

Isilon OneFS

Version 8.2.1

Backup and Recovery Guide

September, 2019

Copyright © 2014-2019 Dell Inc. or its subsidiaries. All rights reserved.

Dell 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.” DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

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

CONTENTS

Chapter 1	Introduction to this guide	9
	About this guide.....	10
	Isilon scale-out NAS overview.....	10
	Where to go for support.....	10
Chapter 2	OneFS backup and recovery	11
	OneFS backup and recovery overview.....	12
	SyncIQ data replication overview.....	12
	NDMP backup and recovery overview.....	12
Chapter 3	Data replication with SyncIQ	13
	Replication policies and jobs.....	14
	Automated replication policies.....	15
	Source and target cluster association.....	16
	Configuring SyncIQ source and target clusters with NAT.....	16
	Full and differential replication.....	17
	Controlling replication job resource consumption.....	18
	Replication policy priority.....	18
	Replication reports.....	18
	Replication snapshots.....	19
	Source cluster snapshots.....	19
	Target cluster snapshots.....	19
	Data failover and failback with SyncIQ.....	20
	Data failover.....	20
	Data failback.....	21
	SmartLock compliance mode failover and failback.....	21
	Replication and backup with SmartLock.....	21
	SmartLock replication limitations.....	22
	Data replication and backup with deduplication.....	23
	Recovery times and objectives for SyncIQ.....	23
	SyncIQ license functionality.....	24
Chapter 4	Protecting data with SyncIQ	25
	Creating replication policies.....	26
	Excluding directories in replication.....	26
	Excluding files in replication.....	27
	File criteria options.....	27
	Configure default replication policy settings.....	29
	Create a replication policy.....	29
	Create a SyncIQ domain.....	35
	Assess a replication policy.....	36
	Managing replication to remote clusters.....	36
	Start a replication job.....	36
	Pause a replication job.....	37
	Resume a replication job.....	37
	Cancel a replication job.....	37
	View active replication jobs.....	37
	Replication job information.....	37

Managing replication policies.....	38
Modify a replication policy.....	38
Delete a replication policy.....	39
Enable or disable a replication policy.....	39
View replication policies.....	39
Replication policy information.....	40
Replication policy settings.....	40
Managing replication to the local cluster.....	43
Cancel replication to the local cluster.....	43
Break local target association.....	44
View replication policies targeting the local cluster.....	44
Remote replication policy information.....	44
Managing replication performance rules.....	45
Create a network traffic rule.....	45
Create a file operations rule.....	45
Modify a performance rule.....	45
Delete a performance rule.....	46
Enable or disable a performance rule.....	46
View performance rules.....	46
Managing replication reports.....	46
Configure default replication report settings.....	46
Delete replication reports.....	47
View replication reports.....	47
Replication report information.....	47
Managing failed replication jobs.....	48
Resolve a replication policy.....	48
Reset a replication policy.....	48
Perform a full or differential replication.....	49
Managing changelists.....	49
Create a changelist.....	50
View a changelist.....	50
Changelist information.....	50
Chapter 5	Configuring SyncIQ Failover/Failback Policies
	53
Initiating data failover and failback with SyncIQ.....	54
Fail over data to a secondary cluster.....	54
Revert a failover operation.....	54
Fail back data to a primary cluster.....	55
Performing disaster recovery for older SmartLock directories.....	56
Recover SmartLock compliance directories on a target cluster.....	56
Migrate SmartLock compliance directories.....	57
Chapter 6	NDMP backup and recovery
	59
NDMP backup and restore models overview.....	60
NDMP two-way backup.....	60
NDMP three-way backup.....	60
Supportability of NDMP sessions on 6th Generation hardware.....	60
Installing the optional 2x10GbE + 2x8GB Fibre Channel NIC cards.....	61
Verify the 2x10GbE + 2x8GB Fibre Channel NIC connectivity.....	61
Setting preferred IPs for NDMP three-way operations.....	63
Configure preferred IP settings for NDMP three-way operations.....	63
NDMP multi-stream backup and recovery.....	63
NDMP file list backup overview.....	64
Snapshot-based incremental backups.....	64

	NDMP backup and restore of SmartLink files.....	65
	NDMP protocol support.....	66
	Supported DMAs.....	67
	NDMP hardware support.....	67
	Sharing tape drives between clusters.....	67
	NDMP backup limitations.....	68
	NDMP performance recommendations.....	68
	Excluding files and directories from NDMP backups.....	69
Chapter 7	Administering NDMP	71
	NDMP backup and recovery tasks.....	72
	Configuring basic NDMP backup settings.....	72
	NDMP backup settings.....	72
	Configure and enable NDMP backup.....	72
	Disable NDMP backup.....	73
	View NDMP backup settings.....	73
	Managing NDMP user accounts.....	73
	Create an NDMP administrator account.....	73
	View NDMP user accounts.....	73
	Modify the password of an NDMP administrator account.....	74
	Delete an NDMP administrator account.....	74
	Managing NDMP environment variables.....	74
	NDMP environment variable settings.....	75
	Add an NDMP environment variable.....	75
	View NDMP environment variables.....	76
	Edit an NDMP environment variable.....	76
	Delete an NDMP environment variable.....	76
	NDMP environment variables.....	76
	Setting environment variables for backup and restore operations.....	82
	Managing NDMP contexts.....	83
	Managing NDMP restartable backups.....	83
	NDMP context settings.....	84
	Configure NDMP restartable backup settings.....	84
	View NDMP contexts.....	85
	Delete an NDMP context.....	85
	Configure NDMP restartable backups for NetWorker.....	85
	Managing NDMP sessions.....	86
	NDMP session information.....	86
	View NDMP sessions.....	89
	Abort an NDMP session.....	89
	Managing NDMP Fibre Channel ports.....	89
	NDMP backup port settings.....	89
	Enable or disable an NDMP backup port.....	90
	View NDMP backup ports.....	90
	Modify NDMP backup port settings.....	90
	Managing NDMP preferred IP settings.....	91
	Create an NDMP preferred IP setting.....	91
	Modify an NDMP preferred IP setting.....	91
	List NDMP preferred IP settings.....	92
	View NDMP preferred IP settings.....	92
	Delete NDMP preferred IP settings.....	92
	Managing NDMP backup devices.....	92
	NDMP backup device settings.....	93
	Detect NDMP backup devices.....	93
	View NDMP backup devices.....	94

Modify the name of an NDMP backup device.....	94
Delete an entry for an NDMP backup device.....	94
NDMP dumpdates file overview.....	95
Managing the NDMP dumpdates file.....	95
NDMP dumpdates file settings.....	95
View entries in the NDMP dumpdates file.....	95
Delete entries from the NDMP dumpdates file.....	96
Managing snapshot based incremental backups.....	96
Enable snapshot-based incremental backups for a directory.....	96
View snapshots for snapshot-based incremental backups.....	96
Delete snapshots for snapshot-based incremental backups.....	96
NDMP restore operations.....	97
NDMP parallel restore operation.....	97
NDMP serial restore operation.....	97
Specify a NDMP serial restore operation.....	97
Managing file list backups.....	97
Format of a backup file list.....	98
Placement of the file list.....	99
Start a file list backup.....	99
Managing cluster performance for NDMP sessions.....	99
Enable NDMP Redirector to manage cluster performance.....	100
NDMP Redirector limitations.....	100
Managing CPU usage for NDMP sessions.....	101
Enable NDMP Throttler.....	101
NDMP Throttler limitations.....	101
Configuring NDMP backups with NetWorker.....	102
Create a group.....	102
Create a policy.....	102
Create a workflow.....	102
Scan for tape devices.....	103
Configure a library.....	103
Create a data media pool.....	104
Label tape devices.....	104
Create a client.....	105
Back up data through NetWorker.....	106
Recover backed up data through NetWorker.....	106
Configuring NDMP backup with Symantec NetBackup.....	107
Add an NDMP host.....	107
Configure storage devices.....	107
Create a volume pool.....	108
Inventory a robot.....	109
Create a NetBackup policy.....	109
Configuring NDMP backup with CommVault Simpana.....	110
Add a NAS client.....	110
Add an NDMP library.....	111
Create a storage policy.....	112
Assign a storage policy and schedule to a client.....	113
Configure backup options.....	113
Restore backed up data.....	114
Configuring NDMP backup with IBM Tivoli Storage Manager.....	115
Initialize an IBM Tivoli Storage Manager server for an Isilon cluster....	115
Configure an IBM Tivoli Storage Manager server for an Isilon cluster.	115

Chapter 8	Sparse file support	123
	Commvault sparse file support: data aging and space reclamation.....	124

Enable or disable SMB sparse file support..... 124

CHAPTER 1

Introduction to this guide

This section contains the following topics:

- [About this guide](#)..... 10
- [Isilon scale-out NAS overview](#)..... 10
- [Where to go for support](#)..... 10

About this guide

This guide describes how to protect data on Isilon clusters through the SyncIQ software module and the Network Data Management Protocol (NDMP).

Your suggestions help us to improve the accuracy, organization, and overall quality of the documentation. Send your feedback to <https://www.research.net/s/isi-docfeedback>. If you cannot provide feedback through the URL, send an email message to docfeedback@isilon.com.

Isilon scale-out NAS overview

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

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

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

Where to go for support

This topic contains resources for getting answers to questions about Isilon products.

