Enterprise Vault

Enterprise Vault Backups

Best Practise

Version 2.2

Steve Blair
Sr. Regional Product Manager

David Smiley
Sr. Principal Business Critical Engineer

Patti Rodgers
Sr. Principal Business Critical Engineer

Aidan Finley
Sr. Regional Product Manager



September 2008

Table of Contents

Overview

Introduction

This best practise document discusses the inherent needs to securely backup operations Enterprise Vault servers to insure best available DR scenarios. With several very important information sources integrated with Enterprise Vault, it is critical that proper backups of this data be performed on a regular basis. This paper considers the impact this has on Enterprise Vault ways customers can easily implement a high degree of coverage for their archived content. It is important to understand that it is a Best Practise to regularly backup the components of your Enterprise Vault environment to insure that you have adequate DR protection as with any Tier-1 business critical application.

Scope of Document

This document is focused on Enterprise Vault, Microsoft SQL, Vault Store settings, and the indexes that are part of the environment as a whole. This document presumes a good working knowledge of your company backup products, schedules, and DR plans. The implementation, planning and use of actual backup software / hardware solution are outside the scope of this document however Best Practise recommendations to insure your DR plans protect your data in Enterprise Vault for use of your backup software are made. This document does not cover the backups of your server environment, but we encourage you to review and make sure as part of your Best Practises that your company uses. In the appendices to updated document are significant coverage to Symantec NetBackup as well as Symantec Backup Exec which are available options to insure your confidence in your investment with Symantec.

This document does not cover strategies for backing up Enterprise Vault File System Archiving Placeholders; for additional information regarding these files, please refer to the Enterprise Vault File System Archiving documentation.

Target Audience

This document is aimed at customers, consultants and support staff and it is assumed the reader has a good understanding about the architecture and operational aspects of an Enterprise Vault server, and their internal network and storage architecture. This document also discusses concepts related to Microsoft Server management, and expects the reader to either be skilled in this area, or have team member(s) who are.

Acknowledgements

We would like to acknowledge the contribution that other individuals made towards making this a successful and informative document. Contributions and feedback came from the following teams: Regional Product Management, Sr. Product Management, Technical Field Enablement, Engineering, Consulting, Business Critical Services, and our Customer Support Teams.

Backup Trends

Companies implementing Enterprise Vault typically purchase to solve one of several needs. They are either trying to offset the explosive growth of individual mailboxes, file server personal / shared areas, or insure relevant regulatory compliance in their country. Backing up your data to insure your company has an adequate safety net has long been the bane of the modern System Administrator, but lack of backups place a company at extreme risk in a DR (Disaster Recovery) situation. Because of implementing Enterprise Vault, a company must undertake a review of their existing backup schedules and plans, as archiving allows customers more flexibility and newer ways to have a safety copy.

Smaller companies with smaller Enterprise Vault environments have relatively straightforward backup requirements. The complexity of the backup requirement grows as the number of servers and amount of stored data grows. Given the advances in both large storage tape technologies such as DLT^2 and LTO, and disk systems that operate like tape or replicate, no two customers studied are performing their backups in a similar fashion. The one unifying similarity of Enterprise Vault customers regardless of size is that they are aware they have to insure for their internal &/or external customers a high degree of data availability. Both Symantec Backup Exec and Symantec NetBackup for the SME and Enterprise customers allow you great ways to maximise your backup coverage.

Enterprise Vault Infrastructure and your environment

As the landscape changes with the introduction of Enterprise Vault into your environment, it does not absolve one of the responsibilities to insure that adequate backups are in place. Given the data contained in Enterprise Vault is still just as valuable, just as important, it is Best Practise that you treat your Enterprise Vault server(s) as any other Tier-1 application server in your network. Because of installing Enterprise Vault at your site, you will observe (in many cases) greatly shortened backup times for the server(s) that Enterprise Vault archives from, but you will still need to make sure your DR position is strong by the scheduled backups of your Enterprise Vault data.

Backups vs. Archiving

Let's begin by reviewing that many people are confused over the two terms "backup" vs. "archiving" and often wonder if there is any discernable difference. After all, is a backup not an archive in time, or a backup stored offsite considered an archive? It is absolutely true that both views are valid, but the line becomes clear when it is time to produce contact back from a backup. Backup media is not the same as a hard disk spinning online, and as such can and does suffer from media failure, inability to read / restore content, and data loss of the physical media. For the purposes of this paper, we focus on Archiving and the relationship with a sound backup strategy in a way we call "active archiving". Active archiving differs from regular backups as it has indexed and represented the content in a way immediately available for recall, review and use without operator interaction. The very activity of having the content still stored, managed, and available replaces many mundane

System Administrator tasks, and supplements end-user experience of their data being highly-available via "active archiving" as delivered with Enterprise Vault. There are two major Best Practises you have available to perform a backup of your important information stored in Enterprise Vault: data-only backups of the archived content, and data and application backups of the whole Enterprise Vault environment. We shall review both in this paper, and allow you as the reader to take a decision on which will be the Best Practise for your environment.

Enterprise Vault Backup Requirement Overview

There are several crucial components in your Enterprise Vault deployment that need to be backed up periodically. Whilst having Enterprise Vault as your archiving solution greatly reduces your backup reliance, it does not absolve you of performing backups.

The items in your environment that must be backed up are:

- SQL Databases (Directory, Audit, Monitoring and vault Stores)
- Indexes, their directories and files
- Vault Store partition (s) (open AND closed partitions)
- EV Server system for registry settings
- EV Server(s) license file (per EV server!)
- Any customisations done to your EV environment:
- Scripts
- Welcome Messages
- Webapp.ini
- PST Migration Templates

Each of the above components being backed up plays a vital role in the Enterprise Vault environment as a whole. It is a Best Practise to insure that good backups of the above environment are performed on a regular basis as part of your total DR plan. The diagram below shows a sample Enterprise Vault deployment and the interdependencies used by the various components in a simple deployment. Whilst this diagram may not accurately reflect the environment you have Enterprise Vault deployed in at your site, it does cover the important components of the Enterprise Vault environment.

The theories and Best Practise guidance done in this paper are the results of years of testing, and actual customer experience to help guide you on making your environment very reliable. Part of the Best Practise of backing up your Enterprise Vault environment is to test the backups periodically to insure that your backups are good. Many customers have relied solely on others or receiving an email "the backups worked" and never test their backup's validity.

As with any Tier-1 application, this is a risk to your business that we strongly encourage you to avoid by performing regular simple checks of your backup solution. In many cases,

customers adhering to these Best Practises have uncovered that there was a problem with their backup solution only after a hardware / server failure thus we strongly encourage you to chose a known industry partner's products for backups of your Enterprise Vault environment.

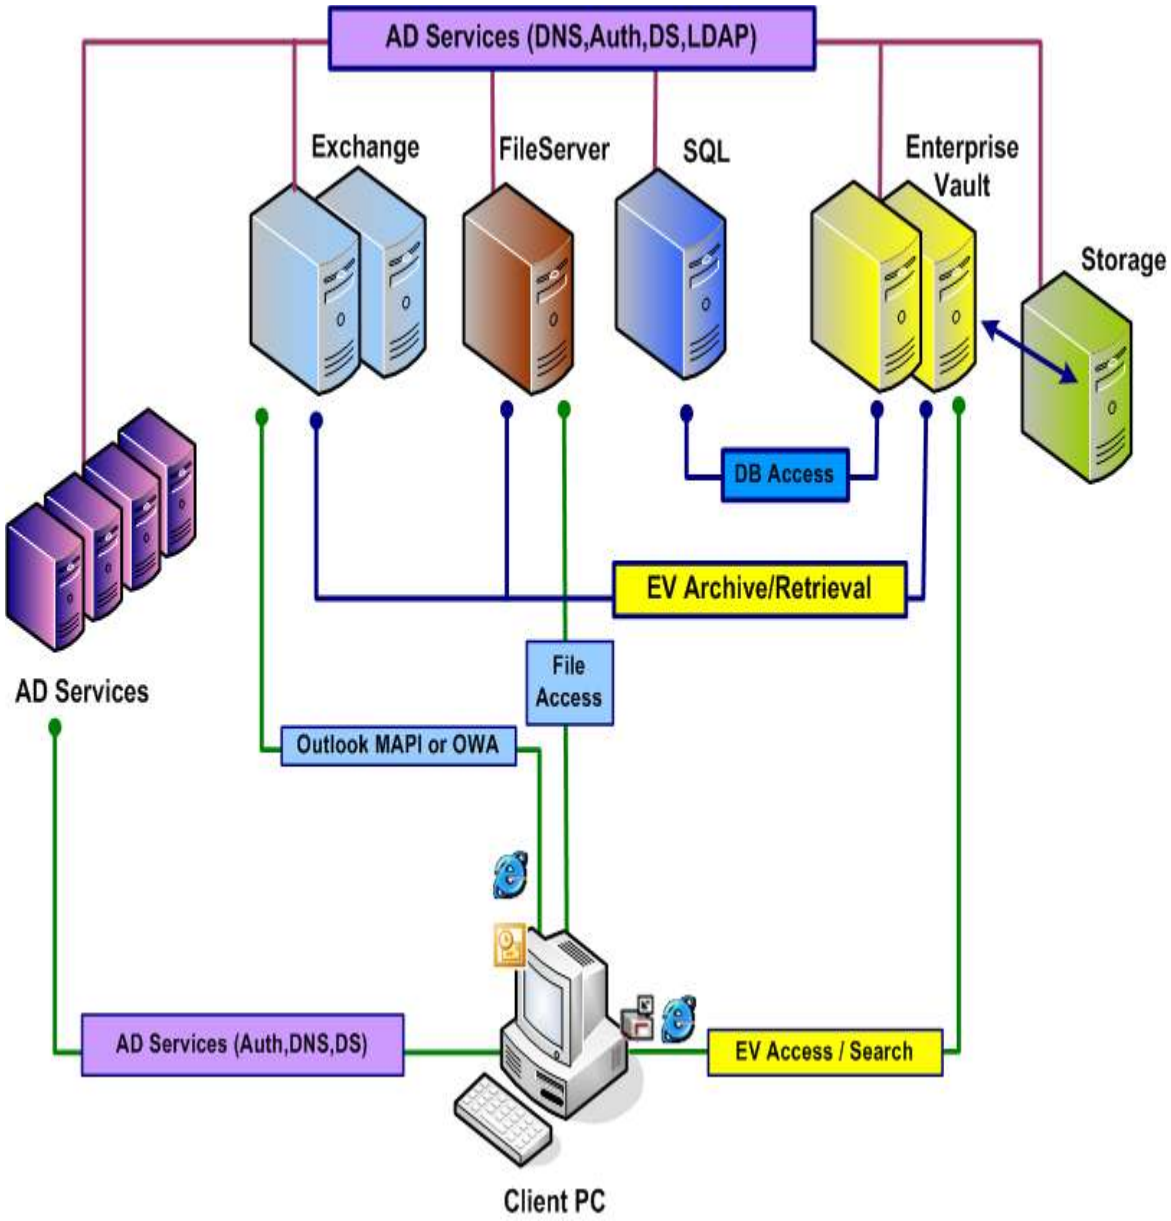

High-level Enterprise Vault Deployment

Enterprise Vault Backup Requirements

SQL Databases

Contained in your Microsoft SQL environment are many databases that Enterprise Vault uses as part of the solution. These databases are essential to insure that your deployment is healthy, able to archive and know where content is when you, or your users access archived content. Because of the number of users, some sites will have a dedicated SQL server for their Enterprise Vault environment; in other smaller sites, shared access to a SQL server. Regardless of SQL installation type, Best Practise dictates that you must make backups of the Enterprise Vault data contained in SQL.

Of particular importance, the following Enterprise Vault databases as part of the Best Practises must be backed up on a regular basis:

Directory Database
Audit Database
Monitoring Database
Database transaction logs

For upgrades and service pack applications to your SQL server, it is a Best Practise to begin by backing up the master and MS-SQL databases as they contain information pointing to the existence of Enterprise Vault on the SQL server.

In many sites, the above parts of your Enterprise Vault infrastructure may be backed up by a separate person, such as a DBA (Database Administrator) which can be sufficient but the scheduling (covered later in this paper) is critical for these components to be properly secured.

Best Practises for Backing Up the SQL Databases

- Backup the Enterprise Vault databases at least once a week
- Truncate the transaction logs at the same time (to maintain reasonable disk space consumption for the overall SQL server health.)
- Backup the database transaction log daily
- Backup the Enterprise Vault databases before / after any changes to your Enterprise Vault environment (Service Pack, new release installation, new target(s) being added)

Depending upon your Enterprise Vault architecture, you may have just one set of databases, or in the case of our Compliance Accelerator / Discovery Accelerator customers you may have additional databases. Both require that Best Practise controls as part of your DR solution that they are be backed up each day as well.

Enterprise Vault Indexes

Contained inside of your Enterprise Vault deployment is one to many AltaVista™ indexes. If you are unsure as to where your indexes are located, you can determine the location in the properties of your Indexing Service in the Vault Admin Console (VAC). These indexes have one of the most important parts of your archived content and must be backed up as a Best Practise on a very regular basis. Scheduling of the backups is (covered later in this paper) very important to insure the system is at a quiescent state for optimal integrity of the indexes whist performing your backup. It is a mandatory Best Practise that you only backup the indexes when Enterprise Vault services are either stopped, or running in "Read Only" mode to insure an accurate backup of them.

