 604
 - file permissions 619
 - file redirection options 634
 - general options 612
 - Hyper-V host 1182
 - Lotus Domino options 1082
 - over existing files 612
 - pre/post commands 352, 399, 618
 - preserve tree option 613
 - redirecting 634
 - redirecting Lotus Domino 1083
 - redirecting SAP data 1373
 - security 612
 - Selections options 610
 - setting defaults 638
- Restore Job Properties dialog box
 - SQL Agent 1289
- Restore Set Details by Resource report 762
- restore to named instance 1306
- Restore Wizard
 - configuring to launch from the Restore button 605
 - launching 604
 - preventing from launching from the Restore button 604
- restoring
 - about restoring data 599
 - ARCserve tapes 625
 - byte count does not match 797
 - creating a selection list 628
 - encrypted data 422
 - Exchange data 1144
 - files
 - using DLO Administration Console 1707
 - using DLO Desktop Agent 1786
 - Lotus Domino Agent 1079
 - media created with other backup software 624
 - media view 600
 - Microsoft clusters
 - cluster quorum for Windows 2000 and Windows Server 2003 840
 - Remote Agent for NetWare Servers volume restrictions 1943
 - resource view 600
 - searching for files 631
 - selecting data 626
 - selection list options 628
 - SQL master database 1301
 - using the media view 627
 - using the resource view 626
 - with Backup Exec Retrieve 1790
- restricted encryption keys
 - defined 417
- restricted logon account 186
- retensioning a tape 484
- retention categories
 - editing 1458
 - overview 1457
- retention category
 - editing default retention category 1495
 - specifying properties 1459
- retention periods for archived items 1457
- retired media
 - defined 221

retired media *(continued)*

- moving damaged media 259

- Retrieve Media from Vault report 763

- returning to a previous configuration 778

revisions

- deleting automatically in DLO 1667

- number to keep

- setting in DLO 1664

- Revision Control tab

- DLO 1664

- revisions (DLO) 1663

RMAN

- backing up with SAP Agent 1373

- restoring with SAP Agent 1375

- using SAP Agent with 1363

- using to protect SAP for Oracle databases 1362

robotic library

- cleaning slot 471

- configuring partitions 476

- creating partitions 475

- example configuration 469

- importing media 489

- initialization on startup 471

- lock front panel 493

- portal support 490

- problem with not displaying 792

- redefining partitions 478

- setting up hardware 468

- unlock front panel 494

- using with Backup Exec 467

robotic library properties

- Configuration tab 471

runtime scripts, for Remote Agent for Linux or UNIX

- Servers 1905

S

SAN Shared Storage Option

- using Remote Media Agent for Linux Servers with 1973

SAN Shared Storage Option (SSO)