Online support	<ul style="list-style-type: none"> • Live Chat • Create a Service Request <p>For questions about accessing online support, send an email to support@emc.com.</p>
Telephone support	<ul style="list-style-type: none"> • United States: 1-800-SVC-4EMC (1-800-782-4362) • Canada: 1-800-543-4782 • Worldwide: 1-508-497-7901 • Local phone numbers for a specific country are available at Dell EMC Customer Support Centers.
Isilon Community Network	<p>The Isilon Community Network connects you to a central hub of information and experts to help you maximize your current storage solution. From this site, you can demonstrate Isilon products, ask questions, view technical videos, and get the latest Isilon product documentation.</p>
Isilon Info Hubs	<p>For the list of Isilon info hubs, see the Isilon Info Hubs page on the Isilon Community Network. Use these info hubs to find product documentation, troubleshooting guides, videos, blogs, and other information resources about the Isilon products and features you're interested in.</p>

CHAPTER 2

OneFS backup and recovery

This section contains the following topics:

- [OneFS backup and recovery overview](#) 12
- [SyncIQ data replication overview](#) 12
- [NDMP backup and recovery overview](#) 12

OneFS backup and recovery overview

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

- The SyncIQ software module enables you to replicate data to an Isilon cluster and recover the backed up data through the failover and failback processes. 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.
- NDMP enables you to backup data to a tape device and recover the backed up data to any Isilon cluster. NDMP tape block sizes of 512 KB are supported.

SyncIQ data replication 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 source directory. Metadata such as access control lists (ACL) and alternate data streams (ADS) are replicated along with data.

SyncIQ enables you to maintain a consistent replica of your data on another Isilon cluster and to control the frequency of data replication. For example, you could configure SyncIQ to back up data from your primary cluster to a secondary cluster once a day at 10 PM. Depending on the size of your data set, the first replication operation could take considerable time. After that, however, replication operations would complete more quickly.

SyncIQ also offers automated failover and failback capabilities so you can continue operations on the secondary Isilon cluster should your primary cluster become unavailable.

NDMP backup and recovery overview

In OneFS, you can back up and recover 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).

Some of the NDMP features are described below:

- NDMP supports two-way and three-way backup models.
- With certain data management applications, NDMP supports backup restartable extension (BRE). The NDMP BRE allows you to resume a failed backup job from the last checkpoint taken prior to the failure. The failed job is restarted automatically and cannot be scheduled or started manually.
- 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. If you back up a SnapshotIQ snapshot, OneFS does not create another snapshot for the backup.
- You can back up WORM domains through NDMP.

CHAPTER 3

Data replication with SyncIQ

This section contains the following topics:

- [Replication policies and jobs](#)14
- [Replication snapshots](#) 19
- [Data failover and failback with SyncIQ](#) 20
- [Replication and backup with SmartLock](#) 21
- [Data replication and backup with deduplication](#) 23
- [Recovery times and objectives for SyncIQ](#) 23
- [SyncIQ license functionality](#) 24

Replication policies and jobs

Data replication is coordinated according to replication policies and replication 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 tree on the source cluster are replicated to a directory tree on the target cluster; these directory trees 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. We recommend that you do not create more than 1,000 policies 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, no more than 50 replication jobs can run on a source cluster at the same time. If more than 50 replication jobs exist on a cluster, the first 50 jobs run while the others are queued to run.

There is no limit to the number of replication jobs that a target cluster can support concurrently. However, because more replication jobs require more cluster resources, replication will slow down as more concurrent jobs are added.

When a replication job runs, OneFS generates workers on the source and target cluster. Workers on the source cluster read and send data while workers on the target cluster receive and write data.

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.

Automated replication policies

You can manually start a replication policy at any time, but you can also configure replication policies to start automatically based on source directory modifications or schedules.

You can configure a replication policy to run according to a schedule, so that you can control when replication is performed. You can also configure policies to replicate the data captured in snapshots of a directory. You can also configure a replication policy to start when SyncIQ detects a modification to the source directory, so that SyncIQ maintains a more current version of your data on the target cluster.

Scheduling a policy can be useful under the following conditions:

- You want to replicate data when user activity is minimal
- You can accurately predict when modifications will be made to the data

If a policy is configured to run according to a schedule, you can configure the policy not to run if no changes have been made to the contents of the source directory since the job was last run. However, if changes are made to the parent directory of the source directory or a sibling directory of the source directory, and then a snapshot of the parent directory is taken, SyncIQ will create a job for the policy, even if no changes have been made to the source directory. Also, if you monitor the cluster through the File System Analytics (FSA) feature of InsightIQ, the FSA job will create snapshots of `/ifs`, which will most likely cause a replication job to start whenever the FSA job is run.

Replicating data contained in snapshots of a directory can be useful under the following conditions:

- You want to replicate data according to a schedule, and you are already generating snapshots of the source directory through a snapshot schedule
- You want to maintain identical snapshots on both the source and target cluster
- You want to replicate existing snapshots to the target cluster
To do this, you must enable archival snapshots on the target cluster. This setting can only be enabled when the policy is created.

If a policy is configured to replicate snapshots, you can configure SyncIQ to replicate only snapshots that match a specified naming pattern.

Configuring a policy to start when changes are made to the source directory can be useful under the following conditions:

- You want to retain a up-to-date copy of your data at all times
- You are expecting a large number of changes at unpredictable intervals

For policies that are configured to start whenever changes are made to the source directory, SyncIQ checks the source directories every ten seconds. SyncIQ checks all files and directories underneath the source directory, regardless of whether those files or directories are excluded from replication, so SyncIQ might occasionally run a replication job unnecessarily. For example, assume that `newPolicy` replicates `/ifs/data/media` but excludes `/ifs/data/media/temp`. If a modification is made to `/ifs/data/media/temp/file.txt`, SyncIQ will run `newPolicy`, even though `/ifs/data/media/temp/file.txt` will not be replicated.

If a policy is configured to start whenever changes are made to the source directory, and a replication job fails, SyncIQ waits one minute before attempting to run the policy again. SyncIQ increases this delay exponentially for each failure up to a maximum of eight hours. You can override the delay by running the policy manually at any time. After a job for the policy completes successfully, SyncIQ will resume checking the source directory every ten seconds.

If a policy is configured to start whenever changes are made to the source directory, you can configure SyncIQ to wait a specified period of time after the source directory is modified before starting a job.

Note: To avoid frequent synchronization of minimal sets of changes, and overtaxing system resources, we strongly advise against configuring continuous replication when the source directory is highly active. In such cases, it is often better to configure continuous replication with a change-triggered delay of several hours to consolidate groups of changes.

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.

Breaking a policy association causes either a full replication 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.

CAUTION Changes to the configuration of the target cluster outside of SyncIQ can introduce an error condition that effectively breaks the association between the source and target cluster. For example, changing the DNS record of the target cluster could cause this problem. If you need to make significant configuration changes to the target cluster outside of SyncIQ, make sure that your SyncIQ policies can still connect to the target cluster.

Configuring SyncIQ source and target clusters with NAT

Source and target clusters can use NAT (network address translation) for SyncIQ failover and failback purposes, but must be configured appropriately.

In this scenario, source and target clusters are typically at different physical locations, use private, non-routable address space, and do not have direct connections to the Internet. Each cluster typically is assigned a range of private IP addresses. For example, a cluster with 12 nodes might be assigned IP addresses 192.168.10.11 to 192.168.10.22.

To communicate over the public Internet, source and target clusters must have all incoming and outgoing data packets appropriately translated and redirected by a NAT-enabled firewall or router.

CAUTION SyncIQ data is not encrypted. Running SyncIQ jobs over the public Internet provides no protection against data theft.

SyncIQ enables you to limit replication jobs to particular nodes within your cluster. For example, if your cluster was made up of 12 nodes, you could limit replication jobs to just three of those nodes. For NAT support, you would need to establish a one-for-one association between the source and target clusters. So, if you are limiting replication jobs to three nodes on your source cluster, you must associate three nodes on your target cluster.

In this instance, you would need to configure static NAT, sometimes referred to as inbound mapping. On both the source and target clusters, for the private address assigned to each node, you would associate a static NAT address. For example:

Source cluster			Target Cluster		
Node name	Private address	NAT address	Node name	Private address	NAT address
source-1	192.168.10.11	10.8.8.201	target-1	192.168.55.101	10.1.2.11
source-2	192.168.10.12	10.8.8.202	target-2	192.168.55.102	10.1.2.12
source-3	192.168.10.13	10.8.8.203	target-3	192.168.55.103	10.1.2.13

To configure static NAT, you would need to edit the `/etc/local/hosts` file on all six nodes, and associate them with their counterparts by adding the appropriate NAT address and node name. For example, in the `/etc/local/hosts` file on the three nodes of the source cluster, the entries would look like:

```
10.1.2.11 target-1
10.1.2.12 target-2
10.1.2.13 target-3
```

Similarly, on the three nodes of the target cluster, you would edit the `/etc/local/hosts` file, and insert the NAT address and name of the associated node on the source cluster. For example, on the three nodes of the target cluster, the entries would look like:

```
10.8.8.201 source-1
10.8.8.202 source-2
10.8.8.203 source-3
```

When the NAT server receives packets of SynclQ data from a node on the source cluster, the NAT server replaces the packet headers and the node's port number and internal IP address with the NAT server's own port number and external IP address. The NAT server on the source network then sends the packets through the Internet to the target network, where another NAT server performs a similar process to transmit the data to the target node. The process is reversed when the data fails back.