Enterprise Vault Shopping Service Locations

When users want to restore items manually contained in Enterprise Vault, their content goes back to a Shopping Service location. This location should receive a regular backup, but it is not necessarily required to back it up on the same nightly schedule. Whilst the Shopping Service is quiescent during your backup schedule, users will still be able to search content, but not restore it so it is a Best Practise that the scheduling of your backups occurs during user off-hours. Unlike other parts of your Enterprise Vault environment, it is not mandatory to backup your Shopping Service locations, but if you do perform backups of them it is a Best Practise that the Enterprise Vault servers are either stopped, or running in "Read Only" mode.

Enterprise Vault Vault Store Partitions

Your archived content stored from its original location into Enterprise Vault will typically be located as many flat files on an NTFS partition (exception: EMC Centera) with a folder directory structure cascading down. If you are unsure as to where your Vault Store Partitions are located, you can determine the location in the properties of them in the VAC Console. As the content in your Vault Store partitions is crucial to your environment, it is a Best Practise that all archived content receives regular backups as disks and disk subsystem solutions can fail. It is a mandatory Best Practise that you only backup the Vault Store partitions when Enterprise Vault services are either stopped or running in "Read Only" mode to insure an accurate backup of them.

If the size of your Vault Store Partitions is very large, you may wish to consider options such as Symantec VVR, disk-replication (EMC Centera), SnapBackup (Network Appliance Filers) or RoboCopy to aid in reducing end-user downtime / access to archived content. Whilst the use of these technologies is outside the technical scope of this paper, they are all very good solutions that should be considered for larger sites as part of a Best Practise scenario for DR protection.

Enterprise Vault Server

Now that we have covered the main components of the archiving infrastructure in the environment, we have to consider the Enterprise Vault server itself. Because changes can and do get made (new archiving policy, new retention rules, new Roles-Based Administration controls), it is a Best Practise that the Enterprise Vault server itself receive a regular backup.

Content on your Enterprise Vault server that should be backed up is:

- System Registry
- Directory where Enterprise Vault is installed
- The entire registry for Administrator including local directories of:
- Local_System
- Current _User
- Vault Store databases
- Any directory where you store any:
    - Custom Forms
    - EVPM scripts
    - Modified IIS configuration files

Whilst the Enterprise Vault server does not require nightly backups, given that changes do happen in production, it is a Best Practise to perform a backup at least once-a-week on the above content. It is very important to understand as we will show later in the backups timing diagrams that the relationship between Vault Store databases and Vault Store partitions being backed up in a particular order makes for Best Practise backups of your environment.

Enterprise Vault Backup Timing and Scheduling

The previous diagrams went into detail at the network level for the various components used within a simple Enterprise Vault environment. Timing is a crucial part of a Best Practise for backing up your archived content to insure optimal safety / DR preparedness for your environment.

Many resources can contend with by the “backup window” in an IT organisation on a nightly basis, and each of these plays a role in determining how, when, and what to backup.

Most customers are aware overall of activities on their infrastructure, but can fail to take into account other activities that impair their ability to regularly successfully get a good backup of their Enterprise Vault data. This can have side-effects during a DR scenario that no one would want to encounter.

Important factors that can affect your backup timing:

- Archiving Window
- The nightly window when Enterprise Vault is converting content
- When are users offline and online
- Network traffic
- Is the LAN or WAN capable of backing up large data sets
- What else is running over the LAN or WAN to the target server(s) to be backed up?
- How much time do you have in a backup window
- Will any servers with Enterprise Vault data be unavailable
- Contention with other network activities
- Exchange backup schedule
- Compress mail stores
- Will the backup be performed to Tape, NTFS disk, Centera replication

Each of the above factors can and do have a direct effect on your ability to perform your Enterprise Vault backups. Failing to counter in these types of network and system activities is against Best Practise and should be avoided wherever possible.

Overview of an evening for backups

In implementing your Best Practises for backups, we recommend you make a chart that details significant activities as discussed in the previous section. A suggested chart to give you good visibility is below.

Activity	Start Time	End Time	Best Practise Recommended Frequency
Archiving Window			Daily
MSMQ Buffer Drain Window			Daily
Exchange Backups			Daily or Weekly – site specific
EV into Read Only Mode			At Start of each time backups are performed
EV Archive Backup			As dictated by local site rules, at least once a week
EV into Normal Mode			At End of each time backups are performed
EV Pending Shortcut Conversion			Can run from end of Backups to 0900 / user arrival time in mornings or as part of archiving task scheduling – site specific
EV SQL Database and Transaction Log Backups			Daily – Transaction Logs, Weekly EV Databases
EV Server Backup			As dictated by local site rules, at least once a month
Exchange Server Online Maintenance			Check Microsoft recommendations to ensure a healthy Information Store

The above chart is very handy to assist you in making sure that you have proper justification, and inform all on your team of the activities to hand. The recommendations in the chart vary slightly if you have Enterprise Vault safety copies enabled, and would allow you to avoid the Pending Shortcut conversion for sites that have chosen that option in the VAC.

Several additional items in your Enterprise Vault environment can have an effect on your backup windows, and should be reviewed periodically as part of Best Practise to insure that you meet your business Service Level Agreement (SLA).

Partition Management for LTO3 users

If you setup your maximum partition size to be slightly less than the current size of one tape (400GB) it will make administration and management better for your time windows.

Backup closed partitions at least:

- Once a month if you have collections enabled as when an item in a collection is no longer retained, the collections can decrease in size.
- Less frequently may be suitable for closed partitions without collections as data is static
- Consider increasing backup frequency as the closed partition becomes eligible for expiration as the partition can decrease in size

Refresh tapes

- In sites where you do not do disk-based backups, we recommend you cycle new tapes into your Enterprise Vault backup schedule as dictated by local site rules.

Scheduling the start and stop of archiving

- Allow your Enterprise Vault server and MSMQ queues sufficient time to finish any current activities so that you have an optimal backup image taken.
- Use a good batch script to run backups
- Insure that all Enterprise Vault tasks are setup to run on the Enterprise Vault server as type “Automatic”
- We recommend as a Best Practise the script below from our Enterprise Vault documentation at a minimum
- Allow a minimum time of 10 – 15 minutes after each script before starting backups.

SAMPLE BATCH FILES AND REGISTRY KEYS FOR BACKUPS

These examples should be customized to each specific environment and server; remove references to services that are not present on the specific Enterprise Vault server and ensure all other relevant services and processes are accounted for before testing. Save Registry files with a .reg extension and ensure the name and paths of the Registry files matches the name and paths in the batch files.

PREBACKUP.BAT:

```
REM --------------------------------
REM prebackup.bat
REM --------------------------------
net stop /y "Enterprise Vault Task Controller Service"
net stop /y "Enterprise Vault Storage Service"
net stop /y "Enterprise Vault Indexing Service"
net stop /y "Enterprise Vault Shopping Service"
regedit /s c:\readonly.reg
net start "Enterprise Vault Storage Service"
net start "Enterprise Vault Indexing Service"
net start "Enterprise Vault Shopping Service"
net start "Enterprise Vault Task Controller Service"
```

POSTBACKUP.BAT:

```
REM --------------------------------
REM postbackup.bat
REM --------------------------------
net stop /y "Enterprise Vault Storage Service"
net stop /y "Enterprise Vault Indexing Service"
net stop /y "Enterprise Vault Task Controller Service"
regedit /s c:\normal.reg
net start "Enterprise Vault Storage Service"
net start "Enterprise Vault Indexing Service"
net start "Enterprise Vault Task Controller Service"
```

Sample Registry File to Place Services in a Read Only Mode

[HKEY_LOCAL_MACHINE\SOFTWARE\KVS\Enterprise Vault\Storage]
"EnableArchive"=dword:00000000
"EnableExpiry"=dword:00000000
"EnableFileWatch"=dword:00000000
"EnableReplayIndex"=dword:00000000
"EnableNSFMigrations"=dword:00000000 ← For Lotus Domino Sites Only!
"EnableRestore"=dword:00000000
"EnableCrawler"=dword:00000000
"EnablePSTMigrations"=dword:00000000

Sample Registry File to Return Services to Normal Mode

[HKEY_LOCAL_MACHINE\SOFTWARE\KVS\Enterprise Vault\Storage]
"EnableArchive"=dword:00000001
"EnableExpiry"=dword:00000001
"EnableFileWatch"=dword:00000001
"EnableReplayIndex"=dword:00000001
"EnableNSFMigrations"=dword:00000001 ← For Lotus Domino Sites Only!
"EnableRestore"=dword:00000001
"EnableCrawler"=dword:00000001
"EnablePSTMigrations"=dword:00000001

Stopping Services on Clustered Enterprise Vault Servers

Special consideration should be taken when creating pre- and post-backup scripts to be used on clustered Enterprise Vault servers. The clustering product should be configured to issue the stop and start commands using its native commands. For example, when Enterprise Vault is clustered using SYMANTEC Cluster Server, the hares command should be used to online and offline the various services. Failure to do so can result in unexpected failovers and resource offlines.

Please refer to the Related Documentation section of this document for links to the appropriate references.

Backup and Recovery of Enterprise Vault with NetBackup

As the Enterprise Vault installation matures, organizations are faced with new challenges in regards to backing up the Enterprise Vault data. Because data accumulates quickly, backup windows become stretched, restore tasks become unmanageable, catalogues bloat, and media costs skyrocket. This section will discuss some advanced strategies for protecting the Enterprise Vault data whilst minimizing backup windows, catalogues growth and costs.

Categories of Enterprise Vault data

Before an effective backup strategy can be designed, it is important to understand the different characteristics of the data to be backed up. In larger environments, it may be more efficient to have several different approaches to the backups, whilst a smaller environment may value a less complicated approach.

Enterprise Vault data can be broken into these categories:

- Small files that change regularly
- Small files that do not change regularly
- Medium files that change regularly
- Medium files that do not change regularly
- Files that may be restored singly
- Files that rarely will be restored singly
- SQL databases and logs
- Subsystem files (registry, license keys, etc)

As an example, the Enterprise Vault indexes are made up of small files that change regularly and will likely not be restored singly. Collections files stored in a closed partition are made up of medium files that do not change regularly and may be restored singly. When dealing with large amounts of data, it may be more efficient to apply different backup methodology to the different types of data.

FSA recalls can provide Netbackup a challenge in that Netbackup could cause more files to be recalled than a System Administrator wishes. Best Practise dictates that the Enterprise Vault Systems Administrator should make sure that **bpbkar32.exe, tar.exe, bpfis.exe** and **bpcd.exe** are in the FSA exclusions list on any File System / Volume / Folder that FSA archives from before enabling NetBackup to do backups

Leveraging the Enterprise Vault Design

Many parts of the Enterprise Vault design have an impact on the choice of backup method and the length of time required performing a backup. Customers using tiered disk or tape migration to NetBackup should pay particular attention to the Netbackup section for best tape migration strategies. If you are using disk migration, refer to the Enterprise Vault Administrator's guide for assistance with that second or third tier storage applicable to your environment.

Vault Store Partitions

A Vault Store Partition is a physical path where the Enterprise Vault savesets can be stored. A Vault Store is a logical collection of the Vault Store Partitions, and only one Vault Store Partition can be open (eligible for write activity) at a time within that Vault Store. The multiple Vault Store Partitions within the Vault Store can be located on the same physical media, or on multiple locations and even storage types.

Vault Stores and Vault Store Partitions can grow very rapidly, especially during the initial stages of deployment when the "backlog" is being ingested. Backup times will increase very rapidly, causing scheduling contention between backups and scheduled archiving runs. One of the easiest ways to control the backup windows is to periodically "close" the Vault Store Partition. This action will open a new Partition, either on the same media or elsewhere, and prevent new data from being added to the closed partition. This has several advantages:

- The open partition is a smaller subset of the data and will back up more quickly than the whole dataset
- The closed partition can be backed up less frequently, as the organization desires; remember to factor in the impact of Storage Expiry when plotting the backup frequency on the closed partitions
- The backup jobs can be configured to target the top-level saveset directory rather than the root of the drive, allowing for concurrent jobs and shortening the overall time required to capture the backup
- Different NetBackup retention policies can be configured depending on the nature of the data being backed up (open or closed partition)

Collections

Enterprise Vault can be configured to consolidate individual saveset files into Collection (cabinet or cab) files when the files reach a certain age. This can help manage backup times as many types of backups handle the larger cab files more efficiently than the many small savesets that go into them. This is particularly beneficial on closed partitions, where new data is not being added.

Enabling Collections is a change that should be carefully considered, like any other environmental change. There are some special considerations for environments where collections are enabled:

- When a saveset has been collected, it must be extracted again when it is recalled. This can cause a delay in the search, restore or export process
- Adequate free disk space on the storage location should be incorporated into the design to accommodate extracting/uncollecting large search result sets (such as Discovery Accelerator exports)
- Like any other data, Enterprise Vault data can become lost due to physical corruption, hardware failure, and other events. If there is no current backup, the data can be lost forever. Whilst this is true for both savesets and collections, the loss of a single collection file can represent hundreds or thousands of savesets. Therefore the integrity of the backup mechanism becomes much more critical
- Should you lose a collections file and not have a backup copy, contact Symantec Technical Support immediately for assistance in removing references to the missing cab and its contents from the Database. Do not attempt to rebuild any indexes until these references have been removed from the Database as the sheer number of invalid database references will cause indexes to fail if rebuilt
- When items within a collections file expire, they are not dynamically purged from disk; rather they are removed from the database and indexes until a minimum number of items remain in the collections file. Only then will the collections file "shrink" from its original size. Disk space utilization projections should account for this behaviour
- Utilizing collections can have an impact on the Single Instance Storage ratio as additional items are archived

Indexing Locations – User Indexes and Journal Indexes

An indexing location is a physical path where the Index data files can be written. Each Enterprise Vault server that runs the Indexing service can have multiple open Indexing locations; in fact, the more open Indexing locations, the greater flexibility in planning your backup strategy.

