



Isilon OneFS

Version 7.1

Backup and recovery guide

Copyright © 2013-2014 EMC Corporation. All rights reserved. Published in USA.

Published March, 2014

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>). For documentation on EMC Data Domain products, go to the EMC Data Domain Support Portal (<https://my.datadomain.com>).

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

Chapter 1	Introduction to this guide	7
	About this guide.....	8
	Isilon scale-out NAS overview.....	8
	Where to go for support.....	8
Chapter 2	OneFS backup and recovery	9
	OneFS backup and recovery overview.....	10
	SyncIQ backup and recovery overview.....	10
	NDMP backup and recovery overview.....	10
Chapter 3	Data replication with SyncIQ	13
	Replication policies and jobs.....	14
	Source and target cluster association.....	15
	Full and differential replication.....	15
	Controlling replication job resource consumption.....	15
	Replication reports.....	16
	Replication snapshots.....	16
	Source cluster snapshots.....	16
	Target cluster snapshots.....	17
	Data failover and failback with SyncIQ.....	17
	Data failover.....	18
	Data failback.....	18
	Replication and backup with SmartLock.....	19
	SmartLock replication and backup limitations.....	19
	Recovery times and objectives for SyncIQ.....	20
	SyncIQ license functionality.....	20
Chapter 4	Backing up data with SyncIQ	21
	Creating replication policies.....	22
	Excluding directories in replication.....	22
	Excluding files in replication.....	23
	File criteria options.....	23
	Configure default replication policy settings.....	25
	Create a replication policy.....	25
	Create a SyncIQ domain.....	30
	Assess a replication policy.....	31
	Managing replication to remote clusters.....	31
	Start a replication job.....	31
	Pause a replication job.....	32
	Resume a replication job.....	32
	Cancel a replication job.....	32
	View active replication jobs.....	32
	View replication performance information.....	32
	Replication job information.....	33
	Managing failed replication jobs.....	33
	Resolve a replication policy.....	33
	Reset a replication policy.....	34

	Perform a full or differential replication.....	34
	Managing replication policies.....	35
	Modify a replication policy.....	35
	Delete a replication policy.....	35
	Enable or disable a replication policy.....	36
	View replication policies.....	36
	Replication policy information.....	36
	Replication policy settings.....	37
	Managing replication to the local cluster.....	39
	Cancel replication to the local cluster.....	40
	Break local target association.....	40
	View replication policies targeting the local cluster.....	40
	Remote replication policy information.....	40
Chapter 5	Recovering data with SyncIQ	43
	Initiating data failover and failback with SyncIQ.....	44
	Fail over data to a secondary cluster.....	44
	Revert a failover operation.....	44
	Fail back data to a primary cluster.....	45
	Performing disaster recovery for SmartLock directories.....	46
	Recover SmartLock directories on a target cluster.....	46
	Migrate SmartLock directories.....	47
Chapter 6	NDMP backup	49
	NDMP two way backup.....	50
	Snapshot-based incremental backups.....	50
	NDMP protocol support.....	51
	Supported DMAs.....	51
	NDMP hardware support.....	52
	NDMP backup limitations.....	52
	NDMP performance recommendations.....	52
	Excluding files and directories from NDMP backups.....	54
Chapter 7	Backing up and recovering data with NDMP	57
	NDMP backup and recovery tasks.....	58
	Configuring basic NDMP backup settings.....	58
	NDMP backup settings.....	58
	View NDMP backup settings.....	58
	Configure and enable NDMP backup.....	58
	Disable NDMP backup.....	59
	Managing NDMP user accounts.....	59
	Create an NDMP user account.....	59
	View NDMP user accounts.....	59
	Modify the password of an NDMP user account.....	59
	Delete an NDMP user account.....	60
	Managing NDMP backup devices.....	60
	NDMP backup device settings.....	60
	Detect NDMP backup devices.....	60
	View NDMP backup devices.....	61
	Modify the name of an NDMP backup device.....	61
	Delete an entry for an NDMP backup device.....	61
	Managing NDMP backup ports.....	62
	NDMP backup port settings.....	62

View NDMP backup ports.....	63
Modify NDMP backup port settings.....	63
Enable or disable an NDMP backup port.....	63
Managing NDMP backup sessions.....	63
NDMP session information.....	63
View NDMP sessions.....	65
End an NDMP session.....	65
Managing restartable backups.....	65
Configure restartable backups.....	65
View restartable backup contexts.....	66
Delete a restartable backup context.....	66
Configure restartable backup settings.....	67
View restartable backup settings.....	67
Sharing tape drives between clusters.....	67
Managing default NDMP settings.....	67
Set default NDMP settings for a directory.....	68
Modify default NDMP settings for a directory.....	68
View default NDMP settings for directories.....	68
NDMP environment variables.....	69
Managing snapshot based incremental backups.....	71
Enable snapshot-based incremental backups for a directory.....	71
View snapshots for snapshot-based incremental backups.....	71
Delete snapshots for snapshot-based incremental backups.....	71
View NDMP backup logs.....	72
Configuring NDMP backups with EMC NetWorker.....	72
Create a group.....	72
Scan for tape devices.....	72
Configure a library.....	73
Create a data media pool.....	73
Label tape devices.....	74
Create a metadata media pool.....	74
Create a client.....	75
Configuring NDMP backup with Symantec NetBackup.....	76
Add an NDMP host.....	76
Configure storage devices.....	77
Create a volume pool.....	78
Inventory a robot.....	78
Create a NetBackup policy.....	79
Configuring NDMP backup with CommVault Simpana.....	80
Add a NAS client.....	80
Add an NDMP library.....	80
Create a storage policy.....	81
Assign a storage policy and schedule to a client.....	82

CHAPTER 1

Introduction to this guide

This section contains the following topics:

- ◆ [About this guide](#).....8
- ◆ [Isilon scale-out NAS overview](#).....8
- ◆ [Where to go for support](#).....8

About this guide

This guide describes how to back up and recover data on Isilon clusters through either the SyncIQ software module or the Network Data Management Protocol (NDMP).

Isilon scale-out NAS overview

The EMC Isilon scale-out NAS storage platform combines modular hardware with unified software to harness unstructured data. Powered by the distributed OneFS operating system, an EMC Isilon cluster delivers a scalable pool of storage with a global namespace.

The platform's unified software provides centralized web-based and command-line administration to manage the following features:

- ◆ A symmetrical cluster that runs a distributed file system
- ◆ Scale-out nodes that add capacity and performance
- ◆ Storage options that manage files, block data, and tiering
- ◆ Flexible data protection and high availability
- ◆ Software modules that control costs and optimize resources

Where to go for support

You can contact EMC Isilon Technical Support for any questions about EMC Isilon products.

Online Support	Live Chat Create a Service Request
Telephone Support	United States: 800-782-4362 (1-800-SVC-4EMC) Canada: 800-543-4782 Worldwide: +1-508-497-7901 For local phone numbers in your country, see EMC Customer Support Centers .
Help with online support	For questions specific to EMC Online Support registration or access, email support@emc.com .

CHAPTER 2

OneFS backup and recovery

This section contains the following topics:

- ◆ [OneFS backup and recovery overview](#)..... 10
- ◆ [SyncIQ backup and recovery overview](#)..... 10
- ◆ [NDMP backup and recovery overview](#)..... 10

OneFS backup and recovery overview

You can back up data stored on Isilon clusters to another Isilon cluster or a tape device.

You can back up data to an Isilon cluster with the SyncIQ software module. SyncIQ enables you to recover backed up data through failover and failback. Failover enables you to access data on the cluster it was backed up to. After you fail over, you can fail back to resume accessing your data on the cluster it was backed up from.

You can back up data to a tape device over NDMP. After you back up data to a tape device, you can restore the data to any Isilon cluster.

SyncIQ backup and recovery overview

OneFS enables you to replicate data from one Isilon cluster to another through the SyncIQ software module. You must activate a SyncIQ license on both Isilon clusters before you can replicate data between them.

You can replicate data at the directory level while optionally excluding specific files and sub-directories from being replicated. SyncIQ creates and references snapshots to replicate a consistent point-in-time image of a root directory. Metadata such as access control lists (ACLs) and alternate data streams (ADS) are replicated along with data.

SyncIQ enables you to maintain a consistent backup copy of your data on another Isilon cluster. SyncIQ offers automated failover and failback capabilities that enable you to continue operations on another Isilon cluster if a primary cluster becomes unavailable.

NDMP backup and recovery overview

In OneFS, you can back up and restore file-system data through the Network Data Management Protocol (NDMP). From a backup server, you can direct backup and recovery processes between an Isilon cluster and backup devices such as tape devices, media servers, and virtual tape libraries (VTLs).

OneFS supports both NDMP three-way backup and NDMP two-way backup. During an NDMP three-way backup operation, a data management application (DMA) on a backup server instructs the cluster to start backing up data to a tape media server that is either attached to the LAN or directly attached to the DMA.

During a two-way NDMP backup, a DMA on a backup server instructs a Backup Accelerator node on the cluster to start backing up data to a tape media server that is attached to the Backup Accelerator node.

NDMP two-way backup is the most efficient method in terms of cluster resource consumption. However, NDMP two-way backup requires that you attach one or more Backup Accelerator nodes to the cluster.

In both the NDMP two-way and three-way backup models, file history data is transferred from the cluster to the backup server. Before a backup begins, OneFS creates a snapshot of the targeted directory, then backs up the snapshot, which ensures that the backup image represents a specific point in time.

You do not need to activate a SnapshotIQ license on the cluster to perform NDMP backups. If you have activated a SnapshotIQ license on the cluster, you can generate a snapshot through the SnapshotIQ tool, and then back up the same snapshot to multiple tape devices. If you back up a SnapshotIQ snapshot, OneFS does not create another snapshot for the backup.

Note

If you are backing up SmartLock directories for compliance purposes, it is recommended that you do not specify autocommit time periods for the SmartLock directories. This is because, depending on the autocommit period, files in the SmartLock directories may still be subject to change.

CHAPTER 3

Data replication with SyncIQ

This section contains the following topics:

- ◆ [Replication policies and jobs](#).....14
- ◆ [Replication snapshots](#)..... 16
- ◆ [Data failover and failback with SyncIQ](#).....17
- ◆ [Replication and backup with SmartLock](#)..... 19
- ◆ [Recovery times and objectives for SyncIQ](#).....20
- ◆ [SyncIQ license functionality](#)..... 20

Replication policies and jobs

Data replication is coordinated according to replication policies and jobs. Replication policies specify what data is replicated, where the data is replicated to, and how often the data is replicated. Replication jobs are the operations that replicate data from one Isilon cluster to another. SyncIQ generates replication jobs according to replication policies.

A replication policy specifies two clusters: the source and the target. The cluster on which the replication policy exists is the source cluster. The cluster that data is being replicated to is the target cluster. When a replication policy starts, SyncIQ generates a replication job for the policy. When a replication job runs, files from a directory on the source cluster are replicated to a directory on the target cluster; these directories are known as source and target directories.

After the first replication job created by a replication policy finishes, the target directory and all files contained in the target directory are set to a read-only state, and can be modified only by other replication jobs belonging to the same replication policy. There is no limit to the number of replication policies that can exist on a cluster.

Note

To prevent permissions errors, make sure that ACL policy settings are the same across source and target clusters.

You can create two types of replication policies: synchronization policies and copy policies. A synchronization policy maintains an exact replica of the source directory on the target cluster. If a file or sub-directory is deleted from the source directory, the file or directory is deleted from the target cluster when the policy is run again.

You can use synchronization policies to fail over and fail back data between source and target clusters. When a source cluster becomes unavailable, you can fail over data on a target cluster and make the data available to clients. When the source cluster becomes available again, you can fail back the data to the source cluster.

A copy policy maintains recent versions of the files that are stored on the source cluster. However, files that are deleted on the source cluster are not deleted from the target cluster. Failback is not supported for copy policies. Copy policies are most commonly used for archival purposes.

Copy policies enable you to remove files from the source cluster without losing those files on the target cluster. Deleting files on the source cluster improves performance on the source cluster while maintaining the deleted files on the target cluster. This can be useful if, for example, your source cluster is being used for production purposes and your target cluster is being used only for archiving.

After creating a job for a replication policy, SyncIQ must wait until the job completes before it can create another job for the policy. Any number of replication jobs can exist on a cluster at a given time; however, only five replication jobs can run on a source cluster at the same time. If more than five replication jobs exist on a cluster, the first five jobs run while the others are queued to run. The number of replication jobs that a single target cluster can support concurrently is dependent on the number of workers available on the target cluster.

You can replicate any number of files and directories with a single replication job. You can prevent a large replication job from overwhelming the system by limiting the amount of cluster resources and network bandwidth that data synchronization is allowed to consume. Because each node in a cluster is able to send and receive data, the speed at which data is replicated increases for larger clusters.

Source and target cluster association

SyncIQ associates a replication policy with a target cluster by marking the target cluster when the job runs for the first time. Even if you modify the name or IP address of the target cluster, the mark persists on the target cluster. When a replication policy is run, SyncIQ checks the mark to ensure that data is being replicated to the correct location.

On the target cluster, you can manually break an association between a replication policy and target directory. Breaking the association between a source and target cluster causes the mark on the target cluster to be deleted. You might want to manually break a target association if an association is obsolete. If you break the association of a policy, the policy is disabled on the source cluster and you cannot run the policy. If you want to run the disabled policy again, you must reset the replication policy.

Note

Breaking a policy association causes either a full or differential replication to occur the next time you run the replication policy. During a full or differential replication, SyncIQ creates a new association between the source and target clusters. Depending on the amount of data being replicated, a full or differential replication can take a very long time to complete.

Full and differential replication

If a replication policy encounters an issue that cannot be fixed (for example, if the association was broken on the target cluster), you might need to reset the replication policy. If you reset a replication policy, SyncIQ performs either a full or differential replication the next time the policy is run. You can specify the type of replication that SyncIQ performs.

During a full replication, SyncIQ transfers all data from the source cluster regardless of what data exists on the target cluster. A full replication consumes large amounts of network bandwidth and can take a very long time to complete. However, a full replication is less strenuous on CPU usage than a differential replication.