With this type of configuration, SynclQ can determine the correct addresses to connect with, so that SynclQ can send and receive data. In this scenario, no SmartConnect zone configuration is required.

For information about the ports used by SynclQ, see the *OneFS Security Configuration Guide* for your OneFS version.

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, SynclQ performs either a full replication or a differential replication the next time the policy is run. You can specify the type of replication that SynclQ performs.


During a full replication, SynclQ 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 replication job to consume.

Controlling replication job resource consumption

You can create rules that limit the network traffic created by replication jobs, the rate at which files are sent by replication jobs, the percent of CPU used by replication jobs, and the number of workers created for replication jobs.

If you limit the percentage of total workers that SyncIQ can create, the limit is applied to the total amount of workers that SyncIQ could create, which is determined by cluster hardware. Workers on the source cluster read and send data while workers on the target cluster receive and write data.

 **Note:** File-operation rules might not work accurately for files that can take more than a second to transfer and for files that are not predictably similar in size.

Replication policy priority

When creating a replication policy, you can configure a policy to have priority over other jobs.

If multiple replication jobs are queued to be run because the maximum number of jobs are already running, jobs created by policies with priority will be run before jobs without priorities. For example, assume that 50 jobs are currently running. A job without priority is created and queued to run; next, a job with priority is created and queued to run. The job with priority will run next, even though the job without priority has been queued for a longer period of time.

SyncIQ will also pause replication jobs without priority to allow jobs with priority to run. For example, assume that 50 jobs are already running, and one of them does not have priority. If a replication job with priority is created, SyncIQ will pause the replication job without priority and run the job with priority.


Replication reports

After a replication job completes, SyncIQ generates a replication 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 software module. 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 your primary cluster goes offline, you can fail over to a secondary Isilon cluster, enabling clients to continue accessing their data. If the primary cluster becomes operational again, you can fail back to the primary cluster.

For the purposes of SyncIQ failover and failback, the cluster originally accessed by clients is referred to as the primary cluster. The cluster that client data is replicated to is referred to as the secondary cluster.

Failover is the process that allows clients to access, view, modify, and delete data on a secondary cluster. Failback is the process that allows clients to resume their workflow on the primary cluster. During failback, any changes made to data on the secondary cluster are copied back to the primary cluster by means of a replication job using a mirror policy.

Failover and failback can be useful in disaster recovery scenarios. For example, if a primary cluster is damaged by a natural disaster, you can migrate clients to a secondary cluster where they can continue normal operations. When the primary cluster is repaired and back online, you can migrate clients back to operations on the primary cluster.

You can fail over and fail back to facilitate scheduled cluster maintenance, as well. 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 supported both for enterprise and compliance SmartLock directories. Compliance SmartLock directories adhere to U.S. Securities and Exchange Commission (SEC) regulation 17a-4(f), which requires securities brokers and dealers to preserve records in a non-rewritable, non-erasable format. SyncIQ properly maintains compliance with the 17a-4(f) regulation during failover and failback.

Data failover

Failover is the process of preparing data on a secondary cluster and switching over to the secondary cluster for normal client operations. After you fail over to a secondary cluster, you can direct clients to access, view, and modify their data on the secondary cluster.

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

If the action of a replication policy is set to copy, any file that was deleted on the primary cluster will still 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.

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

Failback is the process of restoring primary and secondary clusters to the roles that they occupied before a failover operation. After failback is complete, the primary cluster holds the latest data set and resumes normal operations, including hosting clients and replicating data to the secondary cluster through SyncIQ replication policies in place.

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 is preparing the primary cluster to be accessed by clients. The final step is resuming data replication from the primary to the secondary cluster. At the end of the failback process, you can redirect users to resume data access on the primary cluster.

To update the primary cluster with the modifications that were made on the secondary cluster, SyncIQ must create a SyncIQ domain for the source directory.

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

- The policy has been failed over.
- The policy is a synchronization policy (not a copy policy).
- The policy does not exclude any files or directories from replication.

SmartLock compliance mode failover and failback

Using OneFS 8.0.1 and later releases, you can replicate SmartLock compliance mode domains to a target cluster. This support includes failover and failback of these SmartLock domains.

Because SmartLock compliance mode adheres to the U.S. Securities and Exchange Commission (SEC) regulation 17a-4(f), failover and failback of a compliance mode WORM domain requires some planning and setup.

Most importantly, both your primary (source) and secondary (target) clusters must be configured at initial setup as compliance mode clusters. This process is described in the Isilon installation guide for your node model (for example, the *Isilon S210 Installation Guide*).

In addition, both clusters must have directories defined as WORM domains with the compliance type. For example, if you are storing your WORM files in the SmartLock compliance domain `/ifs/financial-records/locked` on the primary cluster, you must have a SmartLock compliance domain on the target cluster to fail over to. Although the source and target SmartLock compliance domains can have the same pathname, this is not required.

In addition, you must start the compliance clock on both clusters.

SyncIQ handles conflicts during failover/failback operations on a SmartLock compliance mode domain by unlinking committed files from the user store and leaving a link of the file in the compliance store. The ComplianceStoreDelete job automatically tracks and removes expired files from the compliance store if they were put there as a result of SyncIQ conflict resolution. The job runs automatically once per month or when started manually. For information about how to start the ComplianceStoreDelete job, see the *Isilon OneFS CLI Administration Guide*.

Replication and backup with SmartLock

OneFS enables both compliance and enterprise SmartLock directories to be replicated or backed up to a target cluster.

If you are replicating SmartLock directories with SyncIQ, we recommend that you configure all nodes on the source and target clusters with Network Time Protocol (NTP) peer mode to ensure that the node clocks are synchronized. For compliance clusters, we recommend that you configure all nodes on the source and target clusters with NTP peer mode before you set the compliance

clocks. This sets the source and target clusters to the same time initially and helps to ensure compliance with U.S. Securities and Exchange Commission rule 17a-4.

Note: If you replicate data to a SmartLock directory, do not configure SmartLock settings for that directory until you are no longer replicating data to the directory. Configuring an autocommit time period for a SmartLock target directory, for example, can cause replication jobs to fail. If the target 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 overwrite the committed file.

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 recover data to a SmartLock directory on the cluster, the metadata persists on the cluster. However, if the directory that you recover data to is not a SmartLock directory, the metadata is lost. You can recover data to a SmartLock directory only if the directory is empty.

For information on the limitations of replicating and failing back SmartLock directories with SyncIQ, see [SmartLock replication limitations](#) on page 22.

SmartLock replication limitations

Be aware of the limitations of replicating and failing back SmartLock directories with SyncIQ.

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

Source directory type	Target directory type	Replication Allowed	Failback allowed
Non-SmartLock	Non-SmartLock	Yes	Yes
Non-SmartLock	SmartLock enterprise	Yes	Yes, unless files are committed to a WORM state on the target cluster
Non-SmartLock	SmartLock compliance	No	No
SmartLock enterprise	Non-SmartLock	Yes; however, retention dates and commit status of files will be lost.	Yes; however the files will not have WORM status
SmartLock enterprise	SmartLock enterprise	Yes	Yes; any newly committed WORM files will be included
SmartLock enterprise	SmartLock compliance	No	No
SmartLock compliance	Non-SmartLock	No	No
SmartLock compliance	SmartLock enterprise	No	No
SmartLock compliance	SmartLock compliance	Yes	Yes; any newly committed WORM files will be included

If you are replicating a SmartLock directory to another SmartLock directory, you must create the target SmartLock directory prior to running the replication policy. Although OneFS will create a

target directory automatically if a target directory does not already exist, OneFS will not create a target SmartLock directory automatically. If you attempt to replicate an enterprise directory before the target directory has been created, OneFS will create a non-SmartLock target directory and the replication job will succeed. If you replicate a compliance directory before the target directory has been created, the replication job will fail.

If you replicate SmartLock directories to another Isilon 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.

Data replication and backup with deduplication

When deduplicated files are replicated to another Isilon cluster or backed up to a tape device, the deduplicated files no longer share blocks on the target Isilon cluster or backup device. However, although you can deduplicate data on a target Isilon cluster, you cannot deduplicate data on an NDMP backup device.

Shadows stores are not transferred to target clusters or backup devices. Because of this, deduplicated files do not consume less space than non-deduplicated files when they are replicated or backed up. To avoid running out of space, you must ensure that target clusters and tape devices have enough free space to store deduplicated data as if the data had not been deduplicated. To reduce the amount of storage space consumed on a target Isilon cluster, you can configure deduplication for the target directories of your replication policies. Although this will deduplicate data on the target directory, it will not allow SyncIQ to transfer shadow stores. Deduplication is still performed by deduplication jobs running on the target cluster.

The amount of cluster resources required to backup and replicate deduplicated data is the same as for non-deduplicated data. You can deduplicate data while the data is being replicated or backed up.

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

Protecting data with SyncIQ

This section contains the following topics:

- [Creating replication policies](#)..... 26
- [Managing replication to remote clusters](#).....36
- [Managing replication policies](#).....38
- [Managing replication to the local cluster](#)..... 43
- [Managing replication performance rules](#).....45
- [Managing replication reports](#).....46
- [Managing failed replication jobs](#).....48
- [Managing changelists](#)..... 49

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:** Failback is not supported for 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 character	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 > SynclQ > Settings**.
2. In the **Default Policy Settings** section, if you want policies to connect only to nodes in a specified SmartConnect zone, select **Connect only to the nodes within the target cluster SmartConnect zone**.