Unlike the Vault Store Partition, a closed Indexing location will still allow for data to be written; closing an Indexing location means that no new indexes will be created there, but existing indexes will continue to grow. So the backup strategies applied to your Vault Store Partition may not be a good fit for your Indexing Locations. Some other considerations:

- When all Indexes reside on the same physical storage, creating multiple Indexing Locations helps "widen" the directory structure and allows an Administrator greater choice when configuring backup jobs.
- As an example, with a single Indexing Location, all indexes will be written to the directory F:\EVIndex and the backup job can only be configured to target F:\EVIndex (unless the Administrator wishes to copy the Index Volume ID from each child directory and configure the backup job in this way, which can be very

time consuming and leaves much room for error). All indexes are backed up in a single job which may take several hours.

- Alternately, by configuring 10 open Indexing Locations in the same path, an additional level is added to the directory structure. This allows the Administrator to configure 10 NetBackup targets using user-friendly names (for example, F:\EVIndex\Folder1, F:\EVIndex\Folder2, and so forth), thereby splitting the backup work into 10 smaller jobs. Each individual backup job will complete faster because each work stream is processing less data; some of these jobs can be configured to run concurrently, depending on the server hardware and the NetBackup architecture. By scheduling multiple jobs to run simultaneously, the overall time to complete the backup can be drastically reduced.
- User Indexes tend to be smaller and more manageable in size; often, restoring a User Index from multiple backups (Full plus Differential, Full plus Incremental) is more labour intensive and slower than restoring from the most recent Full then using the Enterprise Vault IndexVolumeReplay tool to Update the index to a current state. Therefore, some organizations may elect to run a Full backup more frequently, and no Differentials or Incrementals. This is especially important to consider in heavy usage environments where a User Index will change significantly in the course of a day; the Incremental backup may contain as much change as a Full, requiring more time to back up, but may require multiple restore actions to become current.
- Journal Indexes tend to be mission-critical; they are also made up of millions of very small files and tend to be rather large which means long backup times. Practically speaking, it is unlikely that you will ever restore a single file from any Index; generally the entire index will be restored as a whole. Therefore, there is no benefit to recording the file allocation table for a volume where only Journal Indexes reside (you will be restoring the whole volume as one). Selecting a backup method that backs the data up as a whole without recording individual file locations can drastically increase backup speeds.

Leveraging Flash Backup

As mentioned previously Enterprise Vault index data can consist of millions of small files. A traditional NetBackup on a standard Windows file system requires significant overhead to track the metadata from these files. This overhead can greatly increase the size of the NetBackup catalogue as well as cause the EV backups to write to the backup medium very slowly. This increases the time and performance drag the backups create on the EV environment. FlashBackup can be used to greatly speed up these backups by recording a snapshot of what the file system looks like at the time of the backup and then reading the data at the block level and writing it to the backup medium. More information can be found about FlashBackup in the 6.0 NetBackup Advanced Client Guide or the 6.5 Snapshot Client Guide.

FlashBackup capabilities

FlashBackup is a policy type that combines the speed of raw-partition backups with the ability to restore individual files. The features that distinguish FlashBackup from other raw-partition backups and standard file system backups are these:

- Increases backup performance as compared to standard file-ordered backup methods. For example, a FlashBackup of a file system completes faster than other types of backup in the following case:
- the file system contains a large number of files
- and most of the file system blocks are allocated
- Individual files can be restored from raw-partition backups.
- Supports multiple data streams, to further increase the performance of raw-partition backups when multiple devices are in the Backup Selections list.
- Can be used as the backup method against an Instant Recovery snap shot and several other types of array or software snapshot.

FlashBackup restrictions

- FlashBackup policies do not support file systems that HSM manages.
- FlashBackup does not support VxFS storage checkpoints that the VxFS_Checkpoint snapshot method uses.
- FlashBackup supports the following I/O system components: ufs, VxFS, and Windows NTFS file systems, VxVM volumes and LVM volumes, and raw disks. Other components (such as non-SYMANTEC storage replicators or other non-SYMANTEC volume managers) are not supported.
- FlashBackup-Windows policies do not support the backup of Windows system-protected files (the System State, such as the Registry and Active Directory).
- Restores from a FlashBackup image can be slower than a regular file system restore. This is due to the fact that NetBackup has to search more of the tape media to find the section of the snapshot that has the pertinent files. This can be mitigated by doing restores from a disk image. Also FlashBackup will typically be used on EV data that should be restored as a set rather than at the file level so the impact of this restriction is limited.
- FlashBackup is a licensed option as part of the Advanced Client in 6.0 and the Enterprise client in 6.5. Many large customers already have enterprise clients for their backup clients. No separate software installation is required on Windows.
- Note that FlashBackup requires some free disk space to be used as a cache for changes made to the data whilst the backup executes. For other types of applications, the cache space requirement can be quite large, even as large as the data to be backed up. Placing the Enterprise Vault services in read-only mode as discussed earlier will prevent changes to the data whilst the backup executes. This means that the cache space required to execute a FlashBackup for Enterprise Vault is quite small.

Leveraging RAW Backup

Another way to speed up backups of a volume with many small files is to do a raw partition backup. It has many of the same qualities as a FlashBackup but no licensing is required. The downside is that Raw backups can only be restored at the volume level. No tracking is done as to what files were on the volume at the time of the backup. If a drive or volume will only ever be restored in its entirety than a RAW backup may even be more desirable because no catalogue metadata is consumed tracking each EV file in the backup. RAW backups are configured as an MS-Windows-NT backup which a special File directive. More information can be found in the NetBackup Administrators Guide Volume I for Windows.

Leveraging Snapshots and Off-Host Backups

Enterprise Vault data is often stored on external disk arrays that have hardware snapshot capabilities. If the environment uses disk that has snapshot capabilities and is compatible with the NBU hardware HCL then NetBackup can be used to manage those snapshots and enable low impact backups as well as rapid restores of EV data. Furthermore NetBackup can be used to mount the snapshot data on another host and perform a traditional backup or FlashBackup from that off-host location to completely remove backup impact from the production EV data. For environments with substantial EV data that have very small downtime or backup windows this may be the only viable option to backup the data. Keep in mind that NetBackup does not keep the various elements of the EV data store in sync when doing snapshot backups because this is not a "database extension" backup. So care must be taken when setting up snapshot backups that EV is in a consistent state as the snapshot is taken. This can be scripted using methods described in the link at the end of this whitepaper entitled "How to set the Enterprise Vault™ Services to Read Only Mode"

For more information on how to configure snapshot backups using NetBackup and disk arrays that are supported in the NBU hardware HCL please refer to the Snapshot Client link at the end of this whitepaper.

Leveraging Multi-Streaming

Many servers only have enough backplane or disk speed to saturate a 100baseT network interface. Larger modern Windows servers with several SAN disk volumes and a gigabit network interface for backups can push much more backup data. NetBackup can leverage this data moving capacity by simultaneously backing up several streams at once. This is enabled via a checkbox in the NetBackup Policy Attributes. These streams can be directories on the same drive but more typically are separate drive letters or volumes. To get the best performance from enabling multistreaming, ensure that the different streams operating concurrently are reading from different disk spindles. Also ensure there is sufficient network bandwidth on the client to push the backup data across the network to the NBU media server.

Leveraging Synthetic Backup

Synthetic backups can be used to reduce the amount of time backups are performed against the EV servers and thus increase their uptime and performance for users. Synthetic backups are Full or incremental backups that are created from prior backups taken on a client. All of the work is done on the master and media server and the client is not contacted. For a small to medium environment this may work well.

There are several drawbacks to synthetic backups for EV data:

- The EV index data has a high change rate on a daily basis. Synthetics are generally not recommended for data that has a high rate of change from day to day because the Incrementals are nearly as large as the full's and processing the backups to generate the synthetic full can take too long
- Synthetic backups are also only valid for NetBackup policies of the type standard or MS-WindowsNT. As such synthetic backup schedules don't work for FlashBackups or other snapshot backups.
- Synthetics would work well for environments that are setup with partitions and have a requirement to backup all data frequently whether it changes or not
- However, the closed partitions could be backed up using a mixture of regular backups on the weekends and then synthetic backups at other times to limit the amount of work done on the EV server
- Synthetic backups could be a suitable choice for closed Archive partitions which do not change frequently

Special Considerations for NDMP/NetApp Storage

NDMP storage is frequently used to store EV data. If the storage is accessed through a file system it can be backed up through the host server similar to data on other types of disk. There are some limitations to this method including the fact that millions of small files in the EV data set will take a long time to backup. This can be resolved by backing up the NDMP filer directly, bypassing the host server.

With NetBackup there are several ways to accomplish this. Tape drives can be attached directly to the filer and then a NetBackup policy can be configured to issue NDMP commands to the filer to mount media in the drive(s) and backup the filer data. This is considered a Local NDMP backup. The other method is a three-way backup where one NDMP filer uses tape drives mounted to another NDMP filer to perform the backup. For Large NetBackup environments with many filers and the need to share tape drives it is often beneficial to configure the tape drives on NetBackup media servers and then setup a type of three-way backup known as remote NDMP backups.

This type of backup architecture maintains the benefits of backing up the filer directly, bypassing the host server, but it also allows the tape drives to be used easily by other backups when not in use by the filer(s). With any of these types of NDMP backups, the advantage is that there is not as much overhead in backing up many small files because the backup can be done at the volume level. Also the entire filer can be restored as a unit which can speed recovery. Filers that support Direct Access Recovery (DAR) can still restore individual files as well. It is important to check with your filer vendor to see what methods

and techniques are supported. The Symantec web site also has NDMP compatibility information.

On supported filers, the NetBackup catalogue overhead of recording information about each small filer in the EV data set can also be mitigated. This is accomplished by adding the SET HIST=N directive into the policy file list. This setting tells NetBackup to only catalogue the directory structure on the volume and not the individual files. All of the files/data is still backed up but it is not individually catalogued so a file level restore is not possible from this backup data. This can still be tremendously useful with index data from Enterprise Vault where there are potentially millions of small files but they are typically restored in an all or nothing fashion.

Several snapshot techniques can also be used with NDMP filers. NetBackup can control how often the snapshots are taken and catalogue the snapshots to ease the restore process. With the first snapshot type, NetBackup directs the filer to take a snapshot of the data to storage space on the same filer. This is done through the host that mounts the filer via CIFS or NFS. This type of snapshot is known as a NAS snapshot. The second type of snapshot is known as a SnapVault. With compatible filers (NDMP V4) NetBackup can direct the filer to make a snapshot of its data to another filer in effect making a "vault" copy of the data on the first filer in case it totally fails. The tape based backup methods above can then be used to get a copy of the snapshot or SnapVault data. During either type of snapshot the data remains online for use by EV so user interruption is minimized.

Properly architected, an NDMP filer can be used effectively for storing EV data and it can also be efficiently backed up with a restore SLA kept in mind. With NetBackup 6.0 and Advanced Client license is needed for any of the above backup methods. In NetBackup 6.5 these methods are part of the SnapShot client which is included in the Enterprise client license. There is a dedicated guide to setting up NDMP backups entitled "NetBackup for NDMP System Administrator's Guide" Additionally the snapshot methods are documented in the 6.0 Advanced Client guide or the 6.5 Snapshot Client Guide.

NetBackup Agent for SQL Server

Enterprise Vault stores metadata in a Microsoft SQL format. NetBackup provides and agent that allows this data to be backed up in an online fashion to get consistent database data whilst still allowing EV to remain online. Bear in mind that the NetBackup SQL agent does not in any way keep the SQL data quiescent with the other file based data that EV stores. Therefore it is considered a Best Practise to put EV in a read only mode during any backup session. The agent will backup both the SQL data and the transaction logs and works with the Microsoft SQL backup API to commit and truncate the transaction logs back into the database.

The SQL agent is included in the Windows NetBackup client in 6.0 and above. There are configurations scripts needed on the SQL machine to define how the backup should be taken. This includes multi-streaming the SQL databases to improve backup performance.

On the NetBackup Server side, a MS-SQL-Server type backup policy is created. The policies include list references the scripts created on the SQL server to initiate the backup. All SQL backups are considered Full backups in NetBackup. Incremental backups are not performed. There is a dedicated "NetBackup for Microsoft SQL Server System Administrator's Guide" that is dedicated to how to configure SQL backups. Using the NetBackup SQL agent for EV SQL data greatly increases the chances that a successful restore can be performed from a backup taken whilst SQL is up and running.

If the environment does not allow the NetBackup SQL agent to be used then the EV administrators should setup SQL to do a database dump at the desired frequency out to a separate directory from the live data. NetBackup can then be configured to exclude the live SQL database directories and only backup the database dump files. NetBackup scheduling should be setup such that the dumps are not backed up whilst they are still being created or before the nightly dumps have occurred. Failure to do so will either result in corrupt database backups or backing up data from the previous database dump which will likely not meet expected restore SLA's.

Selecting the Appropriate NetBackup Retention Policy

Each organization must evaluate their NetBackup retention policies as they relate to the Enterprise Vault data; inefficient retention policies for the NetBackup media can lead to excessive overhead, both in terms of media costs and administrative overhead. The NetBackup retention is independent of the Enterprise Vault Retention Categories but both sets of retention rules should be considered when making these choices. (Note: each customer should become familiar with the compliance regulations and legal issues affecting their organization before planning any Retention strategy.)