- best practices 2014

- catalog media 1997

- changing the configuration 2005

- device allocation 1996

- device operations 2003

- drive pools

- how to use with SSO 2004

- environment reconfiguration 2009

- hardware errors 2012

SAN Shared Storage Option (SSO) *(continued)*

- installing 1994

- media rotation 1996

- monitoring drives 2005

- NetWare media servers

- configuring for robotic library sharing 2000

- overview 1991

- renaming

- libraries and drives in shared

- environment 2003

- requirements 1993

- resetting the SAN 2013

- robotic library sharing 1998

- prerequisites 1999

- scheduling jobs 1997

- services

- starting and stopping on multiple

- servers 2009

- sharing media 1997

- standby primary database server

- creating 2007

- troubleshooting 2010

- offline devices 2010

- viewing media 2005

- Windows media servers

- configuring for robotic library sharing 2000

SAP Agent

- about disaster recovery 1380

- backing up with RMAN 1373

- before backing up 1367

- database

- system level backup jobs 1371

- database server

- restoring remote server 1381

- disaster recovery requirements 1381

- features 1362

- how it works 1362

- installing 1365

- overview 1362

- privileges 1365

- requirements 1364

- restoring with RMAN 1375

- security 1365

- submitting jobs from remote computers 1372

- using for backups and restores 1371

SAP Agent catalog

- manually migrating 1376

- migrating from _backint.mdb to

- _backint.xml 1376

- SAP Agent catalog (*continued*)
 - preserving integrity 1366
 - restoring 1367
- scan
 - detecting storage arrays 2044
- schedule
 - exclude dates from 366
 - setting the effective date 363
- scheduled jobs
 - about 356
 - about time windows 364
 - changing priority 570
 - configuring 356
 - configuring default options 366
 - deleting 571
 - hold 568
 - list of statuses 565
 - removing hold 569
 - running immediately 568
 - running on a day interval 362
 - running on recurring days of the month 361
 - running on recurring week days 360
 - running on specific days 359
 - running test job 570
 - setting the time window 364
- Scheduled Server Workload report 765
- scheduling
 - backup jobs in Desktop Agent 1773
 - calendar 359
 - notification recipients 677
 - report jobs 698
 - setting for templates 532
- scratch media
 - creating 232
 - defined 221
- Scratch Media Availability report 766
- SCSI
 - address for devices 468
 - information about devices 463
 - setting address for robotic library drives 468
- SCSI bus
 - configuring for tape devices in a Microsoft cluster 829
- SCSI pass-through mode
 - setting for devices 463
- search
 - catalogs 631
 - log file history 1720
- Search Knowledge Base 91
- Section 508 of the Rehabilitation Act
 - compliance 2063
- secured file access 862
- security
 - changing for Windows systems 108
 - restoring 612
- security options
 - configuring 407
 - setting defaults 404
- selecting devices and data to back up 280
- selection lists
 - about creating jobs 544
 - about priority and availability 306
 - copying 302
 - copying to another server 554
 - creating 296
 - creating a custom filter 310
 - creating jobs 545
 - creating separate for each computer or resource 309
 - defined 295
 - deleting 303
 - editing 304
 - excludes 305
 - filtering 313
 - holding jobs that back them up 303
 - merging 300
 - notification 295
 - replacing 301
 - searching 314
 - sending notification when used in jobs 682
 - setting default priority and availability 307
 - setting priority and availability 307
 - viewing 546
 - viewing history 314
 - viewing summaries 315
- selections
 - Desktop Agent
 - backup 1765
 - DLO
 - backup 1658
 - options for backup job 336
 - user-defined
 - adding 291
 - deleting 292
 - using TCP/IP addresses 291
- selections lists
 - set up notification 297

- Serial numbers
 - maintenance contract expiration dates 179
 - viewing maintenance contract information 179
- server list
 - adding and removing servers 168
 - manually update for NetWare 1942
- server properties
 - viewing 212
 - viewing for Remote Media Agent 1975
- service account
 - about 106
 - changing 107, 168
- service credentials in DLO
 - managing 1618
- service state 320
- services
 - starting and stopping 168
- Set Application Defaults
 - Preferences 197
- Set Remote Agent Priority 351
- setting notification recipients for reports 698
- SGMon 810
- Shadow Copy Components
 - about restoring 622
 - file system 320
- Share Your Ideas, described 91
- shared catalogs
 - using 1992
- SharePoint Agent
 - about 1190
 - about restoring SharePoint 2003 resources 1237
 - about restoring SharePoint Foundation 2010 resources 1203
 - about restoring SharePoint Server 2007 resources 1221
 - about restoring SharePoint Server 2010 resources 1203
 - about restoring SharePoint Services 3.0 resources 1221
 - adding a server farm 1191, 1219
 - automatic protection for SharePoint 1194
 - backing up a Microsoft Office SharePoint 2007 Server 1219
 - backing up a Microsoft Office SharePoint 2010 Server 1202
 - backing up a Windows SharePoint Foundation 2010 farm 1202
 - backing up a Windows SharePoint Services 3.0 farm 1219
- SharePoint Agent (*continued*)
 - backing up individual SharePoint 2007 web applications 1220
 - backing up individual SharePoint 2010 web applications 1203
 - backing up SharePoint Portal Server 2003 resources 1236
 - backup options 1246
 - changing the default name of a farm 1195
 - default options 1197
 - deleting a farm 1195
 - disabling or enabling communication between Web servers and Backup Exec 1196
 - installing 1191
 - overview 1190
 - redirecting individual SharePoint 2003 objects to a file path 1243
 - redirecting individual SharePoint 2007 objects to a file path 1228
 - redirecting individual SharePoint 2010 objects to a file path 1210
 - redirecting restore jobs for SharePoint 2003 document library data 1242
 - redirecting restore jobs for SharePoint Portal Server 2003 1241
 - redirecting restore jobs for SharePoint Server 2007 1227
 - redirecting restore jobs for SharePoint Server 2010 1208
 - redirecting the restore of a SharePoint 2003 database to an alternate SQL instance 1244
 - redirecting the restore of a SharePoint 2007 database to an alternate SQL instance 1230
 - redirecting the restore of a SharePoint 2010 database to an alternate SQL instance 1213
 - redirecting the restore of SharePoint 2007 web applications 1229
 - redirecting the restore of SharePoint 2010 web applications 1211
 - redirection options 1250
 - requirements 1190
 - restore options 1247
 - restoring a Microsoft Office SharePoint Server 2007 Shared Services Provider 1224
 - restoring a SharePoint Server 2007 Web application 1225
 - restoring a SharePoint Server 2010 Web application 1207