Note: This option will affect only policies that specify the target cluster as a SmartConnect zone.

3. Specify which nodes you want replication policies to connect to when a policy is run.

Option	Description
To connect policies to all nodes on a source cluster:	Click Run the policy on all nodes in this cluster .
To 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: SynclQ 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 **Save Changes**.

Create a replication policy

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

About this task

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

Note: If you create a replication policy for a SmartLock compliance directory, the SyncIQ and SmartLock compliance domains must be configured at the same root level. A SmartLock compliance domain cannot be nested inside a SyncIQ domain.

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. For the **Action** setting, 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. For the **Run Job** setting, specify whether replication jobs will be run.

Option	Description
Run jobs only when manually initiated by a user.	Click Manually .
Run jobs automatically according to a schedule.	<ol style="list-style-type: none"> a. Click On a schedule. b. Specify a schedule. 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. c. To prevent the policy from being run when the contents of the source directory have not been modified, click Only run if source directory contents are modified. d. To create OneFS events if a specified RPO is exceeded, click Send RPO alerts after... and then specify an RPO. For example, assume you set an RPO of 5 hours; a job starts at 1:00 PM and completes at 3:00 PM; a second job starts at 3:30 PM; if the second job does not complete by 6:00 PM, SyncIQ will create a OneFS event. <p>Note: This option is valid only if RPO alerts have been globally enabled through SyncIQ settings. The events have an event ID of 400040020.</p>
Run jobs automatically every time that a change is made to the source directory.	<ol style="list-style-type: none"> a. Click Whenever the source is modified. b. To configure SyncIQ to wait a specified amount of time after the source directory is modified before starting a replication

Option	Description
	job, click Change-Triggered Sync Job Delay and then specify a delay.
Runs jobs automatically every time that a snapshot is taken of the source directory.	<p>a. Click Whenever a snapshot of the source directory is taken.</p> <p>b. To only replicate only data contained in snapshots that match a specific naming pattern, type a snapshot naming pattern into the Run job if snapshot name matches the following pattern box.</p> <p>c. To replicate data contained in all snapshots that were taken of the source directory before the policy was created, click Sync existing snapshots before policy creation time.</p>


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.

About this task

 **CAUTION** In a SyncIQ replication policy, OneFS enables you to specify a source directory that is a target directory, or is contained within a target directory, from a different replication policy. Referred to as cascading replication, this use case is specifically for backup purposes, and should be used carefully. OneFS does not allow fallback in such cases.

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 directory `/ifs/.snapshot` or a subdirectory of it.
2. (Optional) Prevent specific subdirectories of the source 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.

Option	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.	<p>a. Click Run the policy only on nodes in the specified subnet and pool.</p> <p>b. From the Subnet and pool list, select the subnet and pool.</p>

 **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

- In the **Target Cluster** area, in the **Target Host** field, type one of the following:
 - The fully qualified domain name (FQDN) 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.

- 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, make sure that the directory is not the target of another replication policy. If this is a synchronization policy, make sure that the directory is empty. All files are deleted from the target of a synchronization policy the first time that the policy is run.

If the specified target directory does not already exist on the target cluster, the directory is created the first time that the job is run. We recommend 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, which prevents 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 to specify 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.

About this task

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, select **Enable capture of snapshots on the target cluster**.
2. (Optional) To modify the default alias of the last snapshot that is 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 that are created according to the replication policy, in the **Snapshot Naming Pattern** field, type a naming pattern. Each snapshot that is generated for this replication policy is assigned a name that is 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 for how snapshots should expire:
 - 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 **Priority** field, specify whether the policy has priority.
Selecting **Normal** will cause jobs created by the policy not to have priority. Selecting **High** will give priority to jobs created by the replication policy.
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:

- **Fatal**
- **Error**
- **Notice**
- **Info**
- **Copy**
- **Debug**
- **Trace**

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.

 **Note:** The recommended log level is **Notice**.

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 increase the speed of failback for the policy, click **Prepare policy for accelerated failback performance**.
Selecting this option causes SyncIQ to perform failback configuration tasks the next time that a job is run, rather than waiting to perform those tasks during the failback process. This will reduce the amount of time needed to perform failback operations when failback is initiated.
5. (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.
6. (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.

7. Specify how the policy replicates CloudPools SmartLink files.

If set to **Deny**, SyncIQ replicates all CloudPools SmartLink files to the target cluster as SmartLink files; if the target cluster does not support CloudPools, the job will fail. If set to **Force**, SyncIQ replicates all SmartLink files to the target cluster as regular files. If set to **Allow**, SyncIQ will attempt to replicate SmartLink files to the target cluster as SmartLink files; if the target cluster does not support CloudPools, SyncIQ will replicate the SmartLink files as regular files.

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. In the **Create SyncIQ Policy** dialog box, after all policy settings are as intended, click **Create 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.

About this task

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.

About this task

 **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.

About this task

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.

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 replication policies

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

Modify a replication policy

You can modify the settings of a replication policy.

About this task

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. After 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.

About this task

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**.

Note: The operation will not succeed until SyncIQ can communicate with the target cluster; until then, the policy will not be removed from the **SyncIQ Policies** table. After the connection between the source cluster and target cluster is reestablished, SyncIQ will delete the policy the next time that the job is scheduled to run; if the policy is configured to run only manually, you must manually run the policy again. If SyncIQ is permanently unable to communicate with the target cluster, run the `isi sync policies delete` command with the `--local-only` option. This will delete the policy from the local cluster only and not break the target association on the target cluster. For more information, see the *OneFS CLI Administration Guide*.

Enable or disable a replication policy

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

About this task

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 a snapshot 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:

- Fatal
- Error
- Notice
- Info
- Copy
- Debug
- Trace

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.

 **Note:** Notice is the recommended log level.

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 clusters.

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 allows writes to the target directory but also requires you to reset the replication policy before you can run the policy again.

About this task

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.

Managing replication performance rules

You can manage the impact of replication on cluster performance by creating rules that limit the network traffic created and the rate at which files are sent by replication jobs.

Create a network traffic rule

You can create a network traffic rule that limits the amount of network traffic that replication policies are allowed to generate during a specified time period.

Procedure

1. Click **Data Protection > SyncIQ > Performance Rules**.
2. Click **Create a SyncIQ Performance Rule**.
3. From the **Rule Type** list, select **Bandwidth**.
4. In the **Limit** field, specify the maximum number of kilobits per second that replication policies are allowed to send.
5. In the **Schedule** area, specify the time and days of the week that you want to apply the rule.
6. Click **Create Performance Rule**.

Create a file operations rule

You can create a file-operations rule that limits the number of files that replication jobs can send per second.

Procedure

1. Click **Data Protection > SyncIQ > Performance Rules**.
2. Click **Create a SyncIQ Performance Rule**.
3. From the **Rule Type** list, select **Bandwidth**.
4. In the **Limit** field, specify the maximum number of files per second that replication policies are allowed to send.
5. In the **Schedule** area, specify the time and days of the week that you want to apply the rule.
6. Click **Create Performance Rule**.

Modify a performance rule

You can modify a performance rule.

Procedure

1. Click **Data Protection > SyncIQ > Performance Rules**.
2. In the **SyncIQ Performance Rules**, in the row for the rule you want to modify, click **View/Edit**.
3. Click **Edit Performance Rule**.
4. Modify rule settings, and then click **Save Changes**.

Delete a performance rule

You can delete a performance rule.

Procedure

1. Click **Data Protection > SyncIQ > Performance Rules**.
2. In the **SyncIQ Performance Rules** table, in the row for the rule you want to delete, select **Delete Rule**.
3. In the **Confirm Delete** dialog box, click **Delete**.

Enable or disable a performance rule

You can disable a performance rule to temporarily prevent the rule from being enforced. You can also enable a performance rule after it has been disabled.

Procedure

1. Click **Data Protection > SyncIQ > Performance Rules**.
2. In the **SyncIQ Performance Rules** table, in the row for a rule you want to enable or disable, select either **Enable Rule** or **Disable Rule**.

View performance rules

You can view information about replication performance rules.

Procedure

1. Click **Data Protection > SyncIQ > Performance Rules**.
2. In the **SyncIQ Performance Rules** table, view information about performance rules.

Managing replication reports

In addition to viewing replication reports, you can configure how long reports are retained on the cluster. You can also delete any reports that have passed their expiration period.

Configure default replication report settings

You can configure the default amount of time that SyncIQ retains replication reports for. You can also configure the maximum number of reports that SyncIQ retains for each replication policy.

Procedure

1. Click **Data Protection > SyncIQ > Settings**.
2. In the **Report Settings** area, in the **Keep Reports For** area, specify how long you want to retain replication reports for.

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 you originally enter them. 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 for display.

3. In the **Number of Reports to Keep Per Policy** field, type the maximum number of reports you want to retain at a time for a replication policy.
4. Click **Submit**.

Delete replication reports

Replication reports are routinely deleted by SyncIQ after the expiration date for the reports has passed. SyncIQ also deletes reports after the number of reports exceeds the specified limit. Excess reports are periodically deleted by SyncIQ; however, you can manually delete all excess replication reports at any time. This procedure is available only through the command-line interface (CLI).