Backup Media Containing Index Data

A shorter retention can be assigned to these media, as the Index backups will become less useful as time passes. For example, it is unlikely that you will need to restore an index from a two-year-old backup; a more recent backup will be used, or the index will be rebuilt from scratch. The Enterprise Vault indexes are not very useful on their own, without the corresponding savesets; a two-year-old index must be brought to its current state before it can effectively be used for searching. And because the Indexes contain instructions for accessing Enterprise Vault data, but no data themselves, it is not likely that the indexes themselves will be relevant or needed for legal or compliance reasons. Therefore, many customers will find it useful to retain the Index backups for a shorter period of time before retiring or recycling the media.

Backup Media Containing Archive/Saveset Data

Backup media containing savesets (live, archived data) can become a liability if retained for too long or not long enough. As an example, Company A sets their Enterprise Vault

retention category to 7 year retention, to comply with Regulation B. The NetBackup retention is 6 years, per company policy. A backup captures Saveset C today, which it the day before it reaches its 7-year age. Saveset C is then expired from Enterprise Vault but because the tape retention is set for 6 years, Saveset C is still present in the organization for another 6 years--- for a total of 13-year retention. This may put the organization in a situation where it must produce Saveset C for discovery or compliance reasons. Companies should evaluate their combined retention policies to ensure the desired goals are being met.

Scenarios and Procedures

Enterprise Vault

All of the following procedures assume the customer is running Enterprise Vault 2007. The same procedures should apply to many other current versions of Enterprise Vault. Consult the product documentation for the procedures appropriate to your version of Enterprise Vault. As with any procedure, appropriate testing should be performed before a production rollout.

How to Close and Open Enterprise Vault Partitions

1. Using an account with sufficient privileges, open the Vault Admin Console (VAC) and navigate to the Partition that you would like to close
2. Use the drop-down menu on the General page to change the Status from Open to Closed

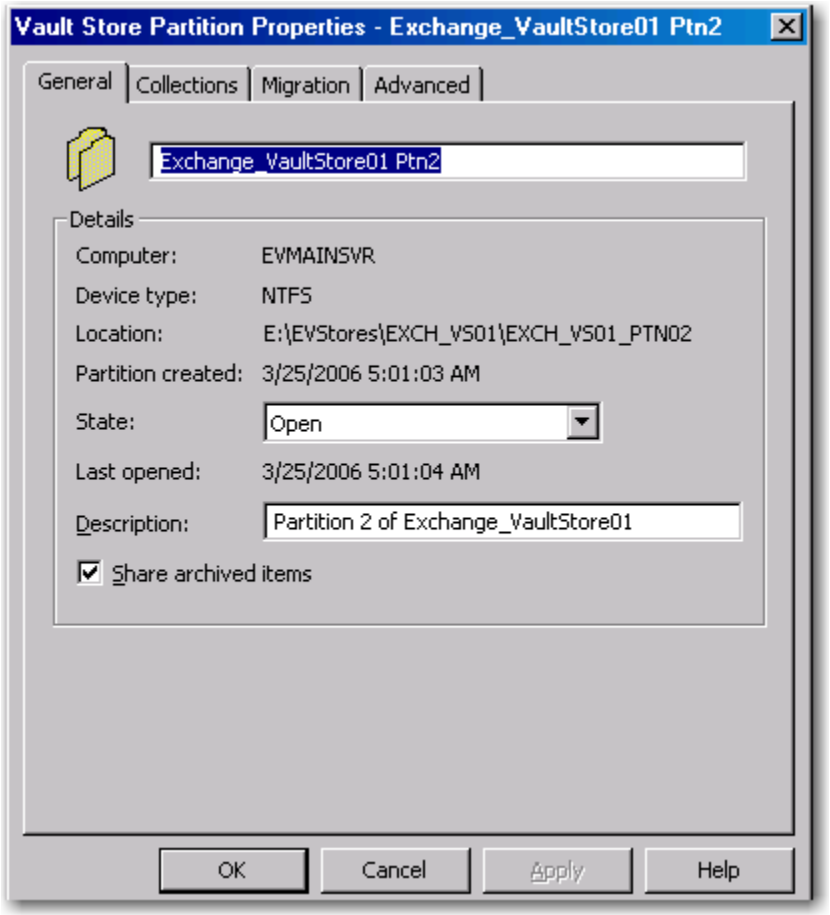

3. You will receive a warning that you are about to close the only open partition. Click OK.

4. Now proceed to open a new Partition by right-clicking the Vault Store that the Partition should belong to
5. Select New-->Partition

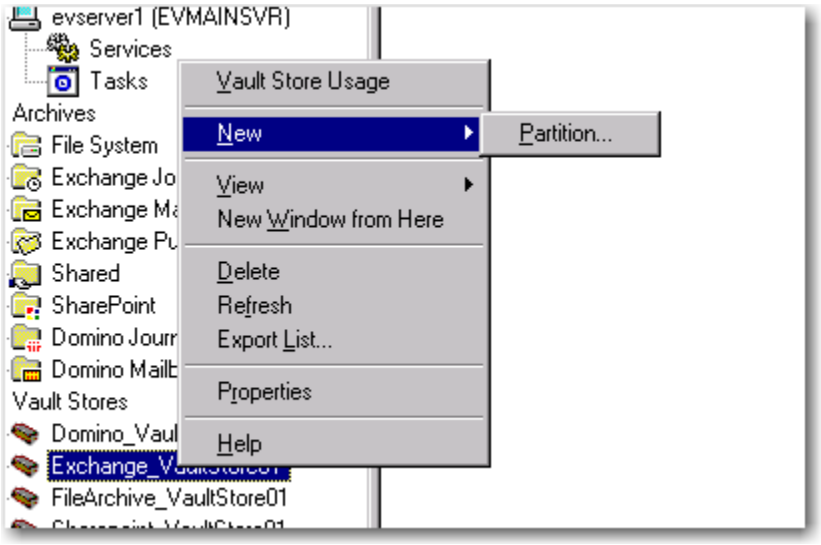

6. A Wizard appears. Step through the wizard; fill in a name and description for the Partition and ensure the Status is set to Open, then select Next
7. Follow through the remaining screens in the Wizard to define the storage type (NTFS, Network Share, NetApp, Centera, etc) and the properties (location, IP, etc) appropriate for your environment

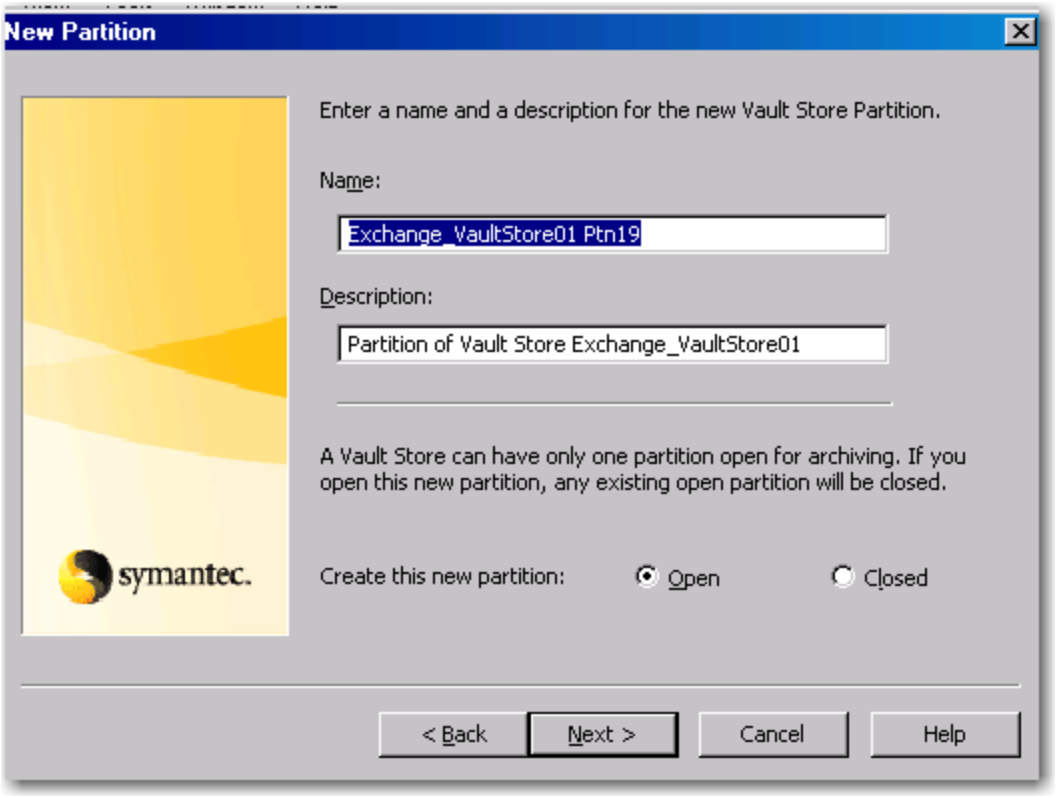

Note: You can also start the process by creating a new Open Partition (step 4-7 above) which will automatically close the existing partition for you.

How to Enable Collections

1. Using an account with sufficient privileges, open the Vault Admin Console (VAC) and navigate to the Partition on which you wish to enable Collections
2. Right-click the Partition and select Properties
3. Select Collections
4. Enable “Use Collections Files”
5. In most environments, the Collector will be Enterprise Vault
6. Configure the appropriate schedule
7. Configure the desired Collection size
8. Configure the “Collect At” age
9. Click Apply, then click OK

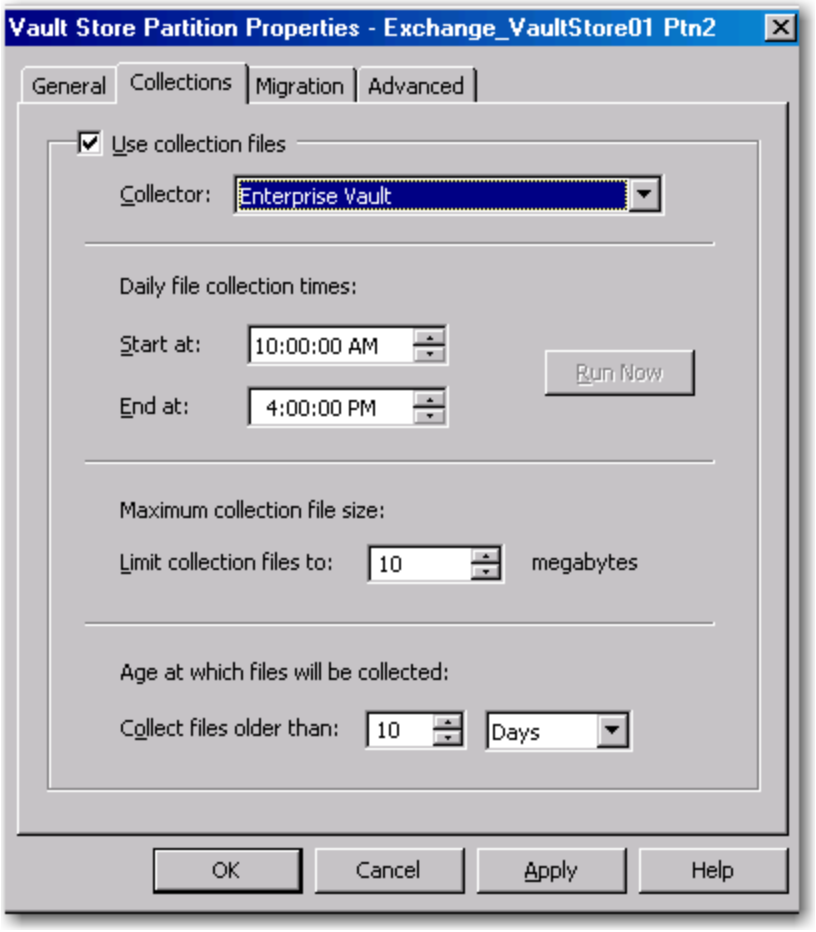

Note: Once enabled on a partition, the use of Collections cannot be disabled

How to Add Indexing Locations

By default, Enterprise Vault only requires a single open Indexing Location. However, distributing indexes into multiple locations can provide additional flexibility, not only for backups, but for future storage migrations. Several indexing locations can point back to the same physical storage; whilst load and I/O can be more fully optimized by using separate physical storage for each indexing location, this is not a requirement. As it is Best Practise to have multiple Indexing Locations, the Indexing Service will create new indexes in each Open Location via round-robin, which will distribute the indexes more or less evenly.

This is best done when the installation is new and there has not been the opportunity for many indexes to be created.