- SharePoint Agent *(continued)*
 - restoring individual SharePoint 2003 objects 1238
 - restoring individual SharePoint 2007 objects 1223
 - restoring individual SharePoint 2010 objects 1205
 - restoring SharePoint 2003 document libraries 1240
 - restoring SharePoint 2003 documents from document library backups 1241
 - restoring SharePoint 2003 resources 1238
 - restoring SharePoint Foundation 2010 resources 1204
 - restoring SharePoint Server 2007 resources 1222
 - restoring SharePoint Server 2010 resources 1204
 - restoring SharePoint Services 3.0 resources 1222
 - selecting SharePoint Portal Server 2003 resources for backup 1236
 - server farms options 1193
 - setting default options for SharePoint Server 2003 and later 1196
 - system requirements 1190
 - using with SharePoint Portal Server 2003 and Windows SharePoint Services 2.0 1235
 - using with SharePoint Server 2007 and Windows SharePoint Services 3.0 1218
 - using with SharePoint Server 2010 and Windows SharePoint Foundation 2010 1200
- sharing a backup-to-disk folder 506
- sharing media 1992
- silent install of DLO 1612
- silent mode installation 102, 153
- Silverlight
 - deploying in your organization 869
- simple recovery model
 - SQL 2000 1255
- simulated tape library
 - creating 1980
 - delete 1983
 - viewing properties 1981
- single block mode
 - setting for devices 462
- single instance backup for NTFS 349
- size
 - cluster 796
 - slot base configuration 472
- Small Business Server Edition of Backup Exec described 67
- SMTP
 - configuring e-mail recipients 667
 - e-mail notification method 662
 - e-mail notification method in DLO 1728
- Snap Start
 - for VSWF volumes 945
- snapshot technology
 - used in Advanced Open File Option 937
 - using with Exchange Agent 1108
- SNMP
 - configuring notification 683
 - configuring system service for Windows 2000 687
 - installing WMI provider 688
 - object identifier prefix 683
 - traps
 - defined 683
 - WMI 687
- SNMP Trap
 - configuring e-mail recipients 1729
 - configuring recipients 1731
 - notifications 1731
- son media rotation strategy 265
- Specified Backup Network feature
 - configuring 404
 - described 402
- specify a date and time for restoring named transaction 1293
- splash screen
 - show at startup 197
- SQL 2000
 - filegroups 1298
 - Guide Me wizard 1280
 - loading state 1278
 - named transaction 1297
 - include 1293
 - restoring to a 1293
 - physical check after backup 1260
 - primary and nonprimary filegroups 1294
 - restoring
 - to named instance 1306
 - simple recovery models 1255
 - standby database 1254
 - standby mode 1278
- SQL Agent
 - ADBO 1264