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster, and log in.
2. Delete excess replication reports by running the following command:

```
isi sync reports rotate
```

View replication reports

You can view replication reports and subreports.

Procedure

1. Click **Data Protection > SyncIQ > Reports**.
2. In the **SyncIQ Reports** table, in the row for a report, click **View Details**.

If a report is composed of subreports, the report is displayed as a folder. Subreports are displayed as files within report folders.

Replication report information

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

Policy Name

The name of the associated policy for the job. You can view or edit settings for the policy by clicking the policy name.

Status

Displays 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.

Finished

The job completed successfully.

Failed

The job failed to complete.

Started

Indicates when the job started.

Ended

Indicates when the job ended.

Duration

Indicates how long the job took to complete.

Transferred

The total number of files that were transferred during the job run, and the total size of all transferred files. For assessed policies, `Assessment` appears.

Source Directory

The path of the source directory on the source cluster.

Target Host

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


Action

Displays any report-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, 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.

About this task

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. You can do this only from the CLI.

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 changelists

You can create and view changelists that describe the differences between two snapshots. You can create a changelist for any two snapshots that have a common root directory.

Changelists are most commonly accessed by applications through the OneFS Platform API. For example, a custom application could regularly compare the two most recent snapshots of a critical directory path to determine whether to back up the directory, or to trigger other actions.

Create a changelist

You can create a changelist to view the differences between two snapshots.

Procedure

1. (Optional) Record the IDs of the snapshots.
 - a. Click **Data Protection > SnapshotIQ > Snapshots**.
 - b. In the row of each snapshot that you want to create a changelist for, click **View Details**, and record the ID of the snapshot.
2. Click **Cluster Management > Job Operations > Job Types**.
3. In the **Job Types** area, in the **ChangelistCreate** row, from the **Actions** column, select **Start Job**.
4. In the **Older Snapshot ID** field, type the ID of the older snapshot.
5. In the **Newer Snapshot ID** field, type the ID of the newer snapshot.
6. Click **Start Job**.

View a changelist

You can view a changelist that describes the differences between two snapshots. This procedure is available only through the command-line interface (CLI).

Procedure

1. View the IDs of changelists by running the following command:

```
isi_changelist_mod -l
```

Changelist IDs include the IDs of both snapshots used to create the changelist. If OneFS is still in the process of creating a changelist, `inprog` is appended to the changelist ID.


2. (Optional) View all contents of a changelist by running the `isi_changelist_mod` command with the `-a` option.

The following command displays the contents of a changelist named `2_6`:

```
isi_changelist_mod -a 2_6
```

Changelist information

You can view the information contained in changelists.

 **Note:** The information contained in changelists is meant to be consumed by applications through the OneFS Platform API.

The following information is displayed for each item in the changelist when you run the `isi_changelist_mod` command:

st_ino

Displays the inode number of the specified item.

st_mode

Displays the file type and permissions for the specified item.

st_size

Displays the total size of the item in bytes.

st_atime

Displays the POSIX timestamp of when the item was last accessed.

st_mtime

Displays the POSIX timestamp of when the item was last modified.

st_ctime

Displays the POSIX timestamp of when the item was last changed.

cl_flags

Displays information about the item and what kinds of changes were made to the item.

01

The item was added or moved under the root directory of the snapshots.

02

The item was removed or moved out of the root directory of the snapshots.

04

The path of the item was changed without being removed from the root directory of the snapshot.

10


The item either currently contains or at one time contained Alternate Data Streams (ADS).

20

The item is an ADS.

40

The item has hardlinks.

 **Note:** These values are added together in the output. For example, if an ADS was added, the code would be `cl_flags=021`.

path

The absolute path of the specified file or directory.

CHAPTER 5

Configuring SyncIQ Failover/Failback Policies

This section contains the following topics:

- [Initiating data failover and failback with SyncIQ](#)..... 54
- [Performing disaster recovery for older SmartLock directories](#)..... 56

Initiating data failover and failback with SyncIQ

You can fail over from one Isilon cluster to another if, for example, your primary cluster becomes unavailable. You can fail back when 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.

Note: Data failover and failback are now supported for both compliance SmartLock directories and enterprise SmartLock directories. Compliance SmartLock directories can be created only on clusters that have been set up as compliance mode clusters during initial configuration.

Fail over data to a secondary cluster

You can fail over to a secondary Isilon cluster if your primary cluster becomes unavailable.

Before you begin

You must have created and successfully run a replication policy on the primary cluster. This action replicates data to the secondary cluster.

About this task

Note: Data failover is supported both for SmartLock enterprise and compliance directories. A SmartLock compliance directory requires its own separate replication policy.

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

Procedure

1. If your primary cluster is still online, complete the following steps:
 - a. Stop all writes to the replication policy's path, including both local and client activity.

This action ensures that new data is not written to the policy path as you prepare for failover to the secondary cluster.
 - b. Modify the replication policy so that it is set to run only manually.

This action prevents the policy on the primary cluster from automatically running a replication job. If the policy on the primary cluster runs a replication job while writes are allowed to the target directory, the job fails and the replication policy is deactivated. If this happens, modify the 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.
2. On the secondary cluster, click **Data Protection > SyncIQ > Local Targets**.
3. In the **SyncIQ Local Targets** table, select **More > Allow Writes** for a replication policy.
4. Re-enable client access, and direct users to begin accessing their data from 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.

About this task

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

Before you can fail back to the primary cluster, you must already have failed over to the secondary cluster. Also, you must ensure that your primary cluster is back online.

About this task

Note: Data failback is supported for SmartLock compliance and enterprise directories. If clients committed new SmartLock files while the secondary cluster was in operation, these SmartLock files are replicated to the primary cluster during failback.

Procedure

1. On the primary cluster, click **Data Protection > SyncIQ > Policies**.
2. In the **SyncIQ Policies** list, for a replication policy, click **More > Resync-prep**.
This action causes SyncIQ to create a mirror policy for the replication policy on the primary cluster and secondary cluster. The mirror policy is placed under **Data Protection > SyncIQ > Local Targets** on the primary cluster. On the secondary cluster, the mirror policy is placed under **Data Protection > SyncIQ > Policies**.
SyncIQ names mirror policies according to the following pattern:

```
<replication-policy-name>_mirror
```
3. Before beginning the failback process, prevent clients from accessing the secondary cluster.
This action ensures that SyncIQ fails back the latest data set, including all changes that users made to data on the secondary cluster while the primary cluster was out of service. We recommend that you wait until client activity is low before preventing access to the secondary cluster.
4. On the secondary cluster, click **Data Protection > SyncIQ > Policies**.
5. In the **SyncIQ Policies** list, for the mirror policy, click **More > Start Job**.
Alternatively, you can edit the mirror policy on the secondary cluster, and specify a schedule for the policy to run.
6. On the primary cluster, click **Data Protection > SyncIQ > Local Targets**.
7. On the primary cluster, in the **SyncIQ Local Targets** list, for the mirror policy, select **More > Allow Writes**.

8. On the secondary cluster, click **Data Protection > SyncIQ > Policies**.
9. On the secondary cluster, in the **SyncIQ Policies** list, click **More > Resync-prep** for the mirror policy.


This puts the secondary cluster back into read-only mode and ensures that the data sets are consistent on both the primary and secondary clusters.

After you finish

Redirect clients to begin accessing their data on the primary cluster. Although not required, it is safe to remove a mirror policy after failback has completed successfully.

Performing disaster recovery for older SmartLock directories

If you replicated a SmartLock compliance directory to a secondary cluster running OneFS 7.2.1 or earlier, you cannot fail back the SmartLock compliance directory to a primary cluster running OneFS 8.0.1 or later. However, you can recover the SmartLock compliance directory stored on the secondary cluster, and migrate it back to the primary cluster.

 **Note:** Data failover and failback with earlier versions of OneFS are supported for SmartLock enterprise directories.

Recover SmartLock compliance directories on a target cluster

You can recover compliance SmartLock directories that you have replicated to a secondary cluster running OneFS 7.2.1 or earlier versions.

About this task

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

Procedure

1. On the secondary cluster, click **Data Protection > SyncIQ > Local Targets**.
2. In the **SyncIQ Local Targets** table, for 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**. Note that you should only break the association if the primary cluster has been taken offline permanently.
3. If you clicked **Break Association**, recover 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 original source cluster would not be scheduled to be committed at the same time on the target cluster. To make sure 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 worm domains modify /ifs/data/smartlock --autocommit-offset 1m
```

After you finish

Redirect clients to begin accessing the target cluster.

Migrate SmartLock compliance directories

You can migrate SmartLock compliance directories from a recovery cluster, either by replicating the directories back to the original source cluster, or to a new cluster. Migration is necessary only when the recovery cluster is running OneFS 7.2.1 or earlier. These OneFS versions do not support failover and failback of SmartLock compliance directories.

Procedure

1. On the recovery cluster, create a replication policy for each SmartLock compliance directory that you want to migrate to another cluster (the original primary cluster or a new cluster).

The policies must meet the following requirements:

- The source directory on the recovery cluster is the SmartLock compliance directory that you are migrating.
- The target directory is an empty SmartLock compliance directory on the cluster to which the data is to be migrated. The source and target directories must both be SmartLock compliance directories.

2. Replicate recovery data to the target directory by running the policies that you created. You can replicate data either by manually starting the policies or by specifying a schedule.
3. (Optional) To ensure that SmartLock protection is enforced for all files, commit all migrated files in the SmartLock target directory to a WORM state.

Because autocommit information is not transferred from the recovery cluster, commit all migrated 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 worm domains modify /ifs/data/smartlock --autocommit-offset 1m
```

This step is unnecessary if you have configured an autocommit time period for the SmartLock directories being migrated.