1. Using an account with sufficient privileges, open the VAC and navigate to the Indexing Service you would like to manage. Right click the Indexing Service and select Properties

2. Click on Index Locations, then click Add
3. Use the Explorer to define the desired location for the Indexes

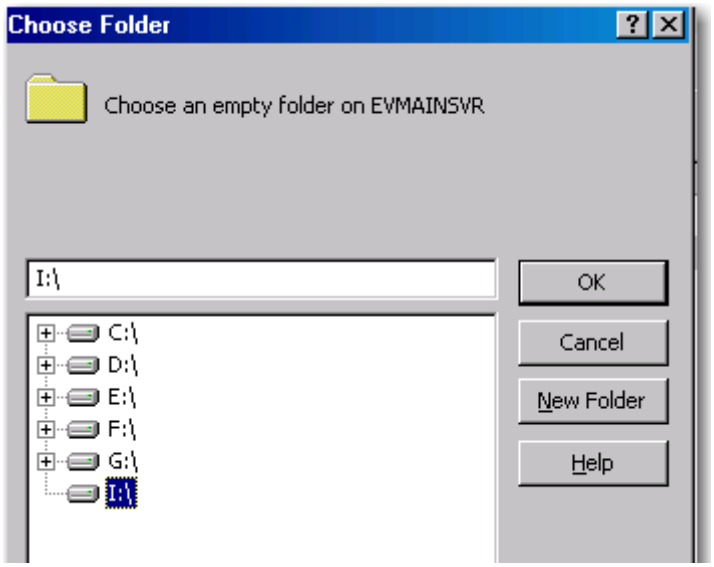

4. If a subdirectory does not exist, click on New Folder to define one (separating the Indexing Locations into subdirectories allows NetBackup jobs to target the subdirectory rather than the root of the drive, thereby dividing one large backup job into several smaller ones)

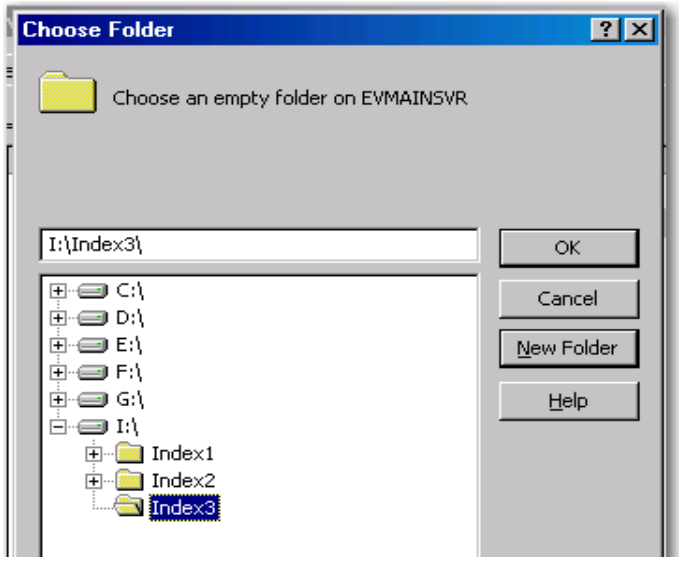

5. Continue defining locations as needed

Note: A minimum of 10 open Indexing Locations is considered Best Practise

To redistribute existing indexes, multiple edits of the Directory database may be required. Please contact Symantec Technical Support for assistance.

NetBackup

All of the following procedures assume the customer is running NetBackup 6.0. The same procedures should apply to NBU 6.5 and beyond. As with any procedure, appropriate testing should be performed before a production rollout.
Maintenance Pack 6 or above is recommended with 6.0.

How to Setup Enterprise Vault for NBU FlashBackup

NBU FlashBackup requires an Advanced Client License for the windows server performing the backup. This license Key is stored on the NetBackup Master Server. Under 6.5 the FlashBackup license is part of the Enterprise Client. The NetBackup administrator can ensure that is in place before starting the configuration.

Configuring FlashBackup for an Enterprise Vault server may require setup steps on both the client and the NetBackup policy on the master server. If the client is running any version of Windows below 2003, then it is recommended to use the Symantec Snapshot Provider or VSP to enable FlashBackup. For any client that is running Windows 2003 with Service pack 1 or above, Microsoft's Volume Snapshot Service or VSS is recommended and is the default behaviour. If VSS is used then configuration steps may be needed on the Enterprise Vault server to configure where the VSS data should reside. For the sake of this example, the environment is an Enterprise Vault server running on Windows 2003 Service Pack 1 and has data stored on the E: drive. The NetBackup environment is a combined master and media server running on 6.0 mp6 running on Windows 2003 as well. Backups will be performed to a basic disk storage unit.

Setting up VSS on the Enterprise Vault Host

1. Log in to the host as an administrator and open up Windows Explorer
2. Right Click on the drive that stores the Enterprise Vault data, in this example the E: drive, and select properties at the bottom of the list.
3. Select the Enterprise Vault data drive, in this example the E: drive and press the Settings… button as shown below:

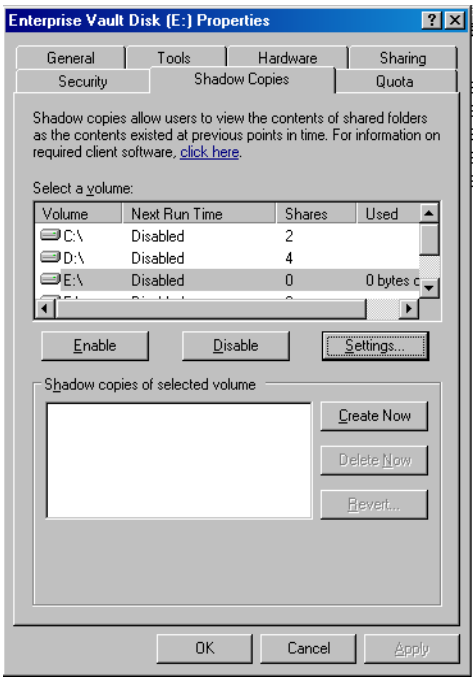

4. In the screen that appears, use the "Storage Area Located on this volume:" drop down to select a drive where the snapshot temporary data should be stored. If the drive, in this example the D: drive, is used for any other purpose, it would be prudent to also set the Maximum size parameter by using the "Use Limit:" setting. In the example shown below 10% of the D: volume is chosen or 400mb.

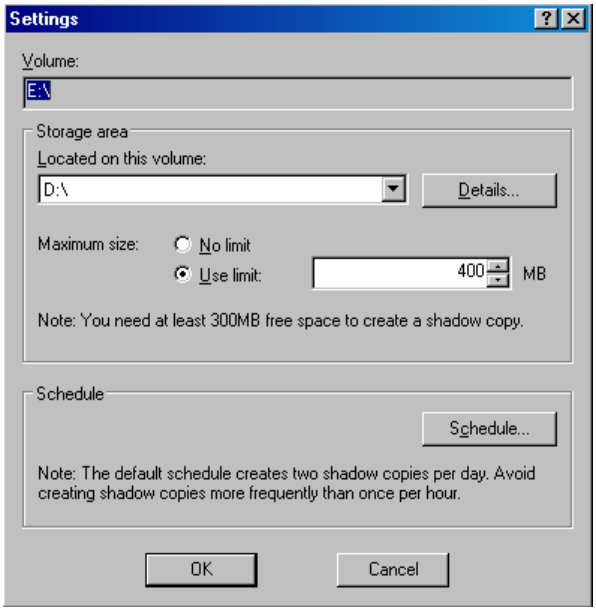

When these settings are adjusted a Snapshot schedule is automatically created to snapshot the entire E: drive. This should typically be disabled unless it will be used for some other purpose. This can be accomplished by pressing the Schedule… button.

5. Once the Schedule dialog appears press the Delete button to remove any schedules listed in the drop down box as shown below:

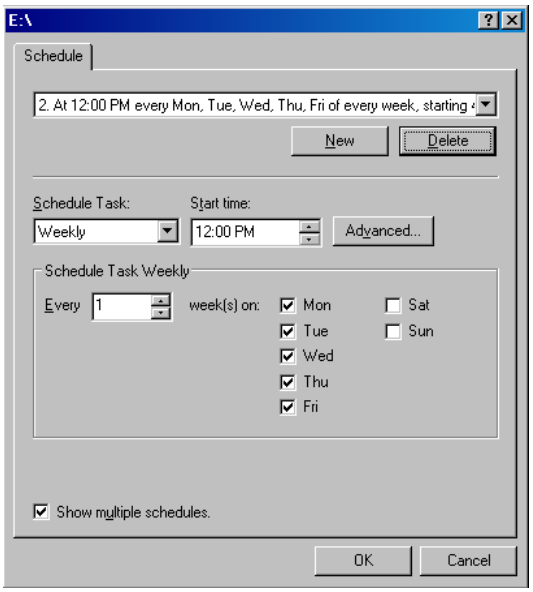

Once all of the schedules have been removed the dialog box should indicate "<Task not scheduled>" as shown below.

6. Select OK to close out of the Schedule dialog and then OK twice more to close out of the Settings Dialog and the Drive Properties dialog to get back to Windows Explorer. The Next Run Time Field in the Shadow Copies dialog for Drive E: (or whatever drive EV is on) should list "Disabled" to ensure an automated Shadow Copy Schedule is not setup to run.

Setting up NetBackup to Backup the Enterprise Vault Data Using FlashBackup

This documentation will not go through every step necessary to setup a proper NetBackup Policy for a new client. It is assumed that the Backup administrators have standard policy schema's as well as data retention and duplication standards in place. The purpose of this whitepaper is to supplement those processes and procedures with the steps necessary to configure FlashBackup for the Enterprise Vault Servers. If there are several Enterprise Vault servers in the environment then the steps can be repeated as necessary to configure the appropriate backup method for each server. If each Enterprise Vault server is setup in an identical fashion then all of the like servers can share a properly configured NBU FlashBackup policy.

1. Login to the NBU administration console using either nbSA or the Windows Administration console. In the examples given here, the Windows Administration console on the Windows master server will be used.
2. Check to ensure that the environment is licensed for the Advanced Client license or the Enterprise client license in 6.5. This can be done via the GUI using the Help Menu and then the License Keys… menu choice or via the command line on UNIX using the get_license_key utility.
3. Click on the policy section of the GUI and create a new policy for the Enterprise Vault client data. The policy name should be descriptive of both the data being backed up as well as the type of backing being a FlashBackup. The NetBackup Administrators will likely have a policy naming convention in place that should be followed
4. In the attributes tab the policy type should be set to FlashBackup-Windows as shown below:

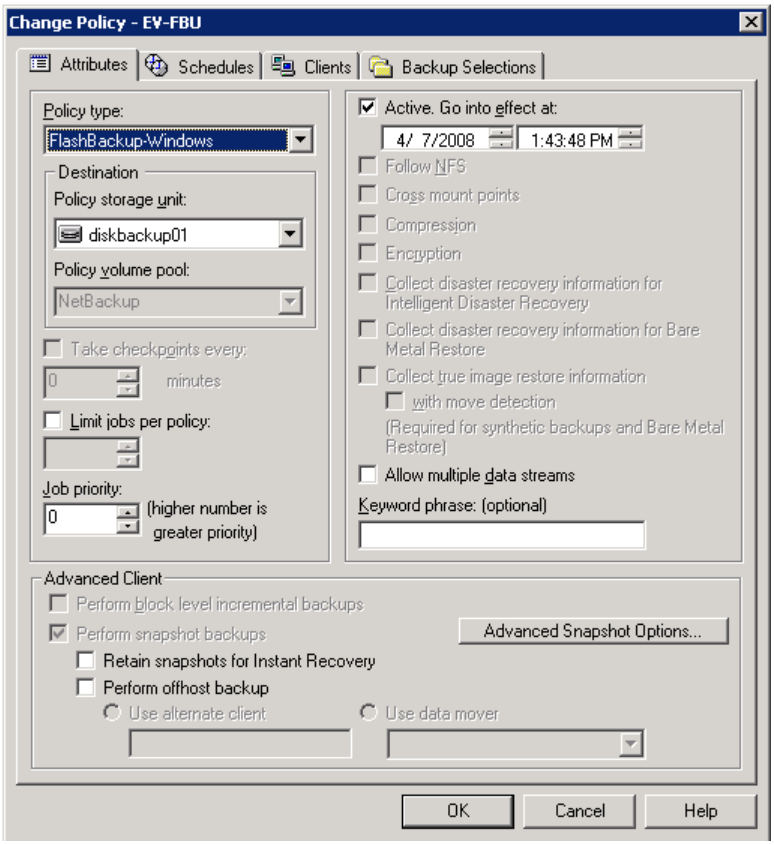

5. Pressing the Advanced Snapshot Options…. Button from the attributes tab of the policy definition will allow the administrator to configure the snapshot method to use VSS rather than Auto. When the dialog opens select VSS from the "Snapshot Method:" drop down. This will save some time initiating the snapshot since NetBackup won't need to determine the snapshot method at run time. An example of this is shown below.

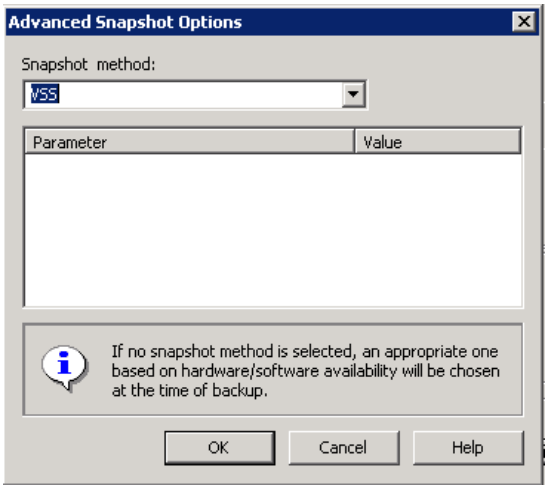

6. Click "OK" to close the Advanced Snapshot Options dialog and change to the Clients tab. The Schedules tab will not be discussed in this document because scheduling the backups is not specific to Enterprise Vault. If the architecture calls

for Enterprise Vault to be shutdown or put in a read only mode during the backup then coordination with the NetBackup staff will be needed to ensure the backup happens during the allowed window.

7. Adding a client in the client's tab is very straightforward. Take care to ensure that the correct client platform and Operating System are selected. Otherwise the FlashBackup policy will not validate once it is setup and will not backup correctly. In this example the client is setup as "PC, Windows Server 2003".