SQL Agent *(continued)*

- AOFO 1263
- backing up
 - backup methods 1274
 - consistency check after backup 1278
 - consistency check recommendations 1260
 - databases 1281
 - filegroups 1284
 - strategies for 1258
 - transaction logs 1286
 - Windows registry 1259
- backing up SQL databases 1283
- backing up SQL filegroups
 - overview 1282
- Backup Job Properties dialog box 1273
- consistency check 1261
 - recommendations 1260
- Database Consistency Check (DBCC)
 - recommendations 1260
- database snapshots
 - backup method 1287
 - creating 1288
 - overview 1286
- default options 1265, 1273
 - setting 1273
- disaster recovery 1312
 - how to prepare 1311
 - manual 1312
 - overview 1310
 - requirements 1311
- displaying filegroups 1284
- features 1254
- installation 881, 1256
- logon accounts 1256
- overview 1254
- redirection options 1306
- requirements 1255
- restore options 1289
 - setting 1289
- restoring
 - automate master database restore 1290
 - create standby database 1290
 - database files to target instance 1308
 - filegroups 1298
 - Guide Me wizard 1293
 - master database 1301
 - point in time log restore option 1296
 - redirecting restores 1305
 - Redirection dialog box 1305

SQL Agent *(continued)*

- restoring *(continued)*
 - Restore Job Properties dialog box 1289
 - specify date and time for named transaction 1293
 - very large databases 1294
- snapshot technology
 - using 1262
- strategy recommendations 1259
- transaction logs 1285
- truncate log on checkpoint option 1285
- SQL database backups
 - restoring from 1295
 - TDE-encrypted database backups 1295
- SQL Server 1607
- SSH (Secure Shell), using to push-install the Remote Agent for Linux or UNIX Servers 1877
- staging data 548
- stalled job status
 - threshold for 595–596
- standby database
 - creating
 - SQL Agent 1290
 - SQL 2000 1254
- starting
 - Desktop and Laptop Option (DLO) 1638
- statistical properties for media 263
- statistics
 - device usage 463
 - devices
 - since cleaning 465
- storage
 - sharing 444
- storage array
 - about identifying physical disks 2045
 - blink 2045
 - capacity 2023
 - configuring 2021–2022
 - configuring virtual disks 2034
 - detecting 2044
 - hardware health 2023
 - hardware status 2023
 - identifying the physical disks 2046
 - renaming 2044
 - status 2023
 - viewing components 2020
 - viewing physical disk properties 2026
 - viewing properties 2023

- Storage Array Configuration Wizard
 - changing or adding hot spares 2043
 - configuring a storage array 2021–2022
 - description 2021
- storage devices
 - about 441
 - installing 103
 - pausing 446
 - pausing a media server 446
 - renaming 447
 - resuming 447
 - resuming a media server 446
- storage limits for user data
 - DLO 1644
- Storage Location (DLO) 1676
 - defined 1604
 - deleting 1682
- Storage Provisioning Alert
 - configuring 2047
 - described 2046
- Storage Provisioning Option
 - configuring alerts for disk space usage 2047
 - configuring in CASO 2019
 - description 2018
 - detecting storage arrays 2044
 - installing 2020
 - predicting disk space usage 2046
 - requirements 2019
 - upgrading 2020
- Symantec Backup Exec 2010
 - described 64
- Symantec Device Driver Installation Wizard 455
- Symantec Endpoint Protection
 - using with Backup Exec 408
 - viewing a summary of 590
- Symantec Knowledge Base 801
 - searching 802
- Symantec Online Storage folder
 - about 2054
 - creating 2054
 - deleting 2061
 - pausing 2056
 - properties 2055
 - resuming 2056
 - sharing 2057
- Symantec Online Storage for Backup Exec
 - about 2051
 - about creating duplicate backup jobs 2057
 - about managing jobs 2060
- Symantec Online Storage for Backup Exec
 - (continued)*
 - about restoring jobs 2062
 - about Symantec Online Storage folders 2054
 - best practices 2052
 - creating a Symantec Online Storage folder 2054
 - creating duplicate backup jobs 2058
 - deleting Symantec Online Storage folders 2061
 - downloading the Symantec Online Storage for Backup Exec Protection Agent 2054
 - erasing files 2060
 - pausing a Symantec Online Storage folder 2056
 - resuming a Symantec Online Storage folder 2056
 - setting up 2053
 - sharing an existing Symantec Online Storage folder 2057
 - signing up 2053
 - Symantec Online Storage folder options 2055
- Symantec Online Storage for Backup Exec Protection Agent 2054
- Symantec Volume Snapshot Provider
 - changing defaults 945
 - with AOFO 949
- synchronization
 - defined in DLO 1756
- synchronization (DLO)
 - create new sets 1779
 - delete synchronized folder 1781
 - overview 1778
- Synchronized Selections view
 - in DLO Desktop Agent 1778
- synchronizing archiving permissions and settings 1493
- synthetic backup
 - baseline 899
 - creating 904, 906
 - encryption
 - requirements for 901
 - requirements 901
 - template rules 910
 - example 906
- system logon account 189
 - creating 193
- system requirements
 - Backup Exec 114
- System State
 - about 620
 - restoring 621