During a differential replication, SyncIQ first checks whether a file already exists on the target cluster and then transfers only data that does not already exist on the target cluster. A differential replication consumes less network bandwidth than a full replication; however, differential replications consume more CPU. Differential replication can be much faster than a full replication if there is an adequate amount of available CPU for the differential replication job to consume.

Controlling replication job resource consumption

You can create rules that limit the network traffic created and the rate at which files are sent by replication jobs. You can also specify the number of workers that are spawned by a replication policy to limit the amount of cluster resources that are consumed. Also, you can restrict a replication policy to connect only to a specific storage pool.

You can create network-traffic rules that control the amount of network traffic generated by replication jobs during specified time periods. These rules can be useful if, for example, you want to limit the amount of network traffic created during other resource-intensive operations.

You can create multiple network traffic rules to enforce different limitations at different times. For example, you might allocate a small amount of network bandwidth during peak business hours, but allow unlimited network bandwidth during non-peak hours.

When a replication job runs, OneFS generates workers on the source and target cluster. Workers on the source cluster send data while workers on the target cluster write data. OneFS generates no more than 40 workers for a replication job. You can modify the maximum number of workers generated per node to control the amount of resources that a replication job is allowed to consume. For example, you can increase the maximum number of workers per node to increase the speed at which data is replicated to the target cluster.

You can also reduce resource consumption through file-operation rules that limit the rate at which replication policies are allowed to send files. However, it is recommended that you only create file-operation rules if the files you intend to replicate are predictably similar in size and not especially large.

Replication reports

After a replication job completes, SyncIQ generates a report that contains detailed information about the job, including how long the job ran, how much data was transferred, and what errors occurred.

If a replication report is interrupted, SyncIQ might create a subreport about the progress of the job so far. If the job is then restarted, SyncIQ creates another subreport about the progress of the job until the job either completes or is interrupted again. SyncIQ creates a subreport each time the job is interrupted until the job completes successfully. If multiple subreports are created for a job, SyncIQ combines the information from the subreports into a single report.

SyncIQ routinely deletes replication reports. You can specify the maximum number of replication reports that SyncIQ retains and the length of time that SyncIQ retains replication reports. If the maximum number of replication reports is exceeded on a cluster, SyncIQ deletes the oldest report each time a new report is created.

You cannot customize the content of a replication report.

Note

If you delete a replication policy, SyncIQ automatically deletes any reports that were generated for that policy.

Replication snapshots

SyncIQ generates snapshots to facilitate replication, failover, and failback between Isilon clusters. Snapshots generated by SyncIQ can also be used for archival purposes on the target cluster.

Source cluster snapshots

SyncIQ generates snapshots on the source cluster to ensure that a consistent point-in-time image is replicated and that unaltered data is not sent to the target cluster.

Before running a replication job, SyncIQ creates a snapshot of the source directory. SyncIQ then replicates data according to the snapshot rather than the current state of the cluster, allowing users to modify source-directory files while ensuring that an exact point-in-time image of the source directory is replicated.

For example, if a replication job of `/ifs/data/dir/` starts at 1:00 PM and finishes at 1:20 PM, and `/ifs/data/dir/file` is modified at 1:10 PM, the modifications are not reflected on the target cluster, even if `/ifs/data/dir/file` is not replicated until 1:15 PM.

You can replicate data according to a snapshot generated with the SnapshotIQ tool. If you replicate data according to a SnapshotIQ snapshot, SyncIQ does not generate another snapshot of the source directory. This method can be useful if you want to replicate identical copies of data to multiple Isilon clusters.

SyncIQ generates source snapshots to ensure that replication jobs do not transfer unmodified data. When a job is created for a replication policy, SyncIQ checks whether it is the first job created for the policy. If it is not the first job created for the policy, SyncIQ compares the snapshot generated for the earlier job with the snapshot generated for the new job.

SyncIQ replicates only data that has changed since the last time a snapshot was generated for the replication policy. When a replication job is completed, SyncIQ deletes the previous source-cluster snapshot and retains the most recent snapshot until the next job is run.

Target cluster snapshots

When a replication job is run, SyncIQ generates a snapshot on the target cluster to facilitate failover operations. When the next replication job is created for the replication policy, the job creates a new snapshot and deletes the old one.

If a SnapshotIQ license has been activated on the target cluster, you can configure a replication policy to generate additional snapshots that remain on the target cluster even as subsequent replication jobs run.

SyncIQ generates target snapshots to facilitate failover on the target cluster regardless of whether a SnapshotIQ license has been configured on the target cluster. Failover snapshots are generated when a replication job completes. SyncIQ retains only one failover snapshot per replication policy, and deletes the old snapshot after the new snapshot is created.

If a SnapshotIQ license has been activated on the target cluster, you can configure SyncIQ to generate archival snapshots on the target cluster that are not automatically deleted when subsequent replication jobs run. Archival snapshots contain the same data as the snapshots that are generated for failover purposes. However, you can configure how long archival snapshots are retained on the target cluster. You can access archival snapshots the same way that you access other snapshots generated on a cluster.

Data failover and failback with SyncIQ

SyncIQ enables you to perform automated data failover and failback operations between Isilon clusters. If a cluster is rendered unusable, you can fail over to another Isilon cluster, enabling clients to access their data on the other cluster. If the unusable cluster becomes accessible again, you can fail back to the original Isilon cluster.

For the purposes of explaining failover and failback procedures, the cluster originally accessed by clients is referred to as the primary cluster, and the cluster that client data is originally replicated to is referred to as the secondary cluster. Failover is the process that allows clients to modify data on a secondary cluster. Failback is the process that allows clients to access data on the primary cluster again and begins to replicate data back to the secondary cluster.

Failover and failback can be useful in disaster recovery procedures. For example, if a primary cluster is damaged by a natural disaster, you can migrate clients to a secondary cluster until the primary cluster is repaired and then migrate the clients back to the primary cluster.

You can fail over and fail back to facilitate scheduled cluster maintenance. For example, if you are upgrading the primary cluster, you might want to migrate clients to a secondary cluster until the upgrade is complete and then migrate clients back to the primary cluster.

Note

Data failover and failback is not supported for SmartLock directories.

Data failover

Data failover is the process of preparing data on a secondary cluster to be modified by clients. After you fail over to a secondary cluster, you can redirect clients to modify their data on the secondary cluster.

Before failover is performed, you must create and run a replication policy on the primary cluster. You initiate the failover process on the secondary cluster. Failover is performed per replication policy; to migrate data that is spread across multiple replication policies, you must initiate failover for each replication policy.

You can use any replication policy to fail over. However, if the action of the replication policy is set to copy, any file that was deleted on the primary cluster will be present on the secondary cluster. When the client connects to the secondary cluster, all files that were deleted on the primary cluster will be available to the client.

If you initiate failover for a replication policy while an associated replication job is running, the failover operation completes but the replication job fails. Because data might be in an inconsistent state, SyncIQ uses the snapshot generated by the last successful replication job to revert data on the secondary cluster to the last recovery point.

If a disaster occurs on the primary cluster, any modifications to data that were made after the last successful replication job started are not reflected on the secondary cluster. When a client connects to the secondary cluster, their data appears as it was when the last successful replication job was started.

Data failback

Data failback is the process of restoring clusters to the roles they occupied before a failover operation. After data failback is complete, the primary cluster hosts clients and replicates data to the secondary cluster for backup.

The first step in the failback process is updating the primary cluster with all of the modifications that were made to the data on the secondary cluster. The next step in the failback process is preparing the primary cluster to be accessed by clients. The final step in the failback process is resuming data replication from the primary to the secondary cluster. At the end of the failback process, you can redirect users to resume accessing their data on the primary cluster.

You can fail back data with any replication policy that meets all of the following criteria:

- ◆ The source directory is not a SmartLock directory.
- ◆ The policy has been failed over.
- ◆ The policy is a synchronization policy.
- ◆ The policy does not exclude any files or directories from replication.

Replication and backup with SmartLock

You must ensure that SmartLock directories remain protected during replication and backup operations.

If you are replicating SmartLock directories with SyncIQ, it is recommended that you configure all nodes on the source and target clusters into Network Time Protocol (NTP) peer mode to ensure that the node clocks are synchronized. For compliance clusters, it is recommended that you configure all nodes on the source and target clusters into NTP peer mode before you set the compliance clock to ensure that the compliance clocks are initially set to the same time.

Note

Do not configure SmartLock settings for a target SmartLock directory unless you are no longer replicating data to the directory. Configuring an autocommit time period for a target SmartLock directory can cause replication jobs to fail. If the target SmartLock directory commits a file to a WORM state, and the file is modified on the source cluster, the next replication job will fail because it cannot update the file.

SmartLock replication and backup limitations

Be aware of the limitations of replicating and backing up SmartLock directories with SyncIQ and NDMP.

If the source or target directory of a SyncIQ policy is a SmartLock directory, replication might not be allowed. For more information, see the following table:

Source directory type	Target directory type	Allowed
Non-SmartLock	Non-SmartLock	Yes
Non-SmartLock	SmartLock enterprise	Yes
Non-SmartLock	SmartLock compliance	No
SmartLock enterprise	Non-SmartLock	Yes; however, retention dates and commit status of files will be lost.
SmartLock enterprise	SmartLock enterprise	Yes
SmartLock enterprise	SmartLock compliance	No
SmartLock compliance	Non-SmartLock	No
SmartLock compliance	SmartLock enterprise	No
SmartLock compliance	SmartLock compliance	Yes

If you replicate SmartLock directories to another cluster with SyncIQ, the WORM state of files is replicated. However, SmartLock directory configuration settings are not transferred to the target directory.

For example, if you replicate a directory that contains a committed file that is set to expire on March 4th, the file is still set to expire on March 4th on the target cluster. However, if the directory on the source cluster is set to prevent files from being committed for more than a year, the target directory is not automatically set to the same restriction.

If you back up data to an NDMP device, all SmartLock metadata relating to the retention date and commit status is transferred to the NDMP device. If you restore data to a SmartLock directory on the cluster, the metadata persists on the cluster. However, if the directory that you restore to is not a SmartLock directory, the metadata is lost. You can restore to a SmartLock directory only if the directory is empty.

Recovery times and objectives for SyncIQ

The Recovery Point Objective (RPO) and the Recovery Time Objective (RTO) are measurements of the impacts that a disaster can have on business operations. You can calculate your RPO and RTO for a disaster recovery with replication policies.

RPO is the maximum amount of time for which data is lost if a cluster suddenly becomes unavailable. For an Isilon cluster, the RPO is the amount of time that has passed since the last completed replication job started. The RPO is never greater than the time it takes for two consecutive replication jobs to run and complete.

If a disaster occurs while a replication job is running, the data on the secondary cluster is reverted to the state it was in when the last replication job completed. For example, consider an environment in which a replication policy is scheduled to run every three hours, and replication jobs take two hours to complete. If a disaster occurs an hour after a replication job begins, the RPO is four hours, because it has been four hours since a completed job began replicating data.

RTO is the maximum amount of time required to make backup data available to clients after a disaster. The RTO is always less than or approximately equal to the RPO, depending on the rate at which replication jobs are created for a given policy.

If replication jobs run continuously, meaning that another replication job is created for the policy before the previous replication job completes, the RTO is approximately equal to the RPO. When the secondary cluster is failed over, the data on the cluster is reset to the state it was in when the last job completed; resetting the data takes an amount of time proportional to the time it took users to modify the data.

If replication jobs run on an interval, meaning that there is a period of time after a replication job completes before the next replication job for the policy starts, the relationship between RTO and RPO depends on whether a replication job is running when the disaster occurs. If a job is in progress when a disaster occurs, the RTO is roughly equal to the RPO. However, if a job is not running when a disaster occurs, the RTO is negligible because the secondary cluster was not modified since the last replication job ran, and the failover process is almost instantaneous.

SyncIQ license functionality

You can replicate data to another Isilon cluster only if you activate a SyncIQ license on both the local cluster and the target cluster.

If a SyncIQ license becomes inactive, you cannot create, run, or manage replication policies. Also, all previously created replication policies are disabled. Replication policies that target the local cluster are also disabled. However, data that was previously replicated to the local cluster is still available.

CHAPTER 4

Backing up data with SyncIQ

This section contains the following topics:

- ◆ [Creating replication policies](#) 22
- ◆ [Managing replication to remote clusters](#) 31
- ◆ [Managing failed replication jobs](#) 33
- ◆ [Managing replication policies](#) 35
- ◆ [Managing replication to the local cluster](#) 39

Creating replication policies

You can create replication policies that determine when data is replicated with SyncIQ.

Excluding directories in replication

You can exclude directories from being replicated by replication policies even if the directories exist under the specified source directory.

Note

You cannot fail back replication policies that exclude directories.

By default, all files and directories under the source directory of a replication policy are replicated to the target cluster. However, you can prevent directories under the source directory from being replicated.

If you specify a directory to exclude, files and directories under the excluded directory are not replicated to the target cluster. If you specify a directory to include, only the files and directories under the included directory are replicated to the target cluster; any directories that are not contained in an included directory are excluded.

If you both include and exclude directories, any excluded directories must be contained in one of the included directories; otherwise, the excluded-directory setting has no effect. For example, consider a policy with the following settings:

- ◆ The root directory is `/ifs/data`
- ◆ The included directories are `/ifs/data/media/music` and `/ifs/data/media/movies`
- ◆ The excluded directories are `/ifs/data/archive` and `/ifs/data/media/music/working`

In this example, the setting that excludes the `/ifs/data/archive` directory has no effect because the `/ifs/data/archive` directory is not under either of the included directories. The `/ifs/data/archive` directory is not replicated regardless of whether the directory is explicitly excluded. However, the setting that excludes the `/ifs/data/media/music/working` directory does have an effect, because the directory would be replicated if the setting was not specified.

In addition, if you exclude a directory that contains the source directory, the exclude-directory setting has no effect. For example, if the root directory of a policy is `/ifs/data`, explicitly excluding the `/ifs` directory does not prevent `/ifs/data` from being replicated.

Any directories that you explicitly include or exclude must be contained in or under the specified root directory. For example, consider a policy in which the specified root directory is `/ifs/data`. In this example, you could include both the `/ifs/data/media` and the `/ifs/data/users/` directories because they are under `/ifs/data`.