4. On the cluster with the migrated data, click **Data Protection > SyncIQ > Local Targets**.
5. In the **SyncIQ Local Targets** table, for each replication policy, select **More > Allow Writes**.
6. (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 on the cluster now containing the migrated data.

7. (Optional) Delete the copy of the SmartLock data on the recovery cluster.

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 recovery cluster.

CHAPTER 6

NDMP backup and recovery

This section contains the following topics:

- [NDMP backup and restore models overview](#) 60
- [NDMP two-way backup](#) 60
- [NDMP three-way backup](#) 60
- [Supportability of NDMP sessions on 6th Generation hardware](#) 60
- [Setting preferred IPs for NDMP three-way operations](#) 63
- [NDMP multi-stream backup and recovery](#) 63
- [NDMP file list backup overview](#) 64
- [Snapshot-based incremental backups](#) 64
- [NDMP backup and restore of SmartLink files](#) 65
- [NDMP protocol support](#) 66
- [Supported DMAs](#) 67
- [NDMP hardware support](#) 67
- [Sharing tape drives between clusters](#) 67
- [NDMP backup limitations](#) 68
- [NDMP performance recommendations](#) 68
- [Excluding files and directories from NDMP backups](#) 69

NDMP backup and restore models overview

NDMP supports both three-way and two-way backup models.

A two-way NDMP backup process is significantly faster than the three-way NDMP backup process. It is also the most efficient method in terms of cluster resource consumption. However, a two-way NDMP backup process requires that you attach one or more Backup Accelerator nodes or Fibre Attached Storage nodes to the cluster. In both the 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.

NDMP two-way backup

The NDMP two-way backup is also known as the local or direct NDMP backup. To perform NDMP two-way backups, you must connect your Isilon cluster to a Backup Accelerator node which is synonymous with a Fibre Attached Storage node, and attach a tape device to that 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 Fibre Attached Storage node. Alternatively, you can connect Fibre Channel switches to the Fibre Channel ports on the Fibre Attached Storage 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 Fibre Attached Storage node and the connected tape and media changer devices.

If you attach tape devices to a Fibre Attached Storage 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.

i Note: Generation 6 nodes added to an InfiniBand back end network are supported with the A100 Backup Accelerator as part of an NDMP 2-way backup solution. The A100 Backup Accelerator is not supported as part of an NDMP two-way backup solution with an all-Generation 6 cluster with an Ethernet back end.

NDMP three-way backup


The NDMP three-way backup is also known as the remote NDMP backup.

During a three-way NDMP 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. The NDMP service runs on one NDMP Server and the NDMP tape service runs on a separate server. Both the servers are connected to each other across the network boundary.

Supportability of NDMP sessions on 6th Generation hardware

You can enable two-way NDMP sessions on 6th Generation nodes by configuring them with Sheba cards. A Sheba card is a Fibre Channel hybrid host bus adapter (HBA) that enables two-way

NDMP sessions over the Fibre Channel port. You must contact Isilon Professional Services to enable the Sheba card support.

 **Note:** Sheba card is not supported with F810 nodes.

Installing the optional 2x10GbE + 2x8GB Fibre Channel NIC cards

Install the optional 2x10GbE + 2x8GB Fibre Channel NICs in a paired configuration. On any chassis where you install a 2x10GbE + 2x8GB Fibre Channel NIC, install either two or four cards in the chassis. The optional 2x10GbE + 2x8GB Fibre Channel NIC replaces the front-end Ethernet NIC with a hybrid 4-port (2x10GbE and 2xFC) NIC or HBA to support simultaneous traffic coming in both through either port.

Verify the 2x10GbE + 2x8GB Fibre Channel NIC connectivity

Before you begin

Follow the steps in this procedure to verify whether the 2x10GbE + 2x8GB Fibre Channel NIC connectivity is established.

On the node on which the 2x10GbE + 2x8GB Fibre Channel NIC is installed, perform the following steps:

Procedure

1. Verify that the 2x10GbE + 2x8GB Fibre Channel NIC cards are available:

```
# ls /dev/ocs*
```

The output should show `ocs0` and `ocs1`.

2. Verify that the cluster Fibre Channel connectivity and zoning are accurate by matching the Fibre Channel WWPNs output displayed with the WWPNs displayed in the DMA user interface:

```
# isi fc settings ls
```

In the following sample output, the WWPNs should match the ones displayed in the DMA:

Port	WWNN	WWPN	State	Topology	Rate
1.1	200000109b12baa0	100000109b12baa0	enable	nport	8
1.2	200000109b12baa1	100000109b12baa1	enable	nport	8
2.1	200000109b12b5ce	100000109b12b5ec	enable	nport	8
2.2	200000109b12b5cd	100000109b12b5ed	enable	nport	8
3.1	200000109b12b610	100000109b12b610	enable	nport	8
3.2	200000109b12b611	100000109b12b611	enable	nport	8
4.1	200000109b12b0b8	100000109b12b0b8	enable	nport	8
4.2	200000109b12b0b9	100000109b12b0b9	enable	none	auto

3. Detect NDMP backup devices:

```
# isi tape rescan --reconcile
```

A sample output of the previous command is shown:

Node	Devicename	Path	Status Report	Serial
1	Tape037	1:1:6	New path added	9382630076
1	Tape037	1:1:25	New path added	9382630076
1	Tape037	1:2:6	New path added	9382630076
1	Tape037	1:1:25	New path added	9382630076
1	Tape038	1:1:6	New path added	9382630077
1	Tape038	1:1:25	New path added	9382630077
1	Tape038	1:2:6	New path added	9382630077
1	Tape038	1:1:25	New path added	9382630077
1	Tape039	1:1:6	New path added	9382630078
1	Tape039	1:1:25	New path added	9382630078
1	Tape039	1:2:6	New path added	9382630078
1	Tape039	1:1:25	New path added	9382630078

```
Cont..
Wwnn                Product
200000218809183    IBM ULTRIUM
200000218809183    IBM ULTRIUM
200000218809183    IBM ULTRIUM
200000218809183    IBM ULTRIUM
00000218809183     IBM ULTRIUM
200000218809183    IBM ULTRIUM
200000218809183    IBM ULTRIUM
200000218809183    IBM ULTRIUM
200000218809183    IBM ULTRIUM
200000218809183    IBM ULTRIUM
200000218809183    IBM ULTRIUM
200000218809183    IBM ULTRIUM
200000218809183    IBM ULTRIUM
```

4. Verify that zoned devices are visible and accessible, and that the NICs in the same zone have multiple paths. The following is an example of the command on tape04:

```
isi tape view tape041
```

The output is similar to the following:

```
Name: tape041
State: closed
WNN: 2000002188009183
Vendor/Model/Revision: <IBM ULTRUYN-TD3 8711>
Serial: 938263007A
OpenCount: 0
Paths
  Path: 2:1:6
  LUN: 32
  State: inactive
  WWP: 2400002188009183
  Portid: A80300
  OpenCount: 0
  DevName: sa83
  PassName: pass12B

  Path: 2:1:25
  LUN: 32
  State: inactive
  WWP: 2410002188009183
  Portid: A81700
  OpenCount: 0
```

```
DevName: sa89
PassName: pass135
```

Results

NDMP backup and restore operations are enabled by using the 2x10GbE + 2x8GB Fibre Channel NIC.

Setting preferred IPs for NDMP three-way operations

If you are using Avamar as your data management application (DMA) for an NDMP three-way operation in an environment with multiple network interfaces, you can apply a preferred IP setting across an Isilon cluster or to one or more subnets that are defined in OneFS. A preferred IP setting is a list of prioritized IP addresses to which a data server or tape server connects during an NDMP three-way operation.

The IP address on the NDMP server that receives the incoming request from the DMA decides the scope and precedence for setting the preference. If the incoming IP address is within a subnet scope that has a preference, then the preference setting is applied. If a subnet-specific preference does not exist but a cluster-wide preference exists, the cluster-wide preference setting is applied. Subnet-specific preference always overrides the cluster-wide preference. If both the cluster-wide and subnet-specific preferences do not exist, the IP addresses within the subnet of the IP address receiving the incoming requests from the DMA are used as the preferred IP addresses.

You can have one preferred IP setting per cluster or per network subnet.

You can specify a list of NDMP preferred IPs through the `isi ndmp settings preferred-ips` command.

Configure preferred IP settings for NDMP three-way operations

You can configure preferred IP settings for NDMP three-way operations and apply the settings to one or more subnets that are defined in OneFS.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the following command to specify a preferred IP setting to a subnet defined in OneFS:

```
isi ndmp settings preferred-ips create <scope> <subnet>
```


NDMP multi-stream backup and recovery

You can use the NDMP multi-stream backup feature, in conjunction with certain data management applications (DMAs), to speed up backups.

With multi-stream backup, you can use your DMA to specify multiple streams of data to back up concurrently. OneFS considers all streams in a specific multi-stream backup operation to be part of the same backup context. A multi-stream backup context is deleted if a multi-stream backup session is successful. If a specific stream fails, the backup context is retained for five minutes after the backup operation completes and you can retry the failed stream within that time period.