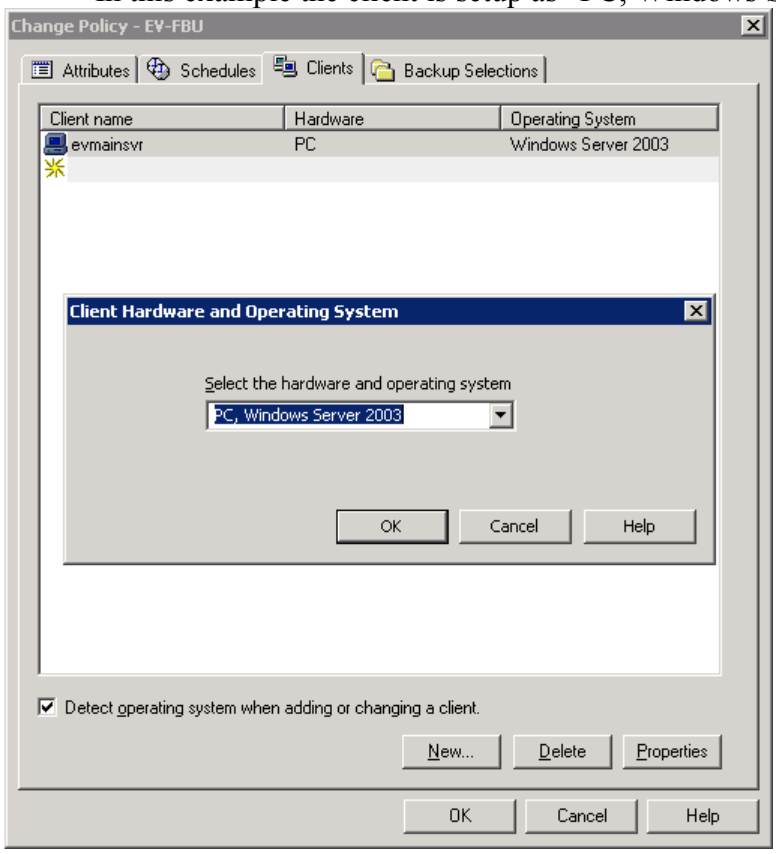

8. Click on the Backup Selections tab. This is where the more important piece of the FlashBackup Configuration comes in. To perform a FlashBackup of the E: a specific directive is needed. **\\.\E:** without the quote marks as shown below.

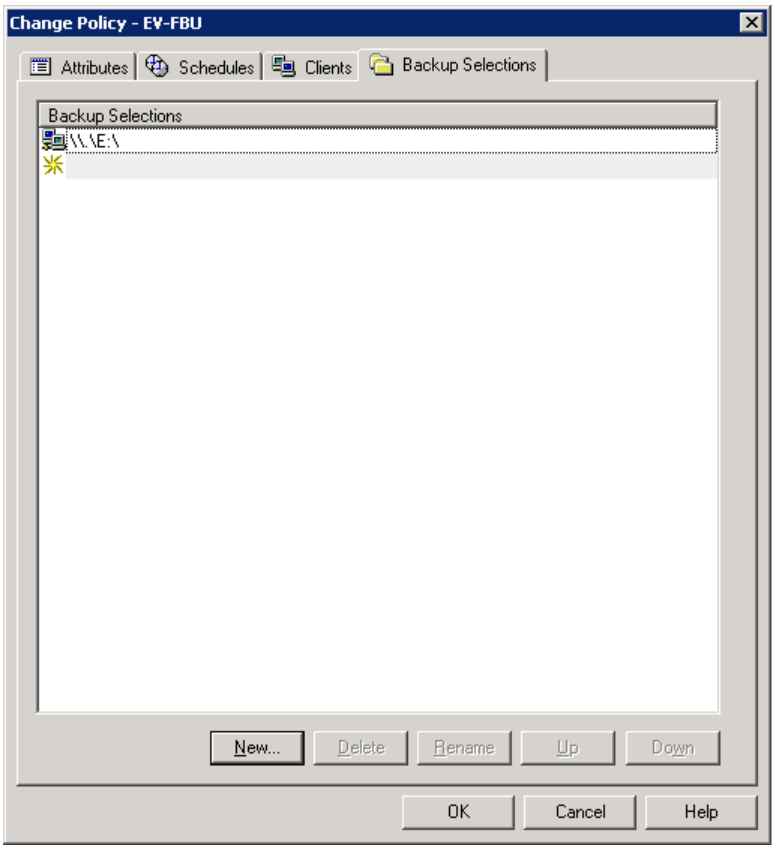

9. Once the directive is setup a dialog may appear that says “The NetBackup Client service… must have a “Log on” account specified other than the system account” as shown below. This warning can be safely ignored unless the security on the Enterprise Vault data drive has been setup such that “Local_System” would not have access.

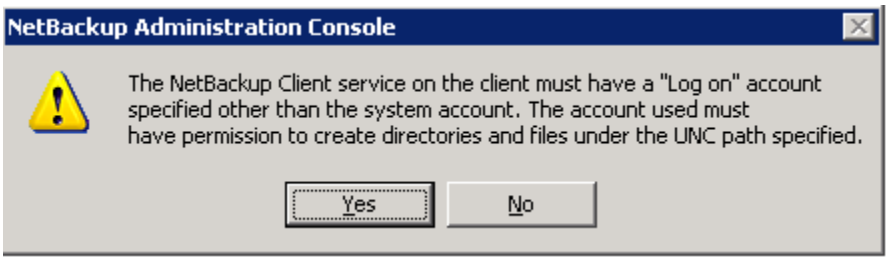

10. If the Enterprise Vault client has multiple data drives that need to be backed up via FlashBackup they can each be configured in the Backup Selections tab at this time. A New_Stream directive should be used to separate the drive letter directives as shown below. This will allow the policy to backup multiple drives at the same time using “multistreaming. Enabling multistreaming is described in the next step.

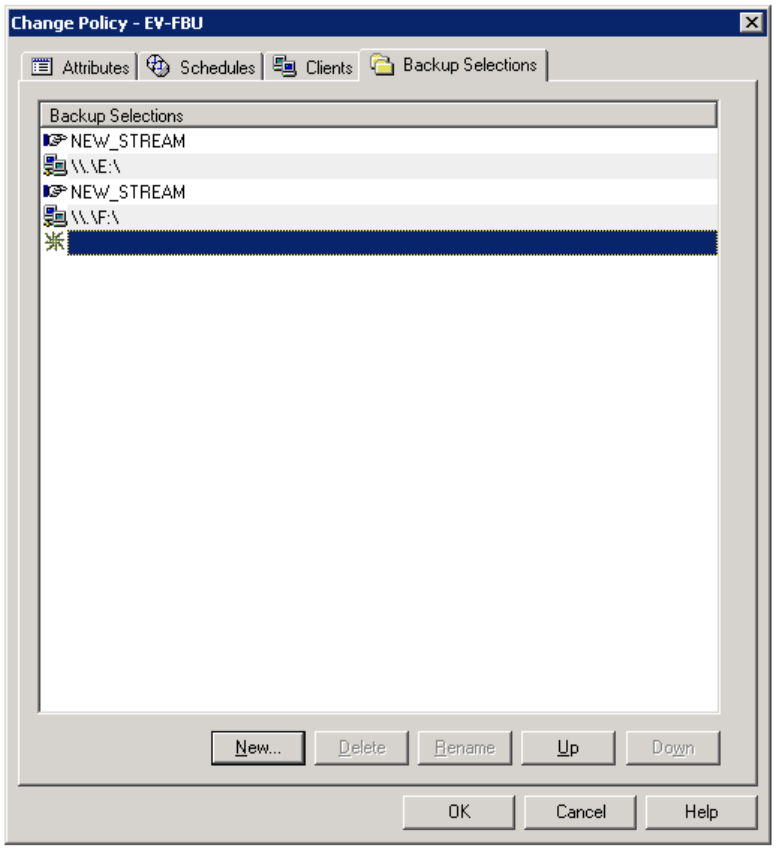

11. In the Attributes tab of the Policy the "Allow Multiple Streams" checkbox should be selected if more than one drive was setup in the prior setup as shown below. For multi-streaming to be effective the underlying LUNs or hard drives that store the EV data should support reading data from both drives at the same time. In other words each drive should be on its own set of "spindles". Also the tape drive or NetBackup disk storage unit must support the rate of data transfer pushed from all client drives combined.

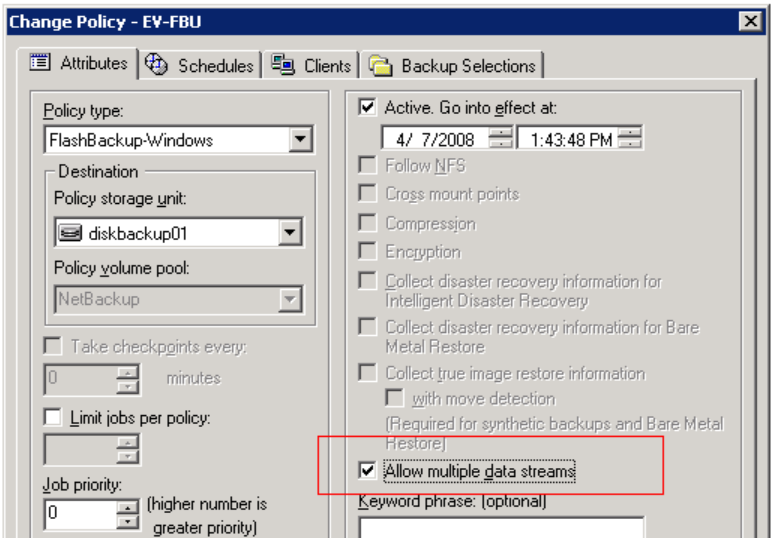

12. Click OK to close the Policy configuration screen. At this time a popup will appear saying that it is validating the advanced client configuration settings as shown below. This popup should disappear after a few minutes and should not result in any errors. If the client cannot be contacted via the NetBackup client this validation process will fail. Open the policy definition at a later time when the client is available and running the NetBackup client software to ensure the advanced client configuration is valid.

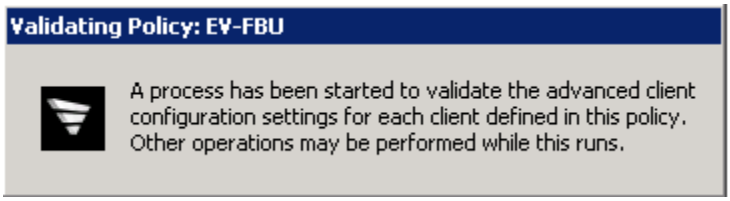

This concludes the steps required to setup a Windows FlashBackup policy to backup an Enterprise Vault Client.

Performance Comparison between a FlashBackup Policy for Enterprise Vault and a Standard OS Backup

The Performance of a FlashBackup backup clearly out performs a standard OS backup of many small files when the volume being backed up is mostly filled with data. This is because FlashBackup only has to interface with the NTFS metadata about the files on the drive once at the beginning of the backup, versus the constant metadata records needed in an OS backup. The Screenshots below demonstrate that on a test Enterprise Vault Environment where only 25% of the volume is allocated with data, the FlashBackup job still takes nearly the exact same time to backup the E drive as the OS backup of the E drive. The job on the right is the FlashBackup job which backs up 4,203,164 kilobytes at 16,762kb/sec. Conversely the OS backup only has 1,019,611 kilobytes to backup but it only backs up at 3,733kb/sec.

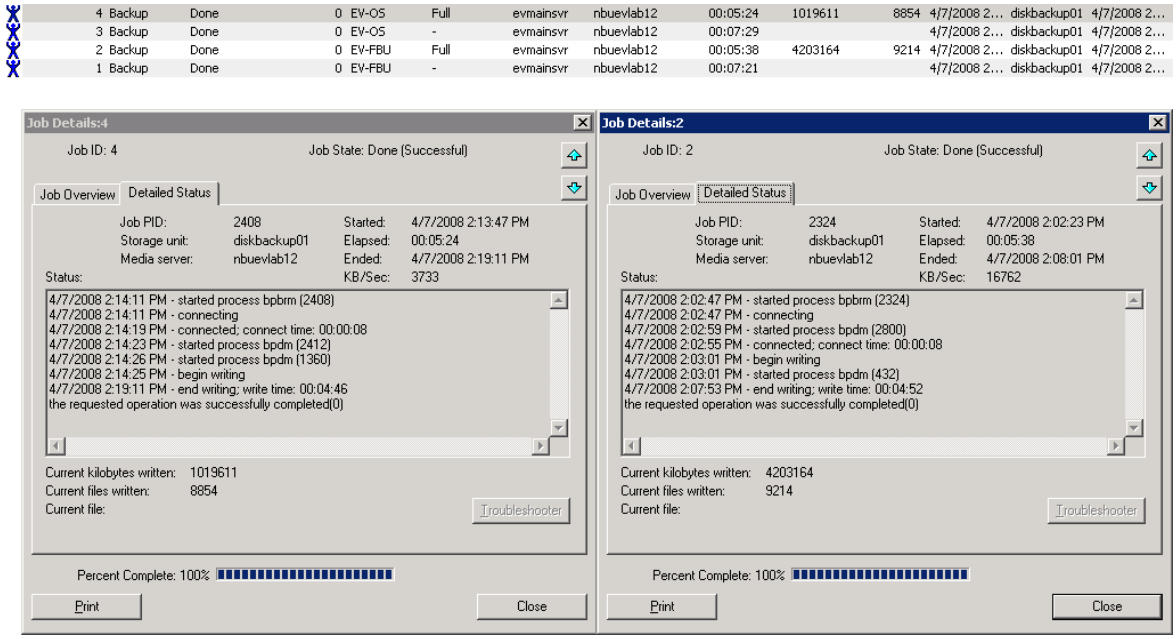

Once the volume of data grows, the performance benefits of using FlashBackup on the millions of tiny files become even more obvious. In the screenshot on the next page the backup on the left is a regular OS backup of one directory on the E drive. It takes more than 40 minutes to backup the sample directory. The EV data is then compressed into cab files and the backup time drops to around 16 minutes in the middle backup. Finally the data is backed up by FlashBackup again. FlashBackup has to back up the entire drive but due to having to only handle metadata up front it is able to backup much faster than the other two methods. This also demonstrates that generating cab files from the EV data can improve backup performance as well.