Excluding directories from a synchronization policy does not cause the directories to be deleted on the target cluster. For example, consider a replication policy that synchronizes `/ifs/data` on the source cluster to `/ifs/data` on the target cluster. If the policy excludes `/ifs/data/media` from replication, and `/ifs/data/media/file` exists on the target cluster, running the policy does not cause `/ifs/data/media/file` to be deleted from the target cluster.

Excluding files in replication

If you do not want specific files to be replicated by a replication policy, you can exclude them from the replication process through file-matching criteria statements. You can configure file-matching criteria statements during the replication-policy creation process.

Note

You cannot fail back replication policies that exclude files.

A file-criteria statement can include one or more elements. Each file-criteria element contains a file attribute, a comparison operator, and a comparison value. You can combine multiple criteria elements in a criteria statement with Boolean "AND" and "OR" operators. You can configure any number of file-criteria definitions.

Configuring file-criteria statements can cause the associated jobs to run slowly. It is recommended that you specify file-criteria statements in a replication policy only if necessary.

Modifying a file-criteria statement will cause a full replication to occur the next time that a replication policy is started. Depending on the amount of data being replicated, a full replication can take a very long time to complete.

For synchronization policies, if you modify the comparison operators or comparison values of a file attribute, and a file no longer matches the specified file-matching criteria, the file is deleted from the target the next time the job is run. This rule does not apply to copy policies.

File criteria options

You can configure a replication policy to exclude files that meet or do not meet specific criteria.

You can specify file criteria based on the following file attributes:

Date created

Includes or excludes files based on when the file was created. This option is available for copy policies only.

You can specify a relative date and time, such as "two weeks ago", or specific date and time, such as "January 1, 2012." Time settings are based on a 24-hour clock.

Date accessed

Includes or excludes files based on when the file was last accessed. This option is available for copy policies only, and only if the global access-time-tracking option of the cluster is enabled.

You can specify a relative date and time, such as "two weeks ago", or specific date and time, such as "January 1, 2012." Time settings are based on a 24-hour clock.

Date modified

Includes or excludes files based on when the file was last modified. This option is available for copy policies only.

You can specify a relative date and time, such as "two weeks ago", or specific date and time, such as "January 1, 2012." Time settings are based on a 24-hour clock.

File name

Includes or excludes files based on the file name. You can specify to include or exclude full or partial names that contain specific text.

The following wildcard characters are accepted:

Note

Alternatively, you can filter file names by using POSIX regular-expression (regex) text. Isilon clusters support IEEE Std 1003.2 (POSIX.2) regular expressions. For more information about POSIX regular expressions, see the BSD man pages.

Table 1 Replication file matching wildcards

Wildcard	Description
*	Matches any string in place of the asterisk. For example, m* matches movies and m123.
[]	Matches any characters contained in the brackets, or a range of characters separated by a dash. For example, b[aei]t matches bat, bet, and bit. For example, 1[4-7]2 matches 142, 152, 162, and 172. You can exclude characters within brackets by following the first bracket with an exclamation mark. For example, b[!ie] matches bat but not bit or bet. You can match a bracket within a bracket if it is either the first or last character. For example, [[c]at matches cat and [at]. You can match a dash within a bracket if it is either the first or last character. For example, car[-s] matches cars and car-.
?	Matches any character in place of the question mark. For example, t?p matches tap, tip, and top.

Path

Includes or excludes files based on the file path. This option is available for copy policies only.

You can specify to include or exclude full or partial paths that contain specified text. You can also include the wildcard characters *, ?, and [].

Size

Includes or excludes files based on their size.

Note

File sizes are represented in multiples of 1024, not 1000.

Type

Includes or excludes files based on one of the following file-system object types:

- ◆ Soft link
- ◆ Regular file
- ◆ Directory

Configure default replication policy settings

You can configure default settings for replication policies. If you do not modify these settings when creating a replication policy, the specified default settings are applied.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Settings**.
2. In the **Default Policy Settings** section, specify how you want replication policies to connect to target clusters by selecting one of the following options:
 - Click **Connect to any nodes in the cluster**.
 - Click **Connect to only the nodes in the subnet and pool if the target cluster name specifies a SmartConnect zone**.
3. Specify which nodes you want replication policies to connect to when a policy is run.

Options	Description
Connect policies to all nodes on a source cluster.	Click Run the policy on all nodes in this cluster .
Connect policies only to nodes contained in a specified subnet and pool.	<ol style="list-style-type: none"> a. Click Run the policy only on nodes in the specified subnet and pool. b. From the Subnet and pool list, select the subnet and pool .

Note

SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.

4. Click **Submit**.

Create a replication policy

You can create a replication policy with SyncIQ that defines how and when data is replicated to another Isilon cluster. Configuring a replication policy is a five-step process.

Configure replication policies carefully. If you modify any of the following policy settings after the policy is run, OneFS performs either a full or differential replication the next time the policy is run:

- ◆ Source directory
- ◆ Included or excluded directories
- ◆ File-criteria statement

- ◆ Target cluster name or address
This applies only if you target a different cluster. If you modify the IP or domain name of a target cluster, and then modify the replication policy on the source cluster to match the new IP or domain name, a full replication is not performed.
- ◆ Target directory

Configure basic policy settings

You must configure basic settings for a replication policy.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Policies**.
2. Click **Create a SyncIQ policy**.
3. In the **Settings** area, in the **Policy name** field, type a name for the replication policy.
4. Optional: In the **Description** field, type a description for the replication policy.
5. In the **Action** area, specify the type of replication policy.

- To copy all files from the source directory to the target directory, click **Copy**.

Note

Failback is not supported for copy policies.

- To copy all files from the source directory to the target directory and delete any files on the target directory that are not in the source directory, click **Synchronize**.
6. In the **Run job** area, specify whether replication jobs will be run.

Options	Description
Run jobs only when manually initiated by a user.	Click Only manually .
Run jobs automatically according to a schedule.	<ol style="list-style-type: none"> a. Click On a schedule. b. Specify a schedule. <p>If you configure a replication policy to run more than once a day, you cannot configure the interval to span across two calendar days. For example, you cannot configure a replication policy to run every hour starting at 7:00 PM and ending at 1:00 AM.</p>
Run jobs automatically every time a change is made to the source directory.	Click Whenever the source is modified .

After you finish

The next step in the process of creating a replication policy is specifying source directories and files.

Specify source directories and files

You must specify the directories and files you want to replicate.

Procedure

1. In the **Source Cluster** area, in the **Source Root Directory** field, type the full path of the source directory that you want to replicate to the target cluster.

You must specify a directory contained in `/ifs`. You cannot specify the `/ifs/.snapshot` directory or subdirectory of it.

2. Optional: Prevent specific subdirectories of the root directory from being replicated.
 - To include a directory, in the **Included Directories** area, click **Add a directory path**.
 - To exclude a directory, in the **Excluded Directories** area, click **Add a directory path**.
3. Optional: Prevent specific files from being replicated by specifying file matching criteria.
 - a. In the **File Matching Criteria** area, select a filter type.
 - b. Select an operator.
 - c. Type a value.

Files that do not meet the specified criteria will not be replicated to the target cluster. For example, if you specify `File Type doesn't match .txt`, SyncIQ will not replicate any files with the `.txt` file extension. If you specify `Created after 08/14/2013`, SyncIQ will not replicate any files created before August 14th, 2013. If you want to specify more than one file matching criterion, you can control how the criteria relate to each other by clicking either **Add an "Or" condition** or **Add an "And" condition**.

4. Specify which nodes you want the replication policy to connect to when the policy is run.

Options	Description
Connect the policy to all nodes in the source cluster.	Click Run the policy on all nodes in this cluster .
Connect the policy only to nodes contained in a specified subnet and pool.	<ol style="list-style-type: none"> a. Click Run the policy only on nodes in the specified subnet and pool. b. From the Subnet and pool list, select the subnet and pool .

Note

SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.

After you finish

The next step in the process of creating a replication policy is specifying the target directory.

Specify the policy target directory

You must specify a target cluster and directory to replicate data to.

Procedure

1. In the **Target Cluster** area, in the **Target Host** field, type one of the following:

- The fully qualified domain name of any node in the target cluster.
- The host name of any node in the target cluster.
- The name of a SmartConnect zone in the target cluster.
- The IPv4 or IPv6 address of any node in the target cluster.
- **localhost**

This will replicate data to another directory on the local cluster.

Note

SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.

2. In the **Target Directory** field, type the absolute path of the directory on the target cluster that you want to replicate data to.

CAUTION

If you specify an existing directory on the target cluster, ensure that the directory is not the target of another replication policy. If this is a synchronization policy, ensure that the directory is empty. All files are deleted from the target of a synchronization policy the first time the policy is run.

If the specified target directory does not already exist on the target cluster, the directory is created the first time the job is run. It is recommended that you do not specify the `/ifs` directory. If you specify the `/ifs` directory, the entire target cluster is set to a read-only state, preventing you from storing any other data on the cluster.

If this is a copy policy, and files in the target directory share the same name as files in the source directory, the target directory files are overwritten when the job is run.

3. If you want replication jobs to connect only to the nodes included in the SmartConnect zone specified by the target cluster, click **Connect only to the nodes within the target cluster SmartConnect Zone**.

After you finish

The next step in the process of creating a replication policy is specifying policy target snapshot settings.

Configure policy target snapshot settings

You can optionally specify how archival snapshots are generated on the target cluster. You can access archival snapshots the same way that you access SnapshotIQ snapshots.

SyncIQ always retains one snapshot on the target cluster to facilitate failover, regardless of these settings.

Procedure

1. To create archival snapshots on the target cluster, in the **Target Snapshots** area, click **Capture snapshots on the target cluster**.
2. Optional: To modify the default alias of the last snapshot created according to the replication policy, in the **Snapshot Alias Name** field, type a new alias.

You can specify the alias name as a snapshot naming pattern. For example, the following naming pattern is valid:

```
%{PolicyName}-on-%{SrcCluster}-latest
```

The previous example produces names similar to the following:

```
newPolicy-on-Cluster1-latest
```

3. Optional: To modify the snapshot naming pattern for snapshots created according to the replication policy, in the **Snapshot Naming Pattern** field, type a naming pattern. Each snapshot generated for this replication policy is assigned a name based on this pattern.

For example, the following naming pattern is valid:

```
%{PolicyName}-from-%{SrcCluster}-at-%H:%M-on-%m-%d-%Y
```

The example produces names similar to the following:

```
newPolicy-from-Cluster1-at-10:30-on-7-12-2012
```

4. Select one of the following options:
 - Click **Snapshots do not expire**.
 - Click **Snapshots expire after...** and specify an expiration period.

After you finish

The next step in the process of creating a replication policy is configuring advanced policy settings.

Configure advanced policy settings

You can optionally configure advanced settings for a replication policy.

Procedure

1. Optional: In the **Worker Threads Per Node** field, specify the maximum number of concurrent processes per node that will perform replication operations.

Note

Do not modify the default setting without consulting Isilon Technical Support.

2. Optional: From the **Log Level** list, select the level of logging you want SyncIQ to perform for replication jobs.

The following log levels are valid, listed from least to most verbose:

- Click **Error**.
 - Click **Notice**.
 - Click **Network Activity**.
 - Click **File Activity**.
3. Optional: If you want SyncIQ to perform a checksum on each file data packet that is affected by the replication policy, select the **Validate File Integrity** check box.

If you enable this option, and the checksum values for a file data packet do not match, SyncIQ retransmits the affected packet.

4. Optional: To modify the length of time SyncIQ retains replication reports for the policy, in the **Keep Reports For** area, specify a length of time.

After the specified expiration period has passed for a report, SyncIQ automatically deletes the report.

Some units of time are displayed differently when you view a report than how they were originally entered. Entering a number of days that is equal to a corresponding value in weeks, months, or years results in the larger unit of time being displayed. For example, if you enter a value of 7 `days`, 1 week appears for that report after it is created. This change occurs because SyncIQ internally records report retention times in seconds and then converts them into days, weeks, months, or years.

5. Optional: Specify whether to record information about files that are deleted by replication jobs by selecting one of the following options:

- Click **Record when a synchronization deletes files or directories**.
- Click **Do not record when a synchronization deletes files or directories**.

This option is applicable for synchronization policies only.

After you finish

The next step in the process of creating a replication policy is saving the replication policy settings.

Save replication policy settings

SyncIQ does not create replication jobs for a replication policy until you save the policy.

Before you begin

Review the current settings of the replication policy. If necessary, modify the policy settings.

Procedure

1. Click **Create Policy**.

After you finish

You can increase the speed at which you can failback a replication policy by creating a SyncIQ domain for the source directory of the policy.

Create a SyncIQ domain

You can create a SyncIQ domain to increase the speed at which failback is performed for a replication policy. Because you can fail back only synchronization policies, it is not necessary to create SyncIQ domains for copy policies.

Failing back a replication policy requires that a SyncIQ domain be created for the source directory. OneFS automatically creates a SyncIQ domain during the failback process. However, if you intend on failing back a replication policy, it is recommended that you create a SyncIQ domain for the source directory of the replication policy while the directory is empty. Creating a domain for a directory that contains less data takes less time.

Procedure

1. Click **Cluster Management** > **Job Operations** > **Job Types**.
2. In the **Job Types** area, in the **DomainMark** row, from the **Actions** column, select **Start Job**.
3. In the **Domain Root Path** field, type the path of a source directory of a replication policy.

4. From the **Type of domain** list, select **SyncIQ**.
5. Ensure that the **Delete domain** check box is cleared.
6. Click **Start Job**.

Assess a replication policy

Before running a replication policy for the first time, you can view statistics on the files that would be affected by the replication without transferring any files. This can be useful if you want to preview the size of the data set that will be transferred if you run the policy.

Note

You can assess only replication policies that have never been run before.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Policies**.
2. In the **SyncIQ Policies** table, in the row of a replication policy, from the **Actions** column, select **Assess Sync**.
3. Click **Data Protection** > **SyncIQ** > **Summary**.
4. After the job completes, in the **SyncIQ Recent Reports** table, in the row of the replication job, click **View Details**.