If you used the NDMP multi-stream backup feature to back data up to tape drives, you can also recover that data in multiple streams, depending on the DMA. In OneFS 8.0.0.0 and later releases, multi-stream backups are supported with CommVault Simpana version 11.0 Service Pack 3 and

NetWorker version 9.0.1. If you back up data using CommVault Simpana, a multi-stream context is created, but data is recovered one stream at a time.

 **Note:** OneFS multi-stream backups are not supported with the NDMP restartable backup feature.

NDMP file list backup overview

The NDMP file list backup process allows you to perform a backup operation based on a list of files.

An NDMP backup session can either be full or incremental. A full backup session backs up every file under a path. An incremental backup session backs up changed files under the same path of a full backup. For both the backup types, you can specify the files to be backed up by generating a file list. A data management application (DMA) can pass the `BACKUP_FILE_LIST` environment variable to specify a file path. The file pointed by the file path can have a list of files relative to the backup directory that need to be backed up. The NDMP file list backup operation processes the files in the file list based on the `BACKUP_FILE_LIST` environment variable setting without traversing through the entire backup directory for backing up the files.

Refer to the sections *Managing file list backups*, *Placement of the file list*, and *Start a file list backup* in this guide for information about how to work with file lists.

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 files 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 the change rate is low.

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.

Set the `BACKUP_MODE` environment variable to `SNAPSHOT` to enable snapshot-based incremental backups. If you enable 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 enable snapshot-based incremental backups, OneFS automatically deletes each snapshot generated after the corresponding backup is completed or canceled.

 **Note:** A snapshot-based incremental backup shares the `dumpdates` entries in `dumpdates` database along with the other level-based backups. Therefore, make sure that you do not run snapshot-based backups and regular level-based backups in the same backup paths. For example, make sure that you do not run a level 0 backup and snapshot-based incremental backup in the same backup path or vice versa.

After setting the `BACKUP_MODE` environment variable, snapshot-based incremental backup works with certain data management applications (DMAs) as listed in the next table.

Table 2 DMA support for snapshot-based incremental backups

DMA	DMA-integrated
Symantec NetBackup	Enabled only through an environment variable.
Networker	Enabled only through an environment variable.
Avamar	Yes
CommVault Simpana	Enabled only through a cluster-based environment variable.
Tivoli Storage Manager	Enabled only through a cluster-based environment variable.
Symantec Backup Exec	Enabled only through a cluster-based environment variable.
NetVault	Enabled only through a cluster-based environment variable.
ASG-Time Navigator	Enabled only through a cluster-based environment variable.

NDMP backup and restore of SmartLink files

You can perform NDMP backup and restore operations on data that has been archived to the cloud.

Backup and restore capabilities with CloudPools data include:

- Archive SmartLink files when backing up from a cluster
- Restore data, including SmartLink files, to the same cluster
- Restore data, including SmartLink files, to another cluster
- Back up version information with each SmartLink file and restore the Smartlink file after verifying the version compatibility on the target cluster.

Note: SmartLink files backed up with OneFS 8.2.0 and later releases cannot be restored to releases earlier than 8.2.0.

You specify how files are backed up and restored by setting the NDMP environment variables `BACKUP_OPTIONS` and `RESTORE_OPTIONS`. See [Administering NDMP](#) for details about configuring the backup settings and managing NDMP environment variables.

Note: DeepCopy and ComboCopy backups recall file data from the cloud. The data is not stored on disks. Recall of file data may incur charges from cloud vendors.

With NDMP backup, by default, CloudPools supports backup of SmartLink files that contain cloud metadata such as location of the object. Other details such as version information, account information, local cache state, and unsynchronized cache data associated with the SmartLink file are also backed up.

To prevent data loss when recovering SmartLink files with incompatible versions, you can use the NDMP combo copy backup option to back up SmartLink files with full data. Full data includes metadata and user data. You can use the NDMP combo copy option by setting the `BACKUP_OPTIONS` environment variable.

When performing an NDMP restore operation on SmartLink files backed up using the combo copy option, you can use one of combo copy, shallow copy, or deep copy restore options to recover

SmartLink files. You can specify these options by setting appropriate values to the `RESTORE_OPTIONS` environment variable:

- The combo copy restore option restores SmartLink files from the backup stream only if their version is compatible with the OneFS version on the target cluster. If the SmartLink file version is incompatible with the OneFS version on the target cluster, a regular file is restored.
- The shallow copy restore operation restores the backed up SmartLink file as a SmartLink file on the target cluster if the version check operation on the target cluster is successful.
- The deep copy restore operation forces the recovery of the SmartLink files as regular files on the target cluster. If the version check operation on the target cluster fails.
- If you do not specify any restore operation, NDMP restores SmartLink files using the combo copy restore operation by default.
- When you specify multiple restore options, the combo copy restore operation has the highest priority followed by the shallow copy restore operation. The deep copy restore operation has the lowest priority.

In CloudPools settings, you can set three retention periods that affect backed up SmartLink files and their associated cloud data:

- Full Backup Retention Period for NDMP takes effect when the SmartLink file is backed up as part of a full backup. The default is five years.
- Incremental Backup Retention Period for Incremental NDMP Backup and SyncIQ takes effect when a SmartLink file is backed up as part of an incremental backup. The default is five years.
- Cloud Data Retention Period defines the duration that data in the cloud is kept when its related SmartLink file is deleted. The default is one week.

CloudPools ensures the validity of a backed-up SmartLink file within the cloud data retention period. It is important for you to set the retention periods appropriately to ensure that when the SmartLink file is restored from tape, it remains valid. CloudPools disallows restoring invalid SmartLink files.

To check whether a backed-up SmartLink file is still valid, CloudPools checks the retention periods stored on tape for the file. If the retention time is past the restore time, CloudPools prevents NDMP from restoring the SmartLink file.

CloudPools also makes sure that the account under which the SmartLink files were originally created has not been deleted. If it has, both NDMP backup and restore of SmartLink files will fail.

NDMP protocol support

You can back up the Isilon 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 1-9) NDMP backups and Incremental Forever (level 10)
 - ① **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
- Dump 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)
- Backup and restore of HDFS attributes

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.

Note: All supported DMAs can connect to an Isilon cluster through the IPv4 protocol. However, only some of the DMAs support the IPv6 protocol for connecting to an Isilon cluster.

NDMP hardware support

OneFS can back up data to and recover data from tape devices and virtual tape libraries (VTLs).

Supported tape devices

For NDMP three-way backups, the data management application (DMA) determines the tape devices that are supported.

Supported tape libraries

For both the two-way and three-way NDMP backups, OneFS supports all of the tape libraries that are supported by the DMA.

Supported virtual tape libraries

For three-way NDMP backups, the DMA determines the virtual tape libraries that will be supported.

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.

Supported DMAs	Tested configurations
<ul style="list-style-type: none"> • NetWorker 8.0 and later • Symantec NetBackup 7.5 and later 	<ul style="list-style-type: none"> • Isilon Backup Accelerator node or Fibre Attached Storage node with a second Backup Accelerator node or Fibre Attached Storage node.

Supported DMAs	Tested configurations
	<ul style="list-style-type: none"> Isilon Backup Accelerator node or Fibre Attached Storage node with a NetApp storage system

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.

NDMP backup limitations

NDMP backups have the following limitations.

- Supports block sizes up to 512 KB.
- Does not support more than 4 KB file path length.
- Does not back up file system configuration data, such as file protection level policies and quotas.
- Does not support recovering data from a file system other than OneFS. However, you can migrate data through the NDMP protocol from a NetApp or Unity storage system to OneFS through the `isi_vol_copy` tools. For more information on these tools, see the *OneFS Built-In Migration Tools Guide*.
- Fibre Attached Storage nodes cannot interact with more than 4096 tape paths.
- The maximum length of the FILESYSTEM environment variable supported for a backup operation is 1024.

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).
- Run a maximum of eight NDMP concurrent sessions per Fibre Attached Storage node and four NDMP concurrent sessions per Isilon IQ Backup Accelerator node to obtain optimal throughput per session.
- NDMP backups result in very high Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs). You can reduce your RPO and RTO by attaching one or more Fibre Attached Storage nodes to the cluster and then running two-way NDMP backups.
- The throughput for an Isilon cluster during the backup and recovery operations is dependent on the dataset and is considerably reduced for small files.
- If you are backing up large numbers of small files, set up a separate schedule for each directory.
- If you are performing NDMP three-way backups, run multiple NDMP sessions on multiple nodes in your Isilon cluster.
- Recover files through Directory DAR (DDAR) if you recover large numbers of files frequently.
- Use the largest tape record size available for your version of OneFS to increase throughput.
- 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.

SmartConnect recommendations

- A two-way NDMP backup session with SmartConnect requires Fibre Attached Storage node for backup and recovery operations. However, a three-way NDMP session with SmartConnect does not require Fibre Attached Storage nodes for these operations.
- For a NDMP two-way backup session with SmartConnect, connect to the NDMP session through a dedicated SmartConnect zone consisting of a pool of Network Interface Cards (NICs) on the Fibre Attached Storage nodes.
- For a two-way NDMP backup session without SmartConnect, initiate the backup session through a static IP address or fully qualified domain name of the Fibre Attached Storage node.
- For a three-way NDMP backup operation, the front-end Ethernet network or the interfaces of the nodes are used to serve the backup traffic. Therefore, it is recommended that you configure a DMA to initiate an NDMP session only using the nodes that are not already overburdened serving other workloads or connections.
- For a three-way NDMP backup operation with or without SmartConnect, initiate the backup session using the IP addresses of the nodes that are identified for running the NDMP sessions.