Backup and Recovery of Enterprise Vault with Backup Exec 12 for Windows Servers

Backup Exec 12 Agent for Enterprise Vault – Best Practices

What is the Backup Exec Agent for Enterprise Vault?

The Symantec Backup Exec for Windows Servers Agent for Enterprise Vault is installed as a separate, add-on component of Backup Exec for Windows Servers. The Agent for Enterprise Vault is designed as a multipurpose Backup Agent that assists customers of both Backup Exec and Enterprise Vault in protecting all critical aspects of the Enterprise Vault infrastructure.

The Agent for Enterprise Vault includes all the functionality needed to protect volumes, files, folders, and databases used by Enterprise Vault. A single Agent enables protection of both the SQL databases and the file system objects necessary for complete protection.

The EV Agent can help provide a periodic data protection and 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 EV Agent enables users to do the following:

- Back up and recovery of Enterprise Vault archives from open or closed Vault Store Partitions
- Back up and recovery of one, several, or all Enterprise Vault Indexes
- Back up individual Enterprise Vault servers from within an Enterprise Vault site
- Back up entire Enterprise Vault sites
- Remove cumbersome scripting that is usually associated with correct backup and recovery of Enterprise Vault components or Sites.

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 Vault Store Databases
- Open and Closed Partitions within a Vault Store
- Enterprise Vault indexing files

How does the Backup Exec Agent for Enterprise Vault work?

Backup Exec combines several aspects of existing technology (Microsoft SQL Server backup, file system protection) with new and innovative integration with Enterprise Vault. This integration and packaging allows Backup Exec to understand the infrastructure in use by Enterprise Vault, while offering complete protection for all EV components.

Backup Exec integrates with Enterprise Vault in the following ways:

- Determines the infrastructure (physical servers, server roles) in use by Enterprise Vault.
- Can start and stop Enterprise Vault's Backup Mode for consistent database, partition, and index backups
- Automatically protects items that depend on each other, for example, selects Open Partitions when the controlling Vault Store Database is selected for backup

Because Enterprise Vault can often be configured as a distributed application, it is important to note that each Enterprise Vault server needs to have an Agent for Enterprise Vault installed on it. The integration between BE and Enterprise Vault will be incomplete and will not show the entire Enterprise Vault infrastructure if only some of the servers participating in the EV infrastructure have the Agent for Enterprise Vault installed.

Licensing Scenarios

The Enterprise Vault Agent should be installed on every machine that participates in the Enterprise Vault infrastructure. It is important to note that while Enterprise Vault s licensed per-user or per-gigabyte, Backup Exec is resource-centric and is licensed per physical or virtual server. Enterprise Vault administrators and Backup Exec Administrators need to be aware of the differences in licensing methodologies between the two products.

Here are some examples scenarios.

Single Server

- If the entire EV infrastructure (Directory database, Vault Store Database, is installed on a single server, a single Agent for Enterprise Vault should be purchased and deployed on that server.

One Server for Directory/Vault Store Databases, One Server for Partitions/Indexes

- In this example environment, a standalone SQL Server machine exists, and hosts the database portions of Enterprise Vault (Directory database and Vault Store

Database), along with a single server which hosts Open/Closed partitions and Index directories. This configuration will require two (2) Agents for Enterprise Vault, one for each machine participating in the infrastructure.

One Server for Directory Database, One Server for Vault Store Databases, Multiple Servers for Vault Store Partitions and Indexes

- In this example, the Enterprise Vault infrastructure is very distributed. The database servers are separated, one server hosting the Directory database, and one server hosting Vault Store Databases. In addition, three (3) additional systems are file servers and contain Open and Closed partitions. This configuration would require five (5) Agents for Enterprise Vault, one for each of the database servers, and one for each file server/index location server that exists in the infrastructure.

Important Note:

The Agent for Enterprise Vault includes both file system as well as Microsoft SQL Server backup and recovery capabilities. The customer doesn't need to know which systems include specific components; she just installed the Agent for Enterprise Vault on all systems participating in the Enterprise Vault infrastructure.

Since the Agent for Enterprise Vault includes Microsoft SQL Server backup and recovery capability, if the user has additional SQL Server databases on the same machine as the Enterprise Vault Directory/Vault Store databases, those other databases can be protected in exactly the same way as the protection provided by the Backup Exec Agent for SQL Server.

About the Backup Exec Enterprise Vault Agent

The Backup Exec Agent for Enterprise Vault provides data protection for key Enterprise Vault components, such as the following:

- Sites
- Servers
- Databases
- Indexes
- Vault Store partitions

The EV Agent can help provide a periodic data protection and 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 EV Agent enables users to do the following:

- Back up and recovery of Enterprise Vault archives from open or closed vault store partitions

- Back up and recovery of one, several, or all Enterprise Vault Indexes
- Back up individual Enterprise Vault servers from within an Enterprise Vault site
- Back up entire Enterprise Vault sites

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 Vault Store Databases
- Open and Closed Partitions within a Vault Store
- Enterprise Vault indexing files

Backup Exec backs up Enterprise Vault by placing the Enterprise Vault services in "backup" or read-only mode. This is one of the advantages to the Agent for Enterprise Vault and eliminates complicated pre-and post backup job scripting to enter and leave Backup Mode. Backup Mode ensures that databases, indexes, and Vault Stores are synchronized when the backup job runs. Because the Enterprise Vault Services are placed in Backup Mode, Enterprise Vault does not archive new data while the backup is occurring. After the backup job successfully completes, Backup Exec takes the Enterprise Vault services out of Backup Mode so that archival operations can continue.

Symantec recommends performing weekly or monthly full backups, with daily incremental/differential backups, to minimize the downtime on the Enterprise Vault infrastructure.

When specific Enterprise Vault components are selected for backup, other Enterprise Vault components are automatically backed up at the same time. This is to maintain consistency between related items (for example, a Vault Store database and an Open Partition), with the overall goal of reducing consistency checks on restored data in the event of a disaster.

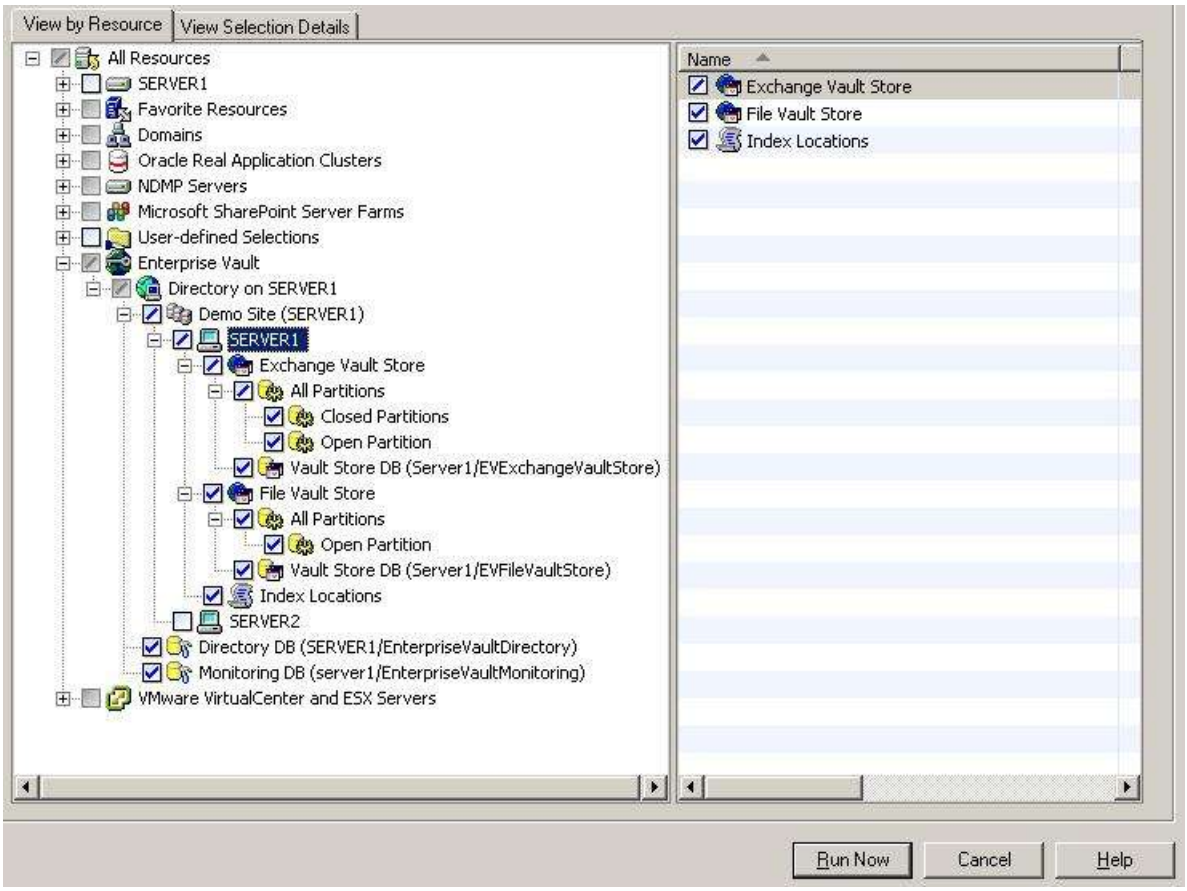

Figure 1 – configuring a backup of Enterprise Vault in Backup Exec, with the Backup Exec Agent for Enterprise Vault providing information for the Directory, Vault Stores, Index Locations, and databases.

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 server or servers	Directory database or databases	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.
Open partition	Vault store database	Backup Exec automatically backs up the vault store database that is associated with the open partition.

Figure 2 – Implicit backup functionality built into the Agent for Enterprise Vault

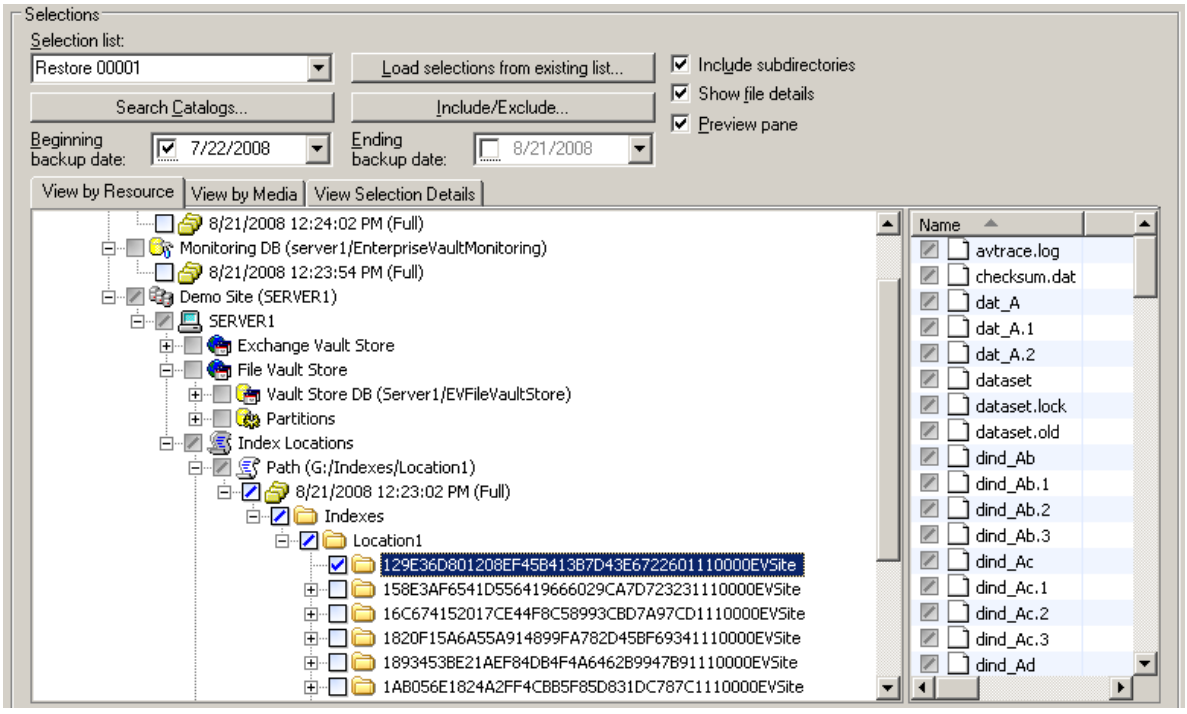

Figure 3 – Restoring an individual Index folder from a Backup Exec backup done with the Agent for Enterprise Vault

Backup Exec Enterprise Vault Agent and NDMP Filers

The EV 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 EV 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 EV 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.

You can find a list of compatible operating systems, platforms, and applications at the following URL: **http://entsupport.Symantec.com/umi/V-269-1**

Important Note about Logon Accounts in Backup Exec and the EV Service Account

The Enterprise Vault Agent for Backup Exec needs to have proper credentials to access, backup, and recover EV components. 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 for Backup Exec. Otherwise, you may have to enter proper credentials for each Enterprise Vault resource that you select for restore.

Restoring Enterprise Vault Components

Backup Exec can restore individual components of Enterprise Vault, like Vault Store databases or individual Index files. However, depending on the scope of the recovery, there are some procedures that you should be aware of before starting a recovery.