The report displays the total amount of data that would have been transferred in the **Total Data** field.

Managing replication to remote clusters

You can manually run, view, assess, pause, resume, cancel, resolve, and reset replication jobs that target other clusters.

After a policy job starts, you can pause the job to suspend replication activities. Afterwards, you can resume the job, continuing replication from the point where the job was interrupted. You can also cancel a running or paused replication job if you want to free the cluster resources allocated for the job. A paused job reserves cluster resources whether or not the resources are in use. A cancelled job releases its cluster resources and allows another replication job to consume those resources. No more than five running and paused replication jobs can exist on a cluster at a time. However, an unlimited number of canceled replication jobs can exist on a cluster. If a replication job remains paused for more than a week, SyncIQ automatically cancels the job.

Start a replication job

You can manually start a replication job for a replication policy at any time.

If you want to replicate data according to an existing snapshot, at the OneFS command prompt, run the `isi sync jobs start` command with the `--source-snapshot` option. You cannot replicate data according to snapshots generated by SyncIQ.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Policies**.
2. In the **SyncIQ Policies** table, in the **Actions** column for a job, select **Start Job**.

Pause a replication job

You can pause a running replication job and then resume the job later. Pausing a replication job temporarily stops data from being replicated, but does not free the cluster resources replicating the data.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Summary**.
2. In the **Active Jobs** table, in the **Actions** column for a job, click **Pause Running Job**.

Resume a replication job

You can resume a paused replication job.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Summary**.
2. In the **Currently Running** table, in the **Actions** column for a job, click **Resume Running Job**.

Cancel a replication job

You can cancel a running or paused replication job. Cancelling a replication job stops data from being replicated and frees the cluster resources that were replicating data. You cannot resume a cancelled replication job. To restart replication, you must start the replication policy again.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Summary**.
2. In the **Active Jobs** table, in the **Actions** column for a job, click **Cancel Running Job**.

View active replication jobs

You can view information about replication jobs that are currently running or paused.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Policies**.
2. In the **Active Jobs** table, review information about active replication jobs.

View replication performance information

You can view information about how many files are sent and the amount of network bandwidth consumed by replication policies.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Performance**.
2. In the **Network Performance** and **File Operations** tables, view performance information.

Replication job information

You can view information about replication jobs through the **Active Jobs** table.

Status

The status of the job. The following job statuses are possible:

Running

The job is currently running without error.

Paused

The job has been temporarily paused.

Policy Name

The name of the associated replication policy.

Started

The time the job started.

Elapsed

How much time has elapsed since the job started.

Transferred

The number of files that have been transferred, and the total size of all transferred files.

Source Directory

The path of the source directory on the source cluster.

Target Host

The target directory on the target cluster.

Actions

Displays any job-related actions that you can perform.

Managing failed replication jobs

If a replication job fails due to an error, SyncIQ might disable the corresponding replication policy. For example SyncIQ might disable a replication policy if the IP or hostname of the target cluster is modified. If a replication policy is disabled, the policy cannot be run.

To resume replication for a disabled policy, you must either fix the error that caused the policy to be disabled, or reset the replication policy. It is recommended that you attempt to fix the issue rather than reset the policy. If you believe you have fixed the error, you can return the replication policy to an enabled state by resolving the policy. You can then run the policy again to test whether the issue was fixed. If you are unable to fix the issue, you can reset the replication policy. However, resetting the policy causes a full or differential replication to be performed the next time the policy is run.

Note

Depending on the amount of data being synchronized or copied, a full and differential replications can take a very long time to complete.

Resolve a replication policy

If SyncIQ disables a replication policy due to a replication error, and you fix the issue that caused the error, you can resolve the replication policy. Resolving a replication policy

enables you to run the policy again. If you cannot resolve the issue that caused the error, you can reset the replication policy.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Policies**.
2. In the **Policies** table, in the row for a policy, select **Resolve**.

Reset a replication policy

If a replication job encounters an error that you cannot resolve, you can reset the corresponding replication policy. Resetting a policy causes OneFS to perform a full or differential replication the next time the policy is run. Resetting a replication policy deletes the latest snapshot generated for the policy on the source cluster.

CAUTION

Depending on the amount of data being replicated, a full or differential replication can take a very long time to complete. Reset a replication policy only if you cannot fix the issue that caused the replication error. If you fix the issue that caused the error, resolve the policy instead of resetting the policy.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Policies**.
2. In the **SyncIQ Policies** table, in the row for a policy, select **Reset Sync State**.

Perform a full or differential replication

After you reset a replication policy, you must perform either a full or differential replication.

Before you begin

Reset a replication policy.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in through the root or compliance administrator account.
2. Specify the type of replication you want to perform by running the `isi sync policies modify` command.

- To perform a full replication, disable the `--target-compare-initial-sync` option.

For example, the following command disables differential synchronization for `newPolicy`:

```
isi sync policies modify newPolicy \
--target-compare-initial-sync false
```

- To perform a differential replication, enable the `--target-compare-initial-sync` option.

For example, the following command enables differential synchronization for `newPolicy`:

```
isi sync policies modify newPolicy \
--target-compare-initial-sync true
```

3. Run the policy by running the `isi sync jobs start` command.

For example, the following command runs newPolicy:

```
isi sync jobs start newPolicy
```

Managing replication policies

You can modify, view, enable and disable replication policies.

Modify a replication policy

You can modify the settings of a replication policy.

If you modify any of the following policy settings after a policy runs, OneFS performs either a full or differential replication the next time the policy runs:

- ◆ Source directory
- ◆ Included or excluded directories
- ◆ File-criteria statement
- ◆ Target cluster

This applies only if you target a different cluster. If you modify the IP or domain name of a target cluster, and then modify the replication policy on the source cluster to match the new IP or domain name, a full replication is not performed.
- ◆ Target directory

Procedure

1. Click **Data Protection** › **SyncIQ** › **Policies**.
2. In the **SyncIQ Policies** table, in the row for a policy, click **View/Edit**.
3. In the **View SyncIQ Policy Details** dialog box, click **Edit Policy**.
4. Modify the settings of the replication policy, and then click **Save Changes**.

Delete a replication policy

You can delete a replication policy. Once a policy is deleted, SyncIQ no longer creates replication jobs for the policy. Deleting a replication policy breaks the target association on the target cluster, and allows writes to the target directory.

If you want to temporarily suspend a replication policy from creating replication jobs, you can disable the policy, and then enable the policy again later.

Procedure

1. Click **Data Protection** › **SyncIQ** › **Policies**.
2. In the **SyncIQ Policies** table, in the row for a policy, select **Delete Policy**.
3. In the confirmation dialog box, click **Delete**.

Enable or disable a replication policy

You can temporarily suspend a replication policy from creating replication jobs, and then enable it again later.

Note

If you disable a replication policy while an associated replication job is running, the running job is not interrupted. However, the policy will not create another job until the policy is enabled.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Policies**.
2. In the **SyncIQ Policies** table, in the row for a replication policy, select either **Enable Policy** or **Disable Policy**.

If neither **Enable Policy** nor **Disable Policy** appears, verify that a replication job is not running for the policy. If an associated replication job is not running, ensure that the SyncIQ license is active on the cluster.

View replication policies

You can view information about replication policies.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Policies**.
2. In the **SyncIQ Policies** table, review information about replication policies.

Replication policy information

You can view information about replication policies through the **SyncIQ Policies** table.

Policy Name

The name of the policy.

State

Whether the policy is enabled or disabled.

Last Known Good

When the last successful job ran.

Schedule

When the next job is scheduled to run. A value of **Manual** indicates that the job can be run only manually. A value of **When source is modified** indicates that the job will be run whenever changes are made to the source directory.

Source Directory

The path of the source directory on the source cluster.

Target Host : Directory

The IP address or fully qualified domain name of the target cluster and the full path of the target directory.

Actions

Any policy-related actions that you can perform.

Replication policy settings

You configure replication policies to run according to replication policy settings.

Policy name

The name of the policy.

Description

Describes the policy. For example, the description might explain the purpose or function of the policy.

Enabled

Determines whether the policy is enabled.

Action

Determines the how the policy replicates data. All policies copy files from the source directory to the target directory and update files in the target directory to match files on the source directory. The action determines how deleting a file on the source directory affects the target. The following values are valid:

Copy

If a file is deleted in the source directory, the file is not deleted in the target directory.

Synchronize

Deletes files in the target directory if they are no longer present on the source. This ensures that an exact replica of the source directory is maintained on the target cluster.

Run job

Determines whether jobs are run automatically according to a schedule or only when manually specified by a user.

Last Successful Run

Displays the last time that a replication job for the policy completed successfully.

Last Started

Displays the last time that the policy was run.

Source Root Directory

The full path of the source directory. Data is replicated from the source directory to the target directory.

Included Directories

Determines which directories are included in replication. If one or more directories are specified by this setting, any directories that are not specified are not replicated.

Excluded Directories

Determines which directories are excluded from replication. Any directories specified by this setting are not replicated.

File Matching Criteria

Determines which files are excluded from replication. Any files that do not meet the specified criteria are not replicated.

Restrict Source Nodes

Determines whether the policy can run on all nodes on the source cluster or run only on specific nodes.

Target Host

The IP address or fully qualified domain name of the target cluster.

Target Directory

The full path of the target directory. Data is replicated to the target directory from the source directory.

Restrict Target Nodes

Determines whether the policy can connect to all nodes on the target cluster or can connect only to specific nodes.

Capture Snapshots

Determines whether archival snapshots are generated on the target cluster.

Snapshot Alias Name

Specifies an alias for the latest archival snapshot taken on the target cluster.

Snapshot Naming Pattern

Specifies how archival snapshots are named on the target cluster.

Snapshot Expiration

Specifies how long archival snapshots are retained on the target cluster before they are automatically deleted by the system.

Workers Threads Per Node

Specifies the number of workers per node that are generated by OneFS to perform each replication job for the policy.

Log Level

Specifies the amount of information that is recorded for replication jobs. More verbose options include all information from less verbose options. The following list describes the log levels from least to most verbose:

Notice

Includes job and process-level activity, including job starts, stops, and worker coordination information. This is the recommended log level.

Error

Includes events related to specific types of failures.

Network Activity

Includes more job-level activity and work-item information, including specific paths and snapshot names.

File Activity

Includes a separate event for each action taken on a file. Do not select this option without first consulting Isilon Technical Support.

Replication logs are typically used for debugging purposes. If necessary, you can log in to a node through the command-line interface and view the contents of the `/var/log/isi_migrate.log` file on the node.

Validate File Integrity

Determines whether OneFS performs a checksum on each file data packet that is affected by a replication job. If a checksum value does not match, OneFS retransmits the affected file data packet.

Keep Reports For

Specifies how long replication reports are kept before they are automatically deleted by OneFS.

Log Deletions on Synchronization

Determines whether OneFS records when a synchronization job deletes files or directories on the target cluster.

The following replication policy fields are available only through the OneFS command-line interface.

Source Subnet

Specifies whether replication jobs connect to any nodes in the cluster or if jobs can connect only to nodes in a specified subnet.

Source Pool

Specifies whether replication jobs connect to any nodes in the cluster or if jobs can connect only to nodes in a specified pool.

Password Set

Specifies a password to access the target cluster.

Report Max Count

Specifies the maximum number of replication reports that are retained for this policy.

Target Compare Initial Sync

Determines whether full or differential replications are performed for this policy. Full or differential replications are performed the first time a policy is run and after a policy is reset.

Source Snapshot Archive

Determines whether snapshots generated for the replication policy on the source cluster are deleted when the next replication policy is run. Enabling archival source snapshots does not require you to activate the SnapshotIQ license on the cluster.

Source Snapshot Pattern

If snapshots generated for the replication policy on the source cluster are retained, renames snapshots according to the specified rename pattern.

Source Snapshot Expiration

If snapshots generated for the replication policy on the source cluster are retained, specifies an expiration period for the snapshots.

Restrict Target Network

Determines whether replication jobs connect only to nodes in a given SmartConnect zone. This setting applies only if the Target Host is specified as a SmartConnect zone.

Target Detect Modifications

Determines whether SyncIQ checks the target directory for modifications before replicating files. By default, SyncIQ always checks for modifications.

Note

Disabling this option could result in data loss. It is recommended that you consult Isilon Technical Support before disabling this option.

Resolve

Determines whether you can manually resolve the policy if a replication job encounters an error.

Managing replication to the local cluster

You can interrupt replication jobs that target the local cluster.

You can cancel a currently running job that targets the local cluster, or you can break the association between a policy and its specified target. Breaking a source and target cluster association causes SyncIQ to perform a full replication the next time the policy is run.

Cancel replication to the local cluster

You can cancel a replication job that is targeting the local cluster.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Local Targets**.
2. In the **SyncIQ Local Targets** table, specify whether to cancel a specific replication job or all replication jobs targeting the local cluster.
 - To cancel a specific job, in the row for a replication job, select **Cancel Running Job**.
 - To cancel all jobs targeting the local cluster, select the check box to the left of **Policy Name** and then select **Cancel Selection** from the **Select a bulk action** list.

Break local target association

You can break the association between a replication policy and the local cluster. Breaking the target association will allow writes to the target directory but will also require you to reset the replication policy before you can run the policy again.

CAUTION

After a replication policy is reset, SyncIQ performs a full or differential replication the next time the policy is run. Depending on the amount of data being replicated, a full or differential replication can take a very long time to complete.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Local Targets**.
2. In the **SyncIQ Local Targets** table, in the row for a replication policy, select **Break Association**.
3. In the **Confirm** dialog box, click **Yes**.

View replication policies targeting the local cluster

You can view information about replication policies that are currently replicating data to the local cluster.

Procedure

1. Click **Data Protection** > **SyncIQ** > **Local Targets**.
2. In the **SyncIQ Local Targets** table, view information about replication policies.

Remote replication policy information

You can view information about replication policies that are currently targeting the local cluster.

The following information is displayed in the **SyncIQ Local Targets** table:

ID

The ID of the replication policy.

Policy Name

The name of the replication policy.

Source Host

The name of the source cluster.

Source Cluster GUID

The GUID of the source cluster.

Coordinator IP

The IP address of the node on the source cluster that is acting as the job coordinator.