- ## T
- tabbed dialog boxes
 - navigation 2072
 - Tape Device Configuration Wizard 453
 - tape devices, configuring 453
 - Tape Library Simulator Utility
 - creating a simulated tape library 1980
 - deleting library 1983
 - overview 1979
 - running from command line 1984
 - viewing properties 1981
 - tapeinst.exe
 - Symantec Device Driver Installation Wizard 455
 - tapes
 - DLT tape drive 794
 - target domain
 - defined 110
 - task pane
 - overview 94
 - tasks available in Backup Exec Retrieve 864
 - TCP/IP
 - adding for user-defined shares 291
 - required for RAMS agent 1912
 - TDE
 - Transparent Database Encryption 1295
 - technical support
 - contacting 802
 - templates
 - about export media templates 536
 - about template rules 540
 - about verify backup sets templates 533
 - adding a duplicate backup sets template to policies 550
 - adding a verify backup sets template to a policy 534
 - adding an export media template to a policy 537
 - adding backup templates to a policy 530
 - backup template file exclusions 532
 - changing template rules 542
 - deleting from a policy 539
 - deleting template rules 543
 - duplicate backup sets template overview 548
 - editing in a policy 539
 - importing into a policy 538
 - overview 521
 - setting template rules 542
 - setting the schedule 532
 - using in policies 529
 - test run job
 - about 385
 - creating 386
 - defined 386
 - general properties 386
 - running for scheduled job 570
 - setting defaults 387
 - Test Run Results report 767
 - testing logon accounts 337
 - ThreatCon levels 408
 - time windows
 - about 364
 - setting 364
 - transaction logs
 - backing up
 - SQL Agent 1285
 - Lotus Domino DBIID 1074
 - overview
 - Lotus Domino 1074
 - recycling
 - Lotus Domino 1071, 1077
 - viewing Lotus Domino 1068
 - Transparent Database Encryption
 - TDE 1295
 - trial version
 - agents and options 166
 - installing Backup Exec 121
 - troubleshooting
 - Backup Exec performance
 - improving 797
 - backup issues 795
 - clusters 857
 - error messages 794
 - hardware-related issues 789
 - Remote Media Agent for Linux Servers 1987
 - restore issues 797
 - true image restore
 - CASO and 916
 - creating a policy for 916
 - icons 918
 - overview 912
 - requirements 915
 - troubleshooting 918
 - true image catalogs 916
 - truncate log on checkpoint option
 - SQL Agent 1285
 - typical installation 116

U

- unconfigured virtual disk
 - configuring 2034
 - hardware health 2035
 - hardware status 2035
 - viewing properties 2035
- uninstalling 876
 - Backup Exec 169
 - Backup Exec from a Microsoft cluster 820
 - using command line 164
 - Windows Management Instrumentation SNMP provider 688
- Unique Message Identifier (UMI) error code
 - viewing 577, 658
- unlock
 - robotic library panel 494
- Update vault using wizard 257
- updating
 - DLO 1632
- upgrading from previous versions of Backup Exec
 - Retrieve 869
- USB tape devices
 - reconnecting 453
- use case 862
- user data folder
 - defined 1756
- User Data Folder in DLO 1676
- User Properties dialog box in DLO 1699
- user-defined media vault 251
 - adding media to 255
 - creating 251
- user-defined selections
 - about 290
 - adding 290
 - deleting 292
- users
 - access
 - disabling/enabling in DLO 1700
 - adding in DLO 1698
 - Desktop Agent access via profile 1755
 - import in CSV file in DLO 1699
 - managing in DLO desktop agent 1696
 - profile
 - defined 1756
 - properties
 - changing in DLO 1699
 - removing from DLO 1700
 - viewing in DLO 1703