Before you restore Enterprise Vault sites, servers or other components, you should have the following items installed on the destination computer:

- Enterprise Vault
- Backup Exec Agent for Enterprise Vault or the Backup Exec Remote Agent for Windows Servers

You can individually restore Enterprise Vault components such as the Directory and Monitoring databases, vault store databases, indexes, and partitions. Before you begin the restore, these components must exist on the destination Enterprise Vault server. If they do not exist, you must create them by using Enterprise Vault before you begin the restore job.

When you restore an entire 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.

After you restore Enterprise Vault, you may 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. See your Enterprise Vault documentation for complete documentation on how to run the Enterprise Vault recovery tools.

How to use "Force the Restore of Directory and Monitoring Databases" Option on Restore Jobs

In the restore job options for an Enterprise Vault recovery, the option "Force the Restore of Directory and Monitoring Databases" exists. This option is for use when you don't want to manually stop the Enterprise Vault Admin and Directory Services on the computer to which you want to restore a specific database. This option takes the shared Directory and Monitoring databases offline, so Backup Exec can replace them during a restore job. This option results in Enterprise Vault Admin and Directory services on all related Enterprise Vault servers to drop the connection to the Directory database that you restore. It also drops the connection to the monitoring database. 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 estored databases and Enterprise Vault begins archival operations again.

About automatic redirection of Enterprise Vault components under an Enterprise Vault server

You can move a vault store database or a partition to a new location that differs from where it was originally backed up. During restores of the vault store database or a partition, the EV Agent detects the location change. It then automatically redirects the restores of these components to the new location.

Additional Information and Resources

You can find additional resources regarding the Backup Exec Agent for Enterprise Servers at the Backup Exec home page at **http://www.BackupExec.com** or in the Backup Exec Administration Guide which ships with every copy of the Backup Exec.

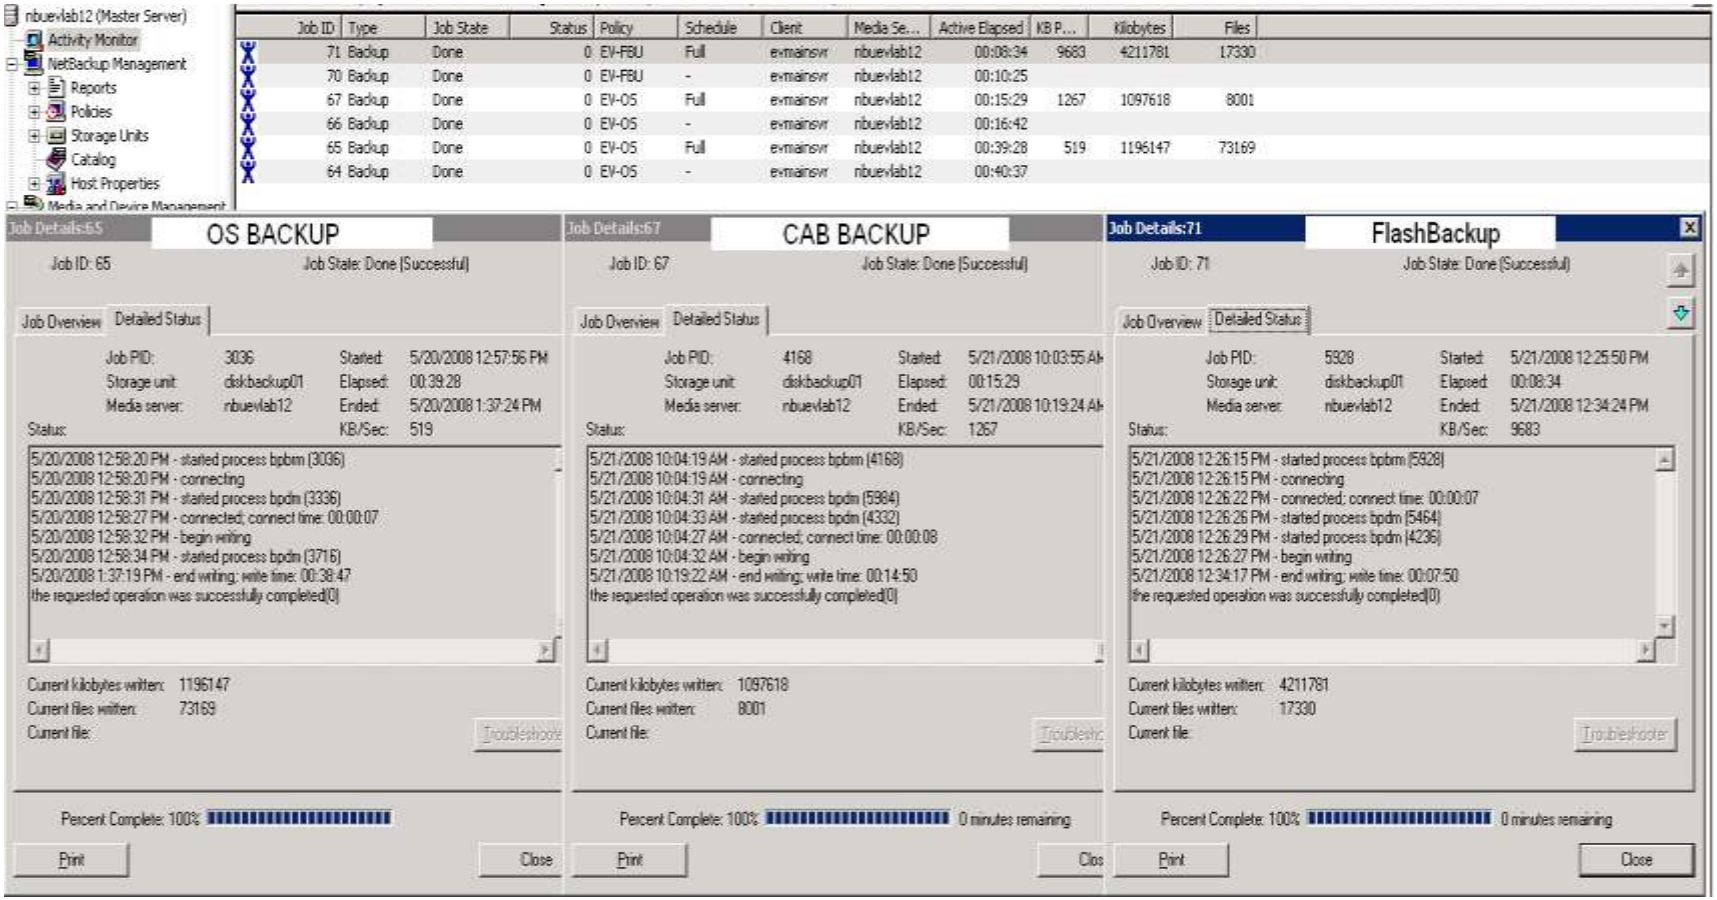
nbuevlab12 (Master Server)
Activity Monitor
NetBackup Management
Reports
Policies
Storage Units
Catalog
Host Properties
Job ID | Type | Job State | Status | Policy | Schedule | Client | Media Se... | Active Elapsed | KB P... | Kilobytes | Files
71 Backup Done 0 EV-FBU Full evmainsvr nbuevlab12 00:08:34 9683 4211781 17330
70 Backup Done 0 EV-FBU - evmainsvr nbuevlab12 00:10:25
67 Backup Done 0 EV-OS Full evmainsvr nbuevlab12 00:15:29 1267 1097618 8001
66 Backup Done 0 EV-OS - evmainsvr nbuevlab12 00:16:42
65 Backup Done 0 EV-OS Full evmainsvr nbuevlab12 00:39:28 519 1196147 73169
64 Backup Done 0 EV-OS - evmainsvr nbuevlab12 00:40:37
Job Details:65
OS BACKUP
Job ID: 65
Job State: Done (Successful)
Job Overview
Detailed Status
Job PID: 3036
Storage unit: diskbackup01
Media server: nbuevlab12
Started: 5/20/2008 12:57:56 PM
Elapsed: 00:39:28
Ended: 5/20/2008 1:37:24 PM
Status:
KB/Sec: 519
5/20/2008 12:58:20 PM - started process bpbrm (3036)
5/20/2008 12:58:20 PM - connecting
5/20/2008 12:58:31 PM - started process bpdm (3336)
5/20/2008 12:58:27 PM - connected; connect time: 00:00:07
5/20/2008 12:58:32 PM - begin writing
5/20/2008 12:58:34 PM - started process bpdm (3716)
5/20/2008 1:37:19 PM - end writing; write time: 00:38:47
the requested operation was successfully completed(0)
Current kilobytes written: 1196147
Current files written: 73169
Current file:
Troubleshoot
Percent Complete: 100%
Print
Close
Job Details:67
CAB BACKUP
Job ID: 67
Job State: Done (Successful)
Job Overview
Detailed Status
Job PID: 4168
Storage unit: diskbackup01
Media server: nbuevlab12
Started: 5/21/2008 10:03:55 AM
Elapsed: 00:15:29
Ended: 5/21/2008 10:19:24 AM
Status:
KB/Sec: 1267
5/21/2008 10:04:19 AM - started process bpbrm (4168)
5/21/2008 10:04:19 AM - connecting
5/21/2008 10:04:31 AM - started process bpdm (5984)
5/21/2008 10:04:33 AM - started process bpdm (4332)
5/21/2008 10:04:27 AM - connected; connect time: 00:00:08
5/21/2008 10:04:32 AM - begin writing
5/21/2008 10:19:22 AM - end writing; write time: 00:14:50
the requested operation was successfully completed(0)
Current kilobytes written: 1097618
Current files written: 8001
Current file:
Troubleshoot
Percent Complete: 100%
0 minutes remaining
Print
Close
Job Details:71
FlashBackup
Job ID: 71
Job State: Done (Successful)
Job Overview
Detailed Status
Job PID: 5928
Storage unit: diskbackup01
Media server: nbuevlab12
Started: 5/21/2008 12:25:50 PM
Elapsed: 00:08:34
Ended: 5/21/2008 12:34:24 PM
Status:
KB/Sec: 9683
5/21/2008 12:26:15 PM - started process bpbrm (5928)
5/21/2008 12:26:15 PM - connecting
5/21/2008 12:26:22 PM - connected; connect time: 00:00:07
5/21/2008 12:26:26 PM - started process bpdm (5464)
5/21/2008 12:26:29 PM - started process bpdm (4236)
5/21/2008 12:26:27 PM - begin writing
5/21/2008 12:34:17 PM - end writing; write time: 00:07:50
the requested operation was successfully completed(0)
Current kilobytes written: 4211781
Current files written: 17330
Current file:
Troubleshooter
Percent Complete: 100%
0 minutes remaining
Print
Close

Supporting Documentation -- NetBackup

NetBackup (tm) 6.0 Network Data Management Protocol (NDMP) System Administrator's Guide for UNIX and Windows
http://seer.entsupport.Symantec.com/docs/279269.htm

NetBackup (tm) 6.0 Advanced Client Configuration
http://seer.entsupport.Symantec.com/docs/279211.htm

NetBackup (tm) 6.0 Advanced Client System Administrator's Guide for UNIX and Windows
http://seer.entsupport.Symantec.com/docs/279289.htm

NetBackup (tm) 6.0 for Microsoft SQL Server System Administrator's Guide
http://seer.entsupport.Symantec.com/docs/279277.htm

NetBackup (tm) 6.0 System Administrator's Guide for Windows, Volume 1
http://seer.entsupport.Symantec.com/docs/279265.htm

NetBackup (tm) 6.5 NDMP Administrator's Guide for UNIX, Linux, and Windows
http://seer.entsupport.Symantec.com/docs/290205.htm

NetBackup (tm) Snapshot Client Configuration
http://seer.entsupport.Symantec.com/docs/288300.htm

NetBackup (tm) 6.5 Snapshot Client Administrator's Guide
http://seer.entsupport.Symantec.com/docs/290224.htm

NetBackup (tm) 6.5 for Microsoft SQL Server Administrator's Guide
http://seer.entsupport.Symantec.com/docs/290212.htm

NetBackup (tm) 6.5 Administrator's Guide for Windows, Volume 1
http://seer.entsupport.Symantec.com/docs/290203.htm

Supporting Documentation – Enterprise Vault

Backup Procedures for Enterprise Vault™
http://seer.entsupport.Symantec.com/docs/284357.htm

Enterprise Vault ™ 7.0 Administrator's Guide
http://support.Symantec.com/docs/286423

Enterprise Vault ™ 7.0 Installing and Configuring Guide
http://support.Symantec.com/docs/287642

Enterprise Vault ™ 7.0 Introduction and Planning Guide
http://support.Symantec.com/docs/287643

Enterprise Vault ™ 2007 Administrators Guide
http://seer.entsupport.Symantec.com/docs/289858.htm

Enterprise Vault ™ 2007 Installing and Configuring Guide
http://support.Symantec.com/docs/289688

Enterprise Vault ™ 2007 Introduction and Planning Guide
http://support.Symantec.com/docs/289860

How to set the Enterprise Vault™ Services to Read Only Mode
http://seer.entsupport.Symantec.com/docs/285283.htm

How to Perform an Online Backup by Setting the Enterprise Vault ™ Services to Read Only
http://seer.entsupport.Symantec.com/docs/284361.htm

Backing up Enterprise Vault™ in a Clustered Environment
http://ftp.support.Symantec.com/pub/support/products/Exchange_Mailbox_Archiving_Unit/285839.pdf

Sample Batch File Template for Backing Up in a Clustered Environment
http://seer.entsupport.Symantec.com/docs/286317.htm

What Happens When Closing an Indexing Location and Opening a New One
http://seer.entsupport.Symantec.com/docs/294679.htm

How to Move Enterprise Vault Indexes that are Managed by the Same Server
http://seer.entsupport.Symantec.com/docs/273141.htm