Updated

The time when data about the policy or job was last collected from the source cluster.

Target Path

The path of the target directory on the target cluster.

Status

The current status of the replication job.

Actions

Displays any job-related actions that you can perform.

CHAPTER 5

Recovering data with SyncIQ

This section contains the following topics:

- ◆ [Initiating data failover and failback with SyncIQ](#)..... 44
- ◆ [Performing disaster recovery for SmartLock directories](#).....46

Initiating data failover and failback with SyncIQ

You can fail over from one Isilon cluster to another if, for example, a cluster becomes unavailable. You can then fail back to a primary cluster if the primary cluster becomes available again. You can revert failover if you decide that the failover was unnecessary, or if you failed over for testing purposes.

If you fail over a scheduled replication policy on the secondary cluster, and the corresponding policy on the primary cluster runs a replication job, the job might fail and the policy might be set to an unrunnable state. To resolve this, modify the replication policy so that it is set to run only manually, resolve the policy, and complete the failback process. After you complete the failback process, you can modify the policy to run according to a schedule again.

Note

Although you cannot fail over or fail back SmartLock directories, you can recover SmartLock directories on a target cluster. After you recover SmartLock directories, you can migrate them back to the source cluster.

Fail over data to a secondary cluster

You can fail over to a secondary Isilon cluster if, for example, a cluster becomes unavailable.

Before you begin

Create and successfully run a replication policy.

Complete the following procedure for each replication policy that you want to fail over.

Procedure

1. On the secondary Isilon cluster, click **Data Protection > SyncIQ > Local Targets**.
2. In the **SyncIQ Local Targets** table, in the row for a replication policy, from the **Actions** column, select **Allow Writes**.

After you finish

Direct clients to begin accessing the secondary cluster.

Revert a failover operation

Failover reversion undoes a failover operation on a secondary cluster, enabling you to replicate data from the primary cluster to the secondary cluster again. Failover reversion is useful if the primary cluster becomes available before data is modified on the secondary cluster or if you failed over to a secondary cluster for testing purposes.

Before you begin

Fail over a replication policy.

Reverting a failover operation does not migrate modified data back to the primary cluster. To migrate data that clients have modified on the secondary cluster, you must fail back to the primary cluster.

Note

Failover reversion is not supported for SmartLock directories.

Complete the following procedure for each replication policy that you want to fail over.

Procedure

1. Run the `isi sync recovery allow-write` command with the `--revert` option.

For example, the following command reverts a failover operation for `newPolicy`:

```
isi sync recovery allow-write newPolicy --revert
```

Fail back data to a primary cluster

After you fail over to a secondary cluster, you can fail back to the primary cluster.

Before you begin

Fail over a replication policy.

Procedure

1. On the primary cluster, click **Data Protection > SyncIQ > Policies**.
2. In the **SyncIQ Policies** table, in the row for a replication policy, from the **Actions** column, select **Resync-prep**.

SyncIQ creates a mirror policy for each replication policy on the secondary cluster.

SyncIQ names mirror policies according to the following pattern:

```
<replication-policy-name>_mirror
```

3. On the secondary cluster, replicate data to the primary cluster by using the mirror policies.

You can replicate data either by manually starting the mirror policies or by modifying the mirror policies and specifying a schedule.
4. Prevent clients from accessing the secondary cluster and then run each mirror policy again.

To minimize impact to clients, it is recommended that you wait until client access is low before preventing client access to the cluster.
5. On the primary cluster, click **Data Protection > SyncIQ > Local Targets**.
6. In the **SyncIQ Local Targets** table, from the **Actions** column, select **Allow Writes** for each mirror policy.
7. On the secondary cluster, click **Data Protection > SyncIQ > Policies**.
8. In the **SyncIQ Policies** table, from the **Actions** column, select **Resync-prep** for each mirror policy.

After you finish

Redirect clients to begin accessing the primary cluster.

Performing disaster recovery for SmartLock directories

Although you cannot fail over or fail back SmartLock directories, you can recover SmartLock directories on a target cluster. After you recover SmartLock directories, you can migrate them back to the source cluster.

Recover SmartLock directories on a target cluster

You can recover SmartLock directories that you have replicated to a target cluster.

Before you begin

Create and successfully run a replication policy.

Complete the following procedure for each SmartLock directory that you want to recover.

Procedure

1. On the target cluster, click **Data Protection > SyncIQ > Local Targets**.
2. In the **SyncIQ Local Targets** table, in the row of the replication policy, enable writes to the target directory of the policy.
 - If the last replication job completed successfully and a replication job is not currently running, select **Allow Writes**.
 - If a replication job is currently running, wait until the replication job completes, and then select **Allow Writes**.
 - If the primary cluster became unavailable while a replication job was running, select **Break Association**.
3. If you clicked **Break Association**, restore any files that are left in an inconsistent state.
 - a. Delete all files that are not committed to a WORM state from the target directory.
 - b. Copy all files from the failover snapshot to the target directory.

Failover snapshots are named according to the following naming pattern:

```
SIQ-Failover-<policy-name>-<year>-<month>-<day>_<hour>-<minute>-<second>
```

Snapshots are stored in the `/ifs/.snapshot` directory.
4. If any SmartLock directory configuration settings, such as an autocommit time period, were specified for the source directory of the replication policy, apply those settings to the target directory.

Because autocommit information is not transferred to the target cluster, files that were scheduled to be committed to a WORM state on the source cluster will not be scheduled to be committed at the same time on the target cluster. To ensure that all files are retained for the appropriate time period, you can commit all files in target SmartLock directories to a WORM state. For example, the following command automatically commits all files in `/ifs/data/smartlock` to a WORM state after one minute.

```
isi smartlock modify --path /ifs/data/smartlock --autocommit 1m
```

After you finish

Redirect clients to begin accessing the target cluster.

Migrate SmartLock directories

You might want to migrate SmartLock directories if you restored the directories on a target cluster, and want to transfer those directories either back to the source cluster or to a new cluster.

Procedure

1. On a cluster, create a replication policy for each policy that you want to migrate.

The policies must meet the following requirements:

- The source directory is the SmartLock directory that you are migrating.
- The target directory is an empty SmartLock directory. The source and target directories must be of the same SmartLock type. For example, if the target directory is a compliance directory, the source must also be a compliance directory.

2. Replicate data to the target cluster by running the policies you created.

You can replicate data either by manually starting the policies or by specifying a policy schedule.

3. Optional: To ensure that SmartLock protection is enforced for all files, commit all files in the SmartLock source directory to a WORM state.

Because autocommit information is not transferred to the target cluster, files that were scheduled to be committed to a WORM state on the source cluster will not be scheduled to be committed at the same time on the target cluster. To ensure that all files are retained for the appropriate time period, you can commit all files in target SmartLock directories to a WORM state.

For example, the following command automatically commits all files in `/ifs/data/smartlock` to a WORM state after one minute.

```
isi smartlock modify --path /ifs/data/smartlock --autocommit 1m
```

This step is unnecessary if you have not configured an autocommit time period for the SmartLock directory being replicated.

4. Prevent clients from accessing the source cluster and run the policy that you created.

To minimize impact to clients, it is recommended that you wait until client access is low before preventing client access to the cluster.

5. On the target cluster, click **Data Protection** > **SyncIQ** > **Local Targets**.

6. In the **SyncIQ Local Targets** table, in the row of each replication policy, from the **Actions** column, select **Allow Writes**.

7. Optional: If any SmartLock directory configuration settings, such as an autocommit time period, were specified for the source directories of the replication policies, apply those settings to the target directories.

8. Optional: Delete the copy of your SmartLock data on the source cluster.

If the SmartLock directories are compliance directories or enterprise directories with the privileged delete functionality permanently disabled, you cannot recover the space consumed by the source SmartLock directories until all files are released from a WORM state. If you want to free the space before files are released from a WORM state, contact Isilon Technical Support for information about reformatting your cluster.

CHAPTER 6

NDMP backup

This section contains the following topics:

- ◆ [NDMP two way backup](#) 50
- ◆ [Snapshot-based incremental backups](#) 50
- ◆ [NDMP protocol support](#) 51
- ◆ [Supported DMAs](#) 51
- ◆ [NDMP hardware support](#) 52
- ◆ [NDMP backup limitations](#) 52
- ◆ [NDMP performance recommendations](#) 52
- ◆ [Excluding files and directories from NDMP backups](#) 54

NDMP two way backup

To perform NDMP two-way backups, you must attach a Backup Accelerator node to your Isilon cluster and attach a tape device to the Backup Accelerator node. You must then use OneFS to detect the tape device before you can back up to that device.

You can connect supported tape devices directly to the Fibre Channel ports of a Backup Accelerator node. Alternatively, you can connect Fibre Channel switches to the Fibre Channel ports on the Backup Accelerator node, and connect tape and media changer devices to the Fibre Channel switches. For more information, see your Fibre Channel switch documentation about zoning the switch to allow communication between the Backup Accelerator node and the connected tape and media changer devices.

If you attach tape devices to a Backup Accelerator node, the cluster detects the devices when you start or restart the node or when you re-scan the Fibre Channel ports to discover devices. If a cluster detects tape devices, the cluster creates an entry for the path to each detected device.

If you connect a device through a Fibre Channel switch, multiple paths can exist for a single device. For example, if you connect a tape device to a Fibre Channel switch, and then connect the Fibre Channel switch to two Fibre Channel ports, OneFS creates two entries for the device, one for each path.

Note

If you perform an NDMP two-way backup operation, you must assign static IP addresses to the Backup Accelerator node. If you connect to the cluster through a data management application (DMA), you must connect to the IP address of a Backup Accelerator node. If you perform an NDMP three-way backup, you can connect to any node in the cluster.

Snapshot-based incremental backups

You can implement snapshot-based incremental backups to increase the speed at which these backups are performed.

During a snapshot-based incremental backup, OneFS checks the snapshot taken for the previous NDMP backup operation and compares it to a new snapshot. OneFS then backs up all data that was modified since the last snapshot was made.

If the incremental backup does not involve snapshots, OneFS must scan the directory to discover which files were modified. OneFS can perform incremental backups significantly faster if snapshots are referenced.

You can perform incremental backups without activating a SnapshotIQ license on the cluster. Although SnapshotIQ offers a number of useful features, it does not enhance snapshot capabilities in NDMP backup and recovery.

If you implement snapshot-based incremental backups, OneFS retains each snapshot taken for NDMP backups until a new backup of the same or lower level is performed. However, if you do not implement snapshot-based incremental backups, OneFS automatically deletes each snapshot generated after the corresponding backup is completed or canceled.

The following table lists whether supported data management applications (DMAs) can perform snapshot-based incremental backups:

Table 2 DMA support for snapshot-based incremental backups

DMA	Supported
Symantec NetBackup	Yes
EMC Networker	Yes
EMC Avamar	No
Commvault Simpana	No
IBM Tivoli Storage Manager	No
Symantec Backup Exec	Yes
Dell NetVault	Yes
ASG-Time Navigator	Yes

NDMP protocol support

You can back up cluster data through version 3 or 4 of the NDMP protocol.

OneFS supports the following features of NDMP versions 3 and 4:

- ◆ Full (level 0) NDMP backups
- ◆ Incremental (levels 1-10) NDMP backups

Note

In a level 10 NDMP backup, only data changed since the most recent incremental (level 1-9) backup or the last level 10 backup is copied. By repeating level 10 backups, you can be assured that the latest versions of files in your data set are backed up without having to run a full backup.

- ◆ Token-based NDMP backups
- ◆ NDMP TAR backup type
- ◆ Path-based and dir/node file history format
- ◆ Direct Access Restore (DAR)
- ◆ Directory DAR (DDAR)
- ◆ Including and excluding specific files and directories from backup
- ◆ Backup of file attributes
- ◆ Backup of Access Control Lists (ACLs)
- ◆ Backup of Alternate Data Streams (ADSs)
- ◆ Backup Restartable Extension (BRE)

OneFS supports connecting to clusters through IPv4 or IPv6.

Supported DMAs

NDMP backups are coordinated by a data management application (DMA) that runs on a backup server.

OneFS supports the following DMAs:

- ◆ Symantec NetBackup
- ◆ EMC NetWorker
- ◆ EMC Avamar
- ◆ Symantec Backup Exec
- ◆ IBM Tivoli Storage Manager
- ◆ Dell NetVault
- ◆ CommVault Simpana (IPv6 protocol only)
- ◆ ASG-Time Navigator

NDMP hardware support

OneFS can backup data to and restore data from tape devices and virtual tape libraries (VTL).

OneFS supports the following types of emulated and physical tape devices:

- ◆ LTO-3
- ◆ LTO-4
- ◆ LTO-5

OneFS supports the following virtual tape libraries (VTLs):

- ◆ FalconStor VTL 5.20
- ◆ Data Domain VTL 5.1.04 or later

NDMP backup limitations

OneFS NDMP backups have the following limitations:

- ◆ OneFS does not back up file system configuration data, such as file protection level policies and quotas.
- ◆ OneFS does not support multiple concurrent backups onto the same tape.
- ◆ OneFS does not support restoring data from a file system other than OneFS. However, you can migrate data via the NDMP protocol from a NetApp or EMC VNX storage system to OneFS.
- ◆ Backup Accelerator nodes cannot interact with more than 1024 device paths, including the paths of tape and media changer devices. For example, if each device has four paths, you can connect 256 devices to a Backup Accelerator node. If each device has two paths, you can connect 512 devices.
- ◆ OneFS does not support more than 64 concurrent NDMP sessions per Backup Accelerator node.

NDMP performance recommendations

Consider the following recommendations to optimize OneFS NDMP backups.

General performance recommendations

- ◆ Install the latest patches for OneFS and your data management application (DMA).
- ◆ If you are backing up multiple directories that contain small files, set up a separate schedule for each directory.