- using RMAN to protect SAP for Oracle
 - databases 1362
- utility jobs
 - overview 480
- utility partitions
 - about performing redirected restores 623
 - backing up in IDR 1815
 - restoring 623
 - selecting for backup 283

V

- vault
 - drag and drop
 - to move media 258
 - finding media 254
 - moving media 257
 - moving media to 254
 - scan bar code labels to move media 255
 - scheduling a job to move media 255
- vault rules for media sets 252
- vault store
 - assigning 1439
 - changing item deletion mode 1448
 - creating 1446
 - deleting 1450
 - deleting items after archiving 1446
 - deleting items after vault store backup 1446
 - editing properties 1447
 - viewing status 1448
- vault store group
 - backing up for Archiving Option 1477
- vault store partition
 - creating 1451
 - editing properties 1452
 - viewing open and closed states 1452
- vault store partitions
 - backing up for Archiving Option 1477
 - open and closed 1451
 - overview 1451
- Vault Store Usage Details 772
- Vault Store Usage Summary Report 773
- vault stores
 - backing up for Archiving Option 1477
 - fingerprint database 1445
 - overview 1445
- Vault wizard 257
- verify
 - after backup completes 347
 - after duplicate backup completes 380, 911

- verify backup sets template
 - about 533
 - adding to a policy 534
- verify job
 - creating 382
 - defined 382
- VERITAS clusters
 - Windows 2003/2008 847
- VHD files
 - about managing 293
 - about redirecting restore jobs to VHD files 636
- view history 314
- viewing
 - alert job log 653
 - alert properties 651
 - Automated User Assignment properties in
 - DLO 1686
 - job monitor 557
 - job workload from the calendar 588
 - Lotus Domino databases 1068
 - Lotus Domino transaction logs 1068
 - users in DLO 1703
- VIM
 - configuring recipients 670, 1730
 - e-mail notification method 665, 1728
- virtual disk 2035
 - See also* unconfigured virtual disk
 - blink 2045
 - capacity 2038
 - concurrent jobs 2038
 - configuring 2034
 - creating 2021–2022
 - description 2029
 - editing default options 2030
 - editing general properties 2037
 - editing global defaults 2032
 - editing low disk space thresholds 2030, 2032
 - hardware health 2038
 - hardware status 2038
 - identifying the physical disks 2046
 - number of files 2038
 - renaming 2044
 - status 2038
- Virtual Disk Service
 - installing for the Storage Provisioning
 - Option 2019
- virtual machines
 - automatic protection for Hyper-V 1174
 - backing up with Agent for VMware 1390

- virtual servers
 - backing up in a Microsoft cluster 838
 - backing up in a VERITAS cluster 849
- virtual tape library
 - DirectCopy to physical devices 381–382
- virus
 - effect on data storage requirements 272
- VMware vCenter Server, adding 1387
- VMware vCenter Server, deleting 1388
- volume level backups
 - automatic exclusion of SQL data 1280
- volume restrictions
 - Remote Agent for NetWare Servers 1943
- VSS
 - perform consistency check before Active
 - Directory backup 887
 - perform consistency check before Exchange
 - backup 1137
 - using to protect Exchange data 1108
- VSS Provider
 - protecting databases and applications 1398

W

- Window menu 2066
- Windows Automated System Recovery (ASR) files
 - in IDR
 - defined 1811
- Windows Change Journal
 - option to use for backup job 346
 - using to determine backed up status 280
- Windows Management Instrumentation (WMI)
 - adding WMI capability 687
- Windows registry
 - backing up with SQL Agent 1259
- Windows Server 2003
 - backing up 320
- Windows Server 2008
 - backing up 320
 - Read Only Domain Controller 114
 - Server Core 114
- wizard for DLO
 - configuration 1640
- WMI
 - installing performance counter provider 687
 - installing SNMP provider 688
 - performance counters 687
 - uninstalling SNMP provider 688
- working set
 - advantages and disadvantages 279

working set (*continued*)
 backups
 defined 277