- ◆ If you are performing three-way NDMP backups, run multiple NDMP sessions on multiple nodes in your Isilon cluster.
- ◆ Restore files through Direct Access Restore (DAR) and Directory DAR (DDAR). This is especially recommended if you restore files frequently. However, it is recommended that you do not use DAR to restore a full backup or a large number of files, as DAR is better suited to restoring smaller numbers of files.
- ◆ Use the largest tape record size available for your version of OneFS. The largest tape record size for OneFS versions 6.5.5 and later is 256 KB. The largest tape record size for versions of OneFS earlier than 6.5.5 is 128 KB.
- ◆ If possible, do not include or exclude files from backup. Including or excluding files can affect backup performance, due to filtering overhead.
- ◆ Limit the depth of nested subdirectories in your file system.
- ◆ Limit the number of files in a directory. Distribute files across multiple directories instead of including a large number of files in a single directory.

Networking recommendations

- ◆ Assign static IP addresses to Backup Accelerator nodes.
- ◆ Configure SmartConnect zones to specify pools of IP address ranges that are exclusive to NDMP backup operations.
- ◆ Connect NDMP sessions only through SmartConnect zones that are exclusively used for NDMP backup.
- ◆ Configure multiple policies when scheduling backup operations, with each policy capturing a portion of the file system. Do not attempt to back up the entire file system through a single policy.

Backup Accelerator recommendations

- ◆ Run a maximum of four concurrent streams per Backup Accelerator node.

Note

This is recommended only if you are backing up a significant amount of data. Running four concurrent streams might not be necessary for smaller backups.

- ◆ Attach more Backup Accelerator nodes to larger clusters. The recommended number of Backup Accelerator nodes is listed in the following table.

Table 3 Nodes per Backup Accelerator node

Node type	Recommended number of nodes per Backup Accelerator node
X-Series	3
NL-Series	3
S-Series	3

- ◆ Attach more Backup Accelerator nodes if you are backing up to more tape devices. The following table lists the recommended number of tape devices per backup accelerator node:

Table 4 Tape devices per Backup Accelerator node

Tape device type	Recommended number of tape devices per Backup Accelerator node
LTO-5	3
LTO-4	4
LTO-3	8

DMA-specific recommendations

- ◆ Apply path-based file history instead of directory/inode (dir/node) file history.
- ◆ Turn on multi-streaming, which enables OneFS to back up data to multiple tape devices at the same time.

Excluding files and directories from NDMP backups

You can exclude files and directories from NDMP backup operations by specifying NDMP environment variables through a data management application (DMA). If you include a file or directory, all other files and directories are automatically excluded from backup operations. If you exclude a file or directory, all files and directories except the excluded one are backed up.

You can include or exclude files and directories by specifying the following character patterns:

Table 5 NDMP file and directory matching wildcards

Character	Description	Example	Includes or excludes the following directories
*	Takes the place of any character or characters	archive*	/ifs/data/archive1 /ifs/data/archive42_a/media
[]	Takes the place of a range of letters or numbers	data_store_[a-f] data_store_[0-9]	/ifs/data/data_store_a /ifs/data/data_store_c /ifs/data/data_store_8
?	Takes the place of any single character	user_?	/ifs/data/user_1 /ifs/data/user_2
\	Includes a blank space	user\ 1	/ifs/data/user 1

Unanchored patterns such as `home` or `user1` target a string of text that might belong to many files or directories. Anchored patterns target specific file pathnames, such as `ifs/data/home`. You can include or exclude either type of pattern.

For example, suppose you want to back up the `/ifs/data/home` directory, which contains the following files and directories:

- ◆ `/ifs/data/home/user1/file.txt`
- ◆ `/ifs/data/home/user2/user1/file.txt`

- ◆ `/ifs/data/home/user3/other/file.txt`
- ◆ `/ifs/data/home/user4/emptydirectory`

If you simply include the `/ifs/data/home` directory, all files and directories, including `emptydirectory` would be backed up.

If you specify both include and exclude patterns, any excluded files or directories under the included directories would not be backed up. If the excluded directories are not found in any of the included directories, the exclude specification would have no effect.

Note

Specifying unanchored patterns can degrade the performance of backups. It is recommended that you avoid unanchored patterns whenever possible.

CHAPTER 7

Backing up and recovering data with NDMP

This section contains the following topics:

◆ NDMP backup and recovery tasks.....	58
◆ Configuring basic NDMP backup settings.....	58
◆ Managing NDMP user accounts.....	59
◆ Managing NDMP backup devices.....	60
◆ Managing NDMP backup ports.....	62
◆ Managing NDMP backup sessions.....	63
◆ Managing restartable backups.....	65
◆ Sharing tape drives between clusters.....	67
◆ Managing default NDMP settings.....	67
◆ Managing snapshot based incremental backups.....	71
◆ View NDMP backup logs.....	72
◆ Configuring NDMP backups with EMC NetWorker.....	72
◆ Configuring NDMP backup with Symantec NetBackup.....	76
◆ Configuring NDMP backup with CommVault Simpana.....	80

NDMP backup and recovery tasks

Before you can back up data with NDMP, you must configure and enable NDMP backup on the cluster. After this, you can configure settings for NDMP backup ports and backup devices. After you start backing up data with NDMP, you can monitor backup sessions.

Configuring basic NDMP backup settings

You can configure NDMP backup settings to control how these backups are performed for the cluster. You can also configure OneFS to interact with a specific data management application (DMA) for NDMP backups.

NDMP backup settings

You can configure the following settings to control how NDMP backups are performed on the cluster.

Port number

The number of the port through which the data management application (DMA) can connect to the cluster.

DMA vendor

The DMA vendor that the cluster is configured to interact with.

View NDMP backup settings

You can view current NDMP backup settings. These settings define whether NDMP backup is enabled, the port through which your data management application (DMA) connects to the cluster, and the DMA vendor that OneFS is configured to interact with.

Procedure

1. Click **Data Protection** > **Backup** > **NDMP Settings** and view NDMP backup settings.
2. In the **Settings** area, review NDMP backup settings.

Configure and enable NDMP backup

OneFS prevents NDMP backups by default. Before you can perform NDMP backups, you must enable NDMP backups and configure NDMP settings.

Procedure

1. Click **Data Protection** > **Backup** > **NDMP Settings**.
2. In the **Service** area, click **Enable**.
3. Optional: To specify a port through which data management applications (DMAs) access the cluster, or the DMA vendor that OneFS is to interact with, in the **Settings** area, click **Edit settings**.
 - In the **Port number** field, type a port number.
 - From the **DMA vendor** list, select the name of the DMA vendor to manage backup operations.

If your DMA vendor is not included in the list, select **generic**. However, note that any vendors not included on the list are not officially supported and might not function as expected.

4. Click **Add administrator** to add an NDMP user account through which your DMA can access the cluster.
 - a. In the **Add Administrator** dialog box, in the **Name** field, type a name for the account.
 - b. In the **Password** and **Confirm password** fields, type a password for the account.
 - c. Click **Submit**.

Disable NDMP backup

You can disable NDMP backup if you no longer want to use this backup method.

Procedure

1. Click **Data Protection** > **Backup** > **NDMP Settings**.
2. In the **Service** area, click **Disable**.

Managing NDMP user accounts

You can create, delete, and modify the passwords of NDMP user accounts.

Create an NDMP user account

Before you can perform NDMP backups, you must create an NDMP user account through which your data management application (DMA) can access the Isilon cluster.

Procedure

1. Click **Data Protection** > **Backup** > **NDMP Settings**.
2. In the **NDMP Administrators** area, click **Add administrator**.
3. In the **Add Administrator** dialog box, in the **Name** field, type a name for the account.
4. In the **Password** and **Confirm password** fields, type a password for the account.
5. Click **Submit**.

View NDMP user accounts

You can view information about NDMP user accounts.

Procedure

1. Click **Data Protection** > **Backup** > **NDMP Settings**.
2. In the **NDMP administrators** area, review information about NDMP user accounts.

Modify the password of an NDMP user account

You can modify the password for an NDMP user account.

Procedure

1. Click **Data Protection** > **Backup** > **NDMP Settings**.
2. In the **NDMP Administrator** table, in the row for an NDMP user account, click **Change password**.
3. In the **Password** and **Confirm password** fields, type a new password for the account.
4. Click **Submit**.

Delete an NDMP user account

You can delete an NDMP user account.

Procedure

1. Click **Data Protection > Backup > NDMP Settings**.
2. In the **NDMP Administrators** table, in the row for an NDMP user account, click **Delete**.
3. In the **Confirm** dialog box, click **Yes**.

Managing NDMP backup devices

After you attach a tape or media changer device to a Backup Accelerator node, you must configure OneFS to detect and establish a connection to the device. After the connection between the cluster and the backup device is established, you can modify the name that the cluster has assigned to the device, or disconnect the device from the cluster.

NDMP backup device settings

OneFS creates a device entry for each device you attach to the cluster through a Backup Accelerator node.

The following table describes the settings available in the **Tape Devices** and **Media Changers** tables:

Table 6 NDMP backup device settings

Setting	Description
Name	A device name assigned by OneFS.
State	Indicates whether the device is in use. If data is currently being backed up to or restored from the device, <i>Read/Write</i> appears. If the device is not in use, <i>Closed</i> appears.
WWN	The world wide node name (WWNN) of the device.
Product	The name of the device vendor and the model name or number of the device.
Serial Number	The serial number of the device.
Paths	The name of the Backup Accelerator node that the device is attached to and the port number or numbers to which the device is connected.
LUN	The logical unit number (LUN) of the device.
Port ID	The port ID of the device that binds the logical device to the physical device.
WWPN	The world wide port name (WWPN) of the port on the tape or media changer device.

Detect NDMP backup devices

If you connect a tape device or media changer to a Backup Accelerator node, you must configure OneFS to detect the device. Only then can OneFS back up data to and restore

data from the device. In OneFS, you can scan a specific node, a specific port, or all ports on all nodes.

Procedure

1. Click **Data Protection** > **Backup** > **Devices**.
2. Click **Discover devices**.
3. Optional: To scan only a specific node for NDMP devices, from the **Nodes** list, select a node.
4. Optional: To scan only a specific port for NDMP devices, from the **Ports** list, select a port.

If you specify a port and a node, only the specified port on the node is scanned. However, if you specify only a port, the specified port will be scanned on all nodes.

5. Optional: To remove entries for devices or paths that have become inaccessible, select the **Delete inaccessible paths or devices** check box.
6. Click **Submit**.

Results

For each device that is detected, an entry is added to either the **Tape Devices** or **Media Changers** tables.

View NDMP backup devices

You can view information about tape and media changer devices that are currently attached to your Isilon cluster.

Procedure

1. Click **Data Protection** > **Backup** > **Devices**.
2. In the **Tape Devices** and **Media Changers** tables, review information about NDMP backup devices.

Modify the name of an NDMP backup device

You can modify the name of an NDMP backup device in OneFS.

Procedure

1. Click **Data Protection** > **Backup** > **Devices**.
2. In the **Tape Devices** table, click the name of a backup device entry.
3. In the **Rename Device** dialog box, in the **Device Name** field, type a new name for the backup device.
4. Click **Submit**.

Delete an entry for an NDMP backup device

If you physically remove an NDMP device from a cluster, OneFS retains the entry for the device. You can delete a device entry for a removed device. You can also remove the device entry for a device that is still physically attached to the cluster; this causes OneFS to disconnect from the device.

If you remove a device entry for a device that is connected to the cluster, and you do not physically disconnect the device, OneFS will detect the device the next time it scans the ports. You cannot remove a device entry for a device that is currently in use.

Procedure

1. Click **Data Protection > Backup > Devices**.
2. In the **Tape Devices** table, in the row for the target device, click **Delete device**.
3. In the **Confirm** dialog box, click **Yes**.

Managing NDMP backup ports

You can manage the Fibre Channel ports that connect tape and media changer devices to a Backup Accelerator node. You can also enable, disable, or modify the settings of an NDMP backup port.

NDMP backup port settings

OneFS assigns default settings to each port on each Backup Accelerator node attached to the cluster. These settings identify each port and specify how the port interacts with NDMP backup devices.

The settings that appear in the **Ports** table are as follows:

Table 7 NDMP backup port settings

Setting	Description
Port	The name of the Backup Accelerator node, and the number of the port.
Topology	<p>The type of Fibre Channel topology that the port is configured to support.. Options are:</p> <p>Point to Point A single backup device or Fibre Channel switch directly connected to the port.</p> <p>Loop Multiple backup devices connected to a single port in a circular formation.</p> <p>Auto Automatically detects the topology of the connected device. This is the recommended setting, and is required if you are using a switched-fabric topology.</p>
WWNN	The world wide node name (WWNN) of the port. This name is the same for each port on a given node.
WWPN	The world wide port name (WWPN) of the port. This name is unique to the port.
Rate	The rate at which data is sent through the port. The rate can be set to 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s, and Auto. 8 Gb/s is available for A100 nodes only. If set to Auto, OneFS automatically negotiates with the DMA to determine the rate. Auto is the recommended setting.

View NDMP backup ports

You can view information about Fibre Channel ports of Backup Accelerator nodes attached to a cluster.

Procedure

1. Click **Data Protection** > **Backup** > **Ports**.
2. In the **Ports** table, review information about NDMP backup ports.

Modify NDMP backup port settings

You can modify the settings of an NDMP backup port.

Procedure

1. Click **Data Protection** > **Backup** > **Ports**.
2. In the **Sessions** table, click the name of a port.
3. In the **Edit Port** dialog box, modify port settings as needed, and then click **Submit**.

Enable or disable an NDMP backup port

You can enable or disable an NDMP backup port.

Procedure

1. Click **Data Protection** > **Backup** > **Ports**.
2. In the **Ports** table, in the row of a port, click **Enable** or **Disable**.

Managing NDMP backup sessions

You can view the status of NDMP backup sessions or terminate a session that is in progress.

NDMP session information

You can view information about active NDMP sessions.

The following information is included in the **Sessions** table, as follows:

Table 8 NDMP session information

Item	Description
Session	The unique identification number that OneFS assigned to the session.
Elapsed	How much time has elapsed since the session started.
Transferred	The amount of data that was transferred during the session.
Throughput	The average throughput of the session over the past five minutes.
Client/Remote	The IP address of the backup server that the data management application (DMA) is running on. If a three-way NDMP backup or restore operation is currently running, the IP address of the remote tape media server also appears.

Table 8 NDMP session information (continued)

Item	Description
Mover/Data	<p>The current state of the data mover and the data server. The first word describes the activity of the data mover. The second word describes the activity of the data server.</p> <p>The data mover and data server send data to and receive data from each other during backup and restore operations. The data mover is a component of the backup server that receives data during backups and sends data during restore operations. The data server is a component of OneFS that sends data during backups and receives information during restore operations.</p> <p>The following states might appear:</p> <p>Active The data mover or data server is currently sending or receiving data.</p> <p>Paused The data mover is temporarily unable to receive data. While the data mover is paused, the data server cannot send data to the data mover. The data server cannot be paused.</p> <p>Idle The data mover or data server is not sending or receiving data.</p> <p>Listen The data mover or data server is waiting to connect to the data server or data mover.</p>
Operation	<p>The type of operation (backup or restore) that is currently in progress. If no operation is in progress, this field is blank.</p> <p>Backup (0-10) Indicates that data is currently being backed up to a media server. The number indicates the level of NDMP backup.</p> <p>Restore Indicates that data is currently being restored from a media server.</p>
Source/Destination	<p>If an operation is currently in progress, specifies the <code>/ifs</code> directories that are affected by the operation. If a backup is in progress, displays the path of the source directory that is being backed up. If a restore operation is in progress, displays the path of the directory that is being restored along with the destination directory to which the tape media server is restoring data. If you are restoring data to the same location that you backed up your data from, the same path appears twice.</p>
Device	<p>The name of the tape or media changer device that is communicating with the cluster.</p>
Mode	<p>How OneFS is interacting with data on the backup media server, as follows:</p> <p>Read/Write OneFS is reading and writing data during a backup operation.</p>

Table 8 NDMP session information (continued)

Item	Description
	<p>Read OneFS is reading data during a restore operation.</p> <p>Raw The DMA has access to tape drives, but the drives do not contain writable tape media.</p>

View NDMP sessions

You can view information about active NDMP sessions.

Procedure

1. Click **Data Protection > Backup > Sessions**.
2. In the **Sessions** table, review information about NDMP sessions.

End an NDMP session

You can end an NDMP backup or restore session at any time.

Procedure

1. Click **Data Protection > Backup > Sessions**.
2. In the **Sessions** table, in the row of the NDMP session that you want to end, click **Kill**.
3. In the **Confirm** dialog box, click **Yes**.

Managing restartable backups

A restartable backup is a type of NDMP backup that you can enable in your data management application (DMA). If a restartable backup fails, for example, because of a power outage, you can restart the backup from a checkpoint close to the point of failure. In contrast, when a non-restartable backup fails, you must back up all data from the beginning, regardless of what was transferred during the initial backup process.

After you enable restartable backups from your DMA, you can manage restartable backup contexts from OneFS. These contexts are the data that OneFS stores to facilitate restartable backups. Each context represents a checkpoint that the restartable backup process can return to if a backup fails.

Restartable backups are supported only for EMC NetWorker 8.1 and later.

Configure restartable backups

You must configure EMC NetWorker to enable restartable backups and, optionally, define the checkpoint interval.

If you do not specify a checkpoint interval, NetWorker uses the default interval of 5 GB.

Procedure

1. Configure the client and the directory path that you want to back up as you would normally.

2. In the **Client Properties** dialog box, enable restartable backups.
 - a. On the **General** screen, click the **Checkpoint enabled** checkbox.
 - b. Specify `File` in the **Checkpoint granularity** drop-down list.
3. In the **Application information** field, type any NDMP variables that you want to specify.
The following specifies a checkpoint interval of 1 GB.
`CHECKPOINT_INTERVAL_IN_BYTES=1GB`
4. Finish configuration and click **OK** in the **Client Properties** dialog box.
5. Start the backup.
6. If the backup is interrupted, for example, because of a power failure, restart it.
 - a. Browse to the **Monitoring** screen, and locate the backup process in the **Groups** list.
 - b. Right-click on the backup process, and in the context menu, click **Restart**.
NetWorker automatically restarts the backup from the last checkpoint.

View restartable backup contexts

You can view restartable backup contexts that have been configured.

Procedure

1. View all backup contexts by running the following command:
`isi ndmp extensions contexts list`
2. To view detailed information about a specific backup context, run the `isi ndmp extensions contexts view` command.

The following command displays detailed information about a backup context with an ID of `792eeb8a-8784-11e2-aa70-0025904e91a4`:

```
isi ndmp extensions contexts view 792eeb8a-8784-11e2-aa70-0025904e91a4
```

Delete a restartable backup context

After a restartable backup context is no longer needed, your data management application (DMA) automatically requests that OneFS delete the context. You can manually delete a restartable backup context before the DMA requests it.

Note

It is recommended that you do not manually delete restartable backup contexts. Manually deleting a restartable backup context requires you to restart the corresponding NDMP backup from the beginning.

Procedure

1. Run the `isi ndmp extensions contexts delete` command.

The following command deletes a restartable backup context with an ID of `792eeb8a-8784-11e2-aa70-0025904e91a4`:

```
isi ndmp extensions contexts delete 792eeb8a-8784-11e2-aa70-0025904e91a4
```

Configure restartable backup settings

You can specify the number of restartable backup contexts that OneFS retains at a time, up to a maximum of 1024 contexts.

Procedure

1. Run the `isi ndmp extensions settings modify` command.

The following command sets the maximum number of restartable backup contexts to 128:

```
isi ndmp extensions settings modify --bre_max_contexts 128
```

View restartable backup settings

You can view the current limit of restartable backup contexts that OneFS retains at one time.

Procedure

1. Run the following command:

```
isi ndmp extensions settings view
```

Sharing tape drives between clusters

Multiple Isilon clusters, or an Isilon cluster and a third-party NAS system, can be configured to share a single tape drive. This helps to maximize the use of the tape infrastructure in your data center.

In your data management application (DMA), you must configure NDMP to control the tape drive and ensure that it is shared properly. The following configurations are supported.

OneFS Versions*	Supported DMAs	Tested configurations
<ul style="list-style-type: none"> • 7.1.0.1 • 7.0.2.5 • 6.6.5.26 	<ul style="list-style-type: none"> • EMC NetWorker 8.0 and later • Symantec NetBackup 7.5 and later 	<ul style="list-style-type: none"> • Isilon Backup Accelerator with a second Backup Accelerator • Isilon Backup Accelerator with a NetApp storage system
* The tape drive sharing function is not supported in the OneFS 7.0.1 release.		

EMC NetWorker refers to the tape drive sharing capability as DDS (dynamic drive sharing). Symantec NetBackup uses the term SSO (shared storage option). Consult your DMA vendor documentation for configuration instructions.

Managing default NDMP settings

In OneFS, you can manage NDMP backup and restore operations by specifying default NDMP environment variables. You can also override default NDMP environment variables

through your data management application (DMA). For more information about specifying NDMP environment variables through your DMA, see your DMA documentation.

Set default NDMP settings for a directory

You can set default NDMP settings for a directory.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Set default NDMP settings by running the `isi ndmp settings variables create` command.

For example, the following command sets the default file history format to path-based format for `/ifs/data/media`:

```
isi ndmp settings variables create /ifs/data/media HIST F
```

Modify default NDMP settings for a directory

You can modify the default NDMP settings for a directory.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Modify default NDMP settings by running the `isi ndmp settings variables modify` command.

For example, the following command sets the default file history format to path-based format for `/ifs/data/media`:

```
isi ndmp settings variables modify /ifs/data/media HIST F
```

3. Optional: To remove a default NDMP setting for a directory, run the `isi ndmp settings variables delete` command:

For example, the following command removes the default file history format for `/ifs/data/media`:

```
isi ndmp settings variables delete /ifs/data/media --name HIST
```

View default NDMP settings for directories

You can view the default NDMP settings for directories.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. View default NDMP settings by running the following command:

```
isi ndmp settings variables list
```

NDMP environment variables

You can specify default settings of NDMP backup and restore operations through NDMP environment variables. You can also specify NDMP environment variables through your data management application (DMA).

Table 9 NDMP environment variables

Environment variable	Valid values	Default	Description
BACKUP_MODE=	TIMESTAMP SNAPSHOT	TIMESTAMP	Enables or disables snapshot-based incremental backups. To enable snapshot-based incremental backups, specify SNAPSHOT. To disable snapshot-based incremental backups, specify TIMESTAMP.
FILESYSTEM=	<i><file-path></i>	None	Specifies the full path of the directory you want to back up. Must be specified by the DMA before starting the backup, or an error is generated.
LEVEL=	<i><integer></i>	0	Specifies the level of NDMP backup to perform. The following values are valid: 0 Performs a full NDMP backup. 1 - 9 Performs an incremental backup at the specified level. 10 Performs unlimited incremental backups.
UPDATE=	Y N	Y	Determines whether OneFS updates the dump dates file. Y OneFS updates the dump dates file. N OneFS does not update the dump dates file.
HIST=	<i><file-history-format></i>	Y	Specifies the file history format. The following values are valid: D Specifies dir/node file history.

Table 9 NDMP environment variables (continued)

Environment variable	Valid values	Default	Description
			<p>F</p> <p>Specifies path-based file history.</p> <p>Y</p> <p>Specifies the default file history format determined by your NDMP backup settings.</p> <p>N</p> <p>Disables file history.</p>
DIRECT=	Y N	N	<p>Enables or disables Direct Access Restore (DAR) and Directory DAR (DDAR). The following values are valid:</p> <p>Y</p> <p>Enables DAR and DDAR.</p> <p>N</p> <p>Disables DAR and DDAR.</p>
FILES=	<i><file-matching-pattern></i>	None	<p>If you specify this option, OneFS backs up only files and directories that meet the specified pattern. Separate multiple patterns with a space.</p>
EXCLUDE=	<i><file-matching-pattern></i>	None	<p>If you specify this option, OneFS does not back up files and directories that meet the specified pattern. Separate multiple patterns with a space.</p>
RESTORE_HARDLINK_BY_TABLE=	Y N	N	<p>Determines whether OneFS recovers hard links by building a hard-link table during restore operations. Specify this option if hard links were incorrectly backed up, and restore operations are failing.</p> <p>If a restore operation fails because hard links were incorrectly backed up, the following message appears in the NDMP backup logs:</p> <pre>Bad hardlink path for <path></pre>
CHECKPOINT_INTERVAL_IN_BYTES=	<i><size></i>	5 GB	<p>Specifies the checkpoint interval for a restartable backup. If a restartable backup fails during the backup process, you can restart the backup from where the</p>

Table 9 NDMP environment variables (continued)

Environment variable	Valid values	Default	Description
			<p>process failed. The <i><size></i> parameter is the space between each checkpoint.</p> <p>Note that this variable can only be set from the DMA. For example, if you specify 2 GB, your DMA would create a checkpoint each time 2 GB of data were backed up.</p> <p>Restartable backups are supported only for EMC NetWorker 8.1 and later.</p>

Managing snapshot based incremental backups

After you enable snapshot-based incremental backups, you can view and delete the snapshots created for these backups.

Enable snapshot-based incremental backups for a directory

You can configure OneFS to perform snapshot-based incremental backups for a directory by default. You can also override the default setting in your data management application (DMA).

Procedure

1. Run the `isi ndmp settings variable create` command.

The following command enables snapshot-based incremental backups for `/ifs/data/media`:

```
isi ndmp settings variables create /ifs/data/media BACKUP_MODE
SNAPSHOT
```

View snapshots for snapshot-based incremental backups

You can view snapshots generated for snapshot-based incremental backups.

Procedure

1. Run the following command:

```
isi ndmp dumpdates list
```

Delete snapshots for snapshot-based incremental backups

You can delete snapshots created for snapshot-based incremental backups.

Note

It is recommended that you do not delete snapshots created for snapshot-based incremental backups. If all snapshots are deleted for a path, the next backup performed for the path is a full backup.

Procedure

1. Run the `isi ndmp dumpdates delete` command.

The following command deletes all snapshots created for backing up `/ifs/data/media`:

```
isi ndmp dumpdates delete /ifs/data/media
```

View NDMP backup logs

You can view information about NDMP backup and restore operations through NDMP backup logs.

Procedure

1. Click **Data Protection > Backup > Logs**.
2. In the **Log Location** area, from the **Node** list, select a node.
3. In the **Log Contents** area, review information about NDMP backup and restore operations.

Configuring NDMP backups with EMC NetWorker

You can configure OneFS and EMC NetWorker to backup data stored on an Isilon cluster. The following procedures explain how to configure NDMP backup with EMC NetWorker 8.0.

Create a group

With EMC NetWorker, you must configure a group to manage backups from an Isilon cluster.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. Click **Configuration**.
3. Right-click **Groups** and then click **New**.
4. In the **Name** field, type a name for the group.
5. Click **OK**.

Scan for tape devices

With EMC NetWorker, you must detect tape devices for backup and restore operations.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. Click **Devices**.
3. Right-click **Libraries**, and then click **Scan for Devices**.
4. Ensure that no existing storage nodes are selected.
5. Click **Create a new Storage Node**.
6. Configure the following settings:

Setting name	Setting value
Storage Node Name	The name of the Isilon cluster you want to back up data from
Device Scan Type	Select ndmp
NDMP User Name	The name of an NDMP user on the Isilon cluster
NDMP Password	The password of the NDMP user

7. Click **Start Scan**.

Configure a library

With EMC NetWorker, you must configure the tape library that contains the tape devices for backup and recovery operations.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. Click **Devices**.
3. Right-click **Libraries** and then click **Refresh**.

The system displays a list of tape libraries that are currently attached to the Isilon cluster.

4. Right-click the name of the tape library you want to configure and then click **Configure Library**.
5. In the **Configure Library** window, click **Check All**.
6. Click **Start Configuration**.

Create a data media pool

With EMC NetWorker, you must create a media pool that specifies the type of backups you want to perform and the tape devices you want to use.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. Click **Media**.
3. Click **Media Pools**.
4. In the **Media Pools** area, right-click and then click **New**.
5. Configure the following settings:

Tab	Setting name	Setting value
Basic	Name	A name for the media pool
	Enabled	Selected
	Groups	The group that you created for the Isilon cluster
Selection Criteria	Levels	Select 1-9, <i>full</i> , and <i>incremental</i>
	Devices	Each tape device that you want to use
Configuration	Max parallelism	The maximum number of tape drives to use for concurrent backups

Label tape devices

With EMC NetWorker, you must label tape devices attached to an Isilon cluster before you can back up data to these devices.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. Click **Devices**.
3. Click the name of the library that you configured.
4. In the device list, highlight all tape drives you want to label.
5. Right-click on the highlighted list, and then click **Label**.
6. In the **Label Library Media** window, from the **Target Media Pool** list, select the name of the media pool you created.
7. Ensure that the **Prompt to Overwrite Existing Label** box is cleared.
8. Click **OK**.

Create a metadata media pool

With EMC NetWorker, you must create a media pool for the metadata you want to back up from an Isilon cluster.

Procedure

1. On your local machine, create a directory to contain your metadata.
2. Connect to the NetWorker server from the NetWorker Management Console Server.
3. Configure a new media pool device.
 - a. Click **Devices**.
 - b. Right-click **Devices** and then click **New**.
 - c. Configure the following settings:

Tab	Setting Name	Setting Value
General	Name	A name for the metadata device
	Media Type	file

- d. Click **OK**.
4. Right-click the name of the device and then click **Label**.
5. In the **Label** window, click **OK**.
6. Configure a new media pool.
 - a. Click **Media**.
 - b. Right-click **Media Pools** and then click **New**.
 - c. Configure the following settings:

Tab	Setting name	Setting value
Basic	Name	A name for the metadata media pool
	Enabled	Selected
	Groups	The group that you created for the Isilon cluster
	Label template	Default
Selection Criteria	Save sets	Type the following text: bootstrap Index:
	Devices	The name of the metadata device

- d. Click **OK**.

Create a client

With EMC NetWorker, you must create a client that specifies the data to be backed up from an Isilon cluster.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. Configure the new client.
 - a. Click **Configuration**.
 - b. Click the name of the group you created.
 - c. In the right pane, right-click and then click **New**.
 - d. In the **Create Client** window, in the **Name** field, type a name for the client.
3. In the **Save set** field, type the full path of the directory that you want to back up.
4. From the **Pool** list, select the name of the data media pool you created.
5. Configure the remote user.
 - a. Click **Apps & Modules**.
 - b. In the **Remote user** field, type the name of an NDMP user you configured on the cluster.
 - c. In the **Password** field, type the password of the NDMP user.
6. Select **NDMP**, and in the **Backup command** field, type the backup command.

Option	Description
With DSA	<code>nsrndmp_save -T -M tar</code>
Without DSA	<code>nsrndmp_save -T tar</code>

7. In the **Application information** field, type any NDMP environment variables that you want to specify.

The following text enables directory-based file history and direct access restores (DAR):

```
DIRECT=Y
HIST=F
```

For a complete list of available options, see [NDMP environment variables on page 69](#).

8. Click **Globals (1 of 2)**.

9. In the **Parallelism** field, specify the client parallelism value.

- If you are not using the Data Service Agent (DSA) feature, specify **1**.
- If you are using the Data Service Agent (DSA) feature, specify a value based on your cluster configuration. For more information about client parallelism values, see the *EMC NetWorker Administration Guide*.

10. In the **Aliases** field, specify the short and fully qualified domain name of the Isilon node that you want to backup data from.

11. Click **Globals (2 of 2)**

12. In the **Storage nodes** field, specify the storage node.

- If you are using the DSA feature, type `nsrserverhost` and then press ENTER.
- If you are not using the DSA feature and performing a two-way NDMP backup, type the hostname of the Isilon node you want to backup data from.
- If you are not using the DSA feature and performing a three-way NDMP backup, type the hostname of the tape server. You can specify multiple tape servers by specifying each tape server on a separate line.

13. In the **Remote access** field, type the name of a user on the Isilon cluster.

- If the cluster has not been upgraded to SmartLock compliance mode, type `root@<cluster-host-name>`.
- If the cluster has been upgraded to SmartLock compliance mode, type `compadmin@<cluster-host-name>`.

Configuring NDMP backup with Symantec NetBackup

You can configure OneFS and Symantec NetBackup to backup data stored on an Isilon cluster. The following procedures explain how to configure NDMP backup with Symantec NetBackup 7.5.

Add an NDMP host

You must add an Isilon cluster as an NDMP host before you can backup data with Symantec NetBackup.

Before you begin

Create an NDMP user account.

Procedure

1. In the **NetBackup Administration Console**, expand **Media and Device Management**.
2. Under **Media and Device Management**, expand **Credentials** and then click **NDMP Hosts**.
3. Click **Actions > New > NDMP Host**.
4. In the **NDMP Host Name** dialog box, specify the cluster you want to backup data from.
 - If you have a single Backup Accelerator node in the cluster, type the fully qualified domain name, host name, IPv4 address, or IPv6 address of the Backup Accelerator node.
 - If you have multiple Backup Accelerator nodes in the cluster, type the name of a SmartConnect zone that contains only the Backup Accelerator nodes.
 - If you are performing a three-way NDMP backup, type the fully qualified domain name, host name, SmartConnect zone, IPv4 address, or IPv6 address of any node in the cluster.
5. Click **OK**.
6. In the **Add NDMP Host** box, click **Use the following credentials for this NDMP host on all media servers**.
7. In the **Username** and **Password** fields, type the username and password of an NDMP user on the cluster.
8. Click **OK**.

Configure storage devices

If you are backing up data to tape devices connected to one or more Backup Accelerator nodes, you must configure Symantec NetBackup to recognize those storage devices.

This procedure is required only if you are performing a two-way NDMP backup.

Procedure

1. In the **NetBackup Administration Console**, click **Media and Device Management**.
2. Click **Configure Storage Devices**.
The **Device Configuration Wizard** appears.
3. Click **Next**.
4. Scan the cluster for attached NDMP devices.
 - a. On the **Device Hosts** page, click **Change**.
 - b. Select **NDMP Host**, and then click **OK**.
 - c. Click **Next**.
 - d. Select the NDMP host you created earlier, and then click **Next**.
 - e. After the wizard completes the scan for devices on the cluster, click **Next**.
5. On the **SAN Clients** page, click **Next**.
6. Specify backup devices that NetBackup should use.
 - a. On the **Backup Devices** page, verify that all attached tape devices are displayed in the table, and then click **Next**.
 - b. On the **Drag and Drop Configuration** page, Select the tape devices that you want NetBackup to backup data to and then click **Next**.

- c. In the confirmation dialog box, click **Yes**.
 - d. On the **Updating Device Configuration** page, click **Next**.
 - e. On the **Configure Storage Units** page, view the name of your storage unit and then click **Next**.
 - f. Click **Finish**.
7. Specify the storage unit to associate with the backup devices.
 - a. Expand **NetBackup Management**.
 - b. Expand **Storage**.
 - c. Click **Storage Units**.
 - d. Double-click the name of the storage unit you created previously.
 - e. In the **Change Storage Unit** window, ensure that **Maximum concurrent write drives** is equal to the number of tape drives connected to your cluster.

Results

A storage unit is created for your cluster and tape devices. You can view all storage units by clicking **Storage Units**.

Create a volume pool

Before you can inventory a robot in NetBackup, you must create a volume pool.

Procedure

1. In the **NetBackup Administration Console**, expand **Media and Device Management**.
2. Expand **Media**.
3. Expand **Volume Pools**.
4. Click **Actions > New > Volume Pool**.
5. In the **Pool name** field, type a name for the volume pool.
6. Optional: In the **Description** field, type a description for the volume pool.
7. Click **OK**.

Inventory a robot

Before you create a NetBackup policy, you must inventory a robot with NetBackup and associate it with a volume pool.

Procedure

1. In the **NetBackup Administration Console**, expand **Media and Device Management**.
2. Inventory a robot.
 - a. Expand **Devices**.
 - b. Click **Robots**.
 - c. Right-click the name of the robot that was added when you configured storage devices, and then click **Inventory Robot**.
3. Associate a volume pool with the robot.
 - a. Click **Update volume configuration**.
 - b. Click **Advanced Options**.

- c. From the **Volume Pool** list, select the volume pool you created previously.
 - d. Click **Start**.
 - e. Click **Yes**.
 - f. Click **Close**.
4. Optional: To verify that the robot has been inventoried successfully, click the name of the media pool you created, and ensure that all media are displayed in the table.

Create a NetBackup policy

You must create a NetBackup policy that specifies how you want to back up data from an Isilon cluster.

Procedure

1. In the **NetBackup Administration Console**, expand **Media and Device Management**.
2. Expand **Policies**.
3. Right-click **Summary of all Policies**, and then click **New Policy**.
4. In the **Policy name** field, type a name for the policy and then click **OK**.
5. Configure the following settings:

Setting name	Setting value	Notes
Policy Type	NDMP	Required
Policy volume pool	The name of the volume pool you created	Required
Allow multiple data streams	Selected	Optional. Enables multistreaming. It is recommended that you enable multistreaming whenever possible to increase the speed of backups.
Clients	The Isilon cluster you want to backup data from	Required
Backup Selections	The full path of a directory on the cluster that you want to backup	Required
	set DIRECT=Y	Optional. Enables direct access restore (DAR) for the directory. It is recommended that you enable DAR for all backups.
	set HIST=F	Optional. Specifies path-based file history. It is recommended that you specify path-based file history for all NetBackup backups.
	set NEW_STREAM The full path of a directory on the cluster that you want to backup	Optional. Backs up the specified directory on another stream. It is recommended that you enable multistreaming whenever possible to increase the speed of backups.

Configuring NDMP backup with CommVault Simpana

You can configure OneFS and CommVault Simpana to backup data stored on an Isilon cluster. The following procedures explain how to configure NDMP backup with CommVault Simpana 10.0.

Add a NAS client

With CommValut Simpana, you must add a NAS client for an Isilon cluster before you can backup data from the cluster.

Procedure

1. In the **CommCell Browser**, right-click **Client Computers** and then click **New Client > File System > NAS**.
2. In the **Add NDMP Server** window, configure the following settings:

Setting name	Setting value
NDMP Server Hostname	<p>The cluster you want to backup data from.</p> <ul style="list-style-type: none"> • If you have a single Backup Accelerator node in the cluster, type the fully qualified domain name, host name, IPv4 address, or IPv6 address of the Backup Accelerator node. • If you have multiple Backup Accelerator nodes in the cluster, type the name of a SmartConnect zone that contains only the Backup Accelerator nodes. • If you are performing a three-way NDMP backup, type the fully qualified domain name, host name, SmartConnect zone, IPv4 address, or IPv6 address of any node in the cluster.
NDMP Login	The name of the NDMP user account you configured on the Isilon cluster.
NDMP Password	The password of the NDMP user account you configured on the Isilon cluster.
Listen port	The number of the port through which data management applications (DMAs) can connect to the cluster. This field must match the Port number setting on the Isilon cluster. The default Port number on the Isilon cluster is 10000.

3. Click **Detect**.
The system populates the **Vendor** and **Firmware Revision** fields.
4. Click **OK**.

Add an NDMP library

With CommVault Simpana, you must add an NDMP library to detect tape devices attached to an Isilon cluster before you can backup data to those devices.

Procedure

1. Add the CommVault Simpana server to the configuration.
 - a. In the **CommCell Browser**, click **Storage > Library and Drive**.
 - b. In the **Select MediaAgents** window, add the Simpana server you are currently using, and then click **OK**.

2. Detect NDMP devices attached to the cluster.
 - a. In the **Library and Drive Configuration** window, click **Start > Detect/Configure Devices**.
 - b. Click **NDMP Devices**.
 - c. Click **OK**.
 - d. In the **Select NDMP Servers to Detect** window, add the Isilon cluster you want to backup data from, and then click **OK**. The system informs you that library services will be stopped during the detection process.
 - e. Click **Yes**.
3. After the detection process is complete, close the **Log** window.
4. In the **Library and Drive Configuration** window, select the media changer that controls the tape drives that you want to back up data to.
 You can view the name of the media changer by right-clicking the media changer and then clicking **Properties**.
5. Right-click the media changer you selected, and then click **Configure**.
6. Click **Library and all drives**, and then click **OK**.
7. In the **Confirm** dialog box, specify whether the library has a barcode reader.
8. In the **Discover Media Options** window, specify the default media type.

Create a storage policy

With Commvault Simpana, you must configure a storage policy that specifies the Isilon cluster with the data you want to back up.

Procedure

1. Add and name a new storage policy.
 - a. In the **CommCell Browser**, expand **Policy**.
 - b. Right-click **Storage Policies**, and then click **New Storage Policy**.
 - c. In the **Create Storage Policy Wizard** window, click **Data Protection and Archiving**, and then click **OK**.
 - d. In the **Storage Policy Name** field, type a name for the storage policy, and then click **Next**.
2. Specify the Isilon cluster containing the data you want to back up.
 - a. From the **Library** list, select the name of the NDMP library you configured previously.
 - b. From the **MediaAgent** list, select the Isilon cluster you want to back up data from.
 - c. Click **Next**.
3. From the **Scratch Pool** list, select **Default Scratch**.
4. Optional: To enable multistreaming, specify the **Number of Device Streams** setting as a number greater than one.
 It is recommended that you enable multistreaming to increase the speed of backup operations.
5. Click **Next**.
6. Select **Hardware Compression**, and then click **Next**.

7. Click **Finish**.

Assign a storage policy and schedule to a client

With Commvault Simpana, you must assign a policy and schedule to a client before you can back up data from an Isilon cluster that is associated with the client.

Procedure

1. In the **CommCell Browser**, expand **Client Computers**, expand *⟨Isilon-cluster-name⟩*, expand **NAS**, and then select the name of a backup set.
2. In the right panel, right-click the name of a subclient, and then click **Properties**.
3. Ensure that the following settings are configured:

Tab	Setting name	Setting value	Notes
Storage Device	Storage Policy	The name of the storage policy you created	Required
Content	Backup Content Path	The full path of the directory that you want to back up	Required

4. Right-click the subclient you configured, and then click **Backup**.
5. In the **Select Backup Type** area, select the type of backup.
6. Click **Schedule**, and then click **Configure**.
7. In the **Schedule Details** window, specify the times that you want to back up data, and then click **OK**.
8. Click **OK**.