Fibre Attached Storage recommendations

- Assign static IP addresses to Fibre Attached Storage nodes.
- Attach more Fibre Attached Storage nodes to larger clusters. The recommended number of Fibre Attached Storage nodes is listed in the following table.

Table 3 Nodes per Fibre Attached Storage node

Node type	Recommended number of nodes per Fibre Attached Storage node
X-Series	3
NL-Series	3
S-Series	3
HD-Series	3

- Attach more Fibre Attached Storage nodes if you are backing up to more tape devices.

DMA-specific recommendations

- Enable parallelism for the DMA if the DMA supports this option. This allows 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. The examples given in the table are valid only if the backup path is `/ifs/data`.

Table 4 NDMP file and directory matching wildcards

Character	Description	Example	Includes or excludes the following directories
*	Takes the place of any character or characters	archive*	archive1 src/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
//	Takes the place of a single slash (/)	ifs//data//archive	/ifs/data/archive
***	Takes the place of a single asterisk (*)		
..	Ignores the pattern if it is at the beginning of a path	../home/john	home/john

Note: " " are required for Symantec NetBackup when multiple patterns are specified. The patterns are not limited to directories.

Unanchored patterns such as `home` or `user1` target a string of text that might belong to many files or directories. If a pattern contains `/`, it is an anchored pattern. An anchored pattern is always matched from the beginning of a path. A pattern in the middle of a path is not matched. Anchored patterns target specific file pathnames, such as `ifs/data/home`. You can include or exclude either types of patterns.

If you specify both the include and exclude patterns, the include pattern is first processed followed by the exclude pattern.

If you specify both the 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

Administering NDMP

This chapter contains the following topics:

• NDMP backup and recovery tasks.....	72
• Configuring basic NDMP backup settings.....	72
• Managing NDMP user accounts.....	73
• Managing NDMP environment variables.....	74
• Managing NDMP contexts.....	83
• Managing NDMP sessions.....	86
• Managing NDMP Fibre Channel ports.....	89
• Managing NDMP preferred IP settings.....	91
• Managing NDMP backup devices.....	92
• NDMP dumpdates file overview.....	95
• Managing snapshot based incremental backups.....	96
• NDMP restore operations.....	97
• Managing file list backups.....	97
• Managing cluster performance for NDMP sessions.....	99
• Managing CPU usage for NDMP sessions.....	101
• Configuring NDMP backups with NetWorker.....	102
• Configuring NDMP backup with Symantec NetBackup.....	107
• Configuring NDMP backup with CommVault Simpana.....	110
• Configuring NDMP backup with IBM Tivoli Storage Manager.....	115

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 on the Isilon 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 Isilon 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.

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 > NDMP > NDMP Settings**.
2. In the **Service** area, click **Enable NDMP Service**.
3. In the **Port number** field, specify a port number through which a data management application (DMA) can connect to the Isilon cluster. The default port number is 10000.
4. (Optional) 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.
5. In the **NDMP Administrators** area, click **Add an NDMP Administrator** to add a new administrator.

The **Add NDMP Administrator** dialog appears.
6. Enter an administrator name and password, confirm the password, and click **Add NDMP Administrator**.
7. Click **Save Changes** to save all the settings. Alternatively, click **Revert Changes** to undo the changes and revert back to the previous settings.

Disable NDMP backup

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

Procedure

1. Click **Data Protection > NDMP > NDMP Settings**.
2. In the **Service** area, clear the **Enable NDMP service** check box to disable NDMP backup.

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 Isilon cluster, and the DMA vendor that OneFS is configured to interact with.

Procedure

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

Managing NDMP user accounts

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

Create an NDMP administrator account


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

Procedure

1. Click **Data Protection > NDMP > NDMP Settings**.
2. In the **NDMP Administrators** area, click **Add an NDMP Administrator**.

The **Add NDMP Administrator** dialog appears.

3. In the **Add NDMP Administrator** dialog box, in the **Name** field, type a name for the account.

 **Note:** The NDMP administrator that you create in this step is applicable only for NDMP operations. You cannot link this NDMP administrator to any other user, group, or identity on the cluster.

4. In the **Password** and **Confirm password** fields, type the password for the account.

 **Note:** There are no special password policy requirements for an NDMP administrator.

5. Click **Add NDMP Administrator**.

View NDMP user accounts

You can view information about NDMP user accounts.

Procedure

1. Click **Data Protection > NDMP > NDMP Settings**.

2. In the **NDMP Administrators** area, review information about an NDMP administrator by selecting the check box corresponding to an administrator and clicking **View/Edit**.

Modify the password of an NDMP administrator account

You can modify the password of an NDMP administrator account.

Procedure

1. Click **Data Protection > NDMP > NDMP Settings**.
2. In the **NDMP Administrators** area, select the check box next to the desired administrator name and click **View/Edit**.

The **View NDMP Administrator Details** dialog box appears.

3. Click **Edit**.

The **Edit NDMP Administrator Details** dialog box appears.

4. Type a new password, confirm the password, and then click **Save Changes**.

Delete an NDMP administrator account

You can delete an NDMP administrator account.

Procedure

1. Click **Data Protection > NDMP > NDMP Settings**.
2. In the **NDMP Administrators** area, select the check box next to the desired administrator name and click **Delete**.

The administrator name is removed from the list of NDMP administrators.

Managing NDMP environment variables

In OneFS, you can manage NDMP backup and recovery operations by specifying default NDMP environment variables. You can also override default NDMP environment variables through your data management application (DMA).

You can add, view, edit, and delete environment variables. The environment variables can be managed on a per-backup-path basis. They are appended to the environment variables passed from a DMA in a backup or recovery session.

The following table lists the DMAs that allow you to directly set environment variables:

Table 5 DMA support for environment variable setting

DMA	Supported directly on the DMA	Supported through OneFS command-line interface
Symantec NetBackup	Yes	Yes
Networker	Yes	Yes
Avamar	No	Yes
CommVault Simpana	No	Yes
IBM Tivoli Storage Manager	No	Yes
Symantec Backup Exec	No	Yes

Table 5 DMA support for environment variable setting (continued)

DMA	Supported directly on the DMA	Supported through OneFS command-line interface
Dell NetVault	No	Yes
ASG-Time Navigator	No	Yes

In case you cannot set an environment variable directly on a DMA for your NDMP backup or recovery operation, log in to an Isilon cluster through an SSH client and set the environment variable on the cluster through the `isi ndmp settings variables create` command.

NDMP environment variable settings

You can view the NDMP environment variable settings and manage them as necessary.

The following settings appear in the **Variables** table:

Setting	Description
Add Variables	Add new path environment variables along with their values.
Path	The path under the <code>/ifs</code> directory to store new environment variables. If Path is set to <code>"/BACKUP"</code> , the environment variable is applied to all the backup operations. If Path is set to <code>"/RESTORE"</code> , the environment variable is applied to all the restore operations.
Add Name/Value	Add a name and value for the new environment variable.
Name	Name of the environment variable.
Value	Value set for the environment variable
Action	Edit, view, or delete an environment variable at a specified path.

Add an NDMP environment variable

You can add environment variables at a specified path that can be applied per-backup-path or globally.

Procedure

1. Click **Data Protection > NDMP > Environment Settings**.
2. Click **Add Variables** to open the **Add Path Variables** dialog box.
3. In the **Variable Settings** area, specify the following parameters:
 - a. Specify or browse to a path under `/ifs` to store the environment variable.

Note:

- To set a global environment variable for backup and recovery operations, specify the `/BACKUP` path for a backup operation and the `/RESTORE` path for a recovery operation.

- The backup path must include `.snapshot/<snapshot name>` when running a backup of a user-created snapshot.

- b. Click **Add Name/Value**, specify an environment variable name and value, and then click **Create Variable**.

View NDMP environment variables

You can view details about the NDMP environment variables

Procedure

1. Click **Data Protection > NDMP > Environment Settings**.
2. In the **Variables** table, click the check box corresponding to an environment variable and then click **View/Edit**.
3. In the **Display Path Variables** dialog box, review the details.

Edit an NDMP environment variable

You can edit an NDMP environment variable.

Procedure

1. Click **Data Protection > NDMP > Environment Settings**.
2. In the **Variables** table, click the check box corresponding to an environment variable and then click **View/Edit**.
3. In the **Display Path Variables** dialog box, click **Edit Path Variables**.
4. In the **Edit Variables** dialog box, click **Add Name/Value** and specify a new name and value for the environment variable.

Delete an NDMP environment variable

You can delete an NDMP environment variable.

Procedure

1. Click **Data Protection > NDMP > Environment Settings**.
2. In the **Variables** table, click the check box corresponding to an environment variable and then click **Delete**.
3. In the confirmation dialog box, click **Delete**.

Results

You can also delete an NDMP environment variable through the **Edit Variables** dialog box that appears when you click **View/Edit** and then click **Edit Path Variables**.

NDMP environment variables

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

Symantec NetBackup and NetWorker are the only two DMAs that allow you to directly set environment variables and propagate them to OneFS.

Table 6 NDMP environment variables

Environment variable	Valid values	Default	Description
BACKUP_FILE_LIST	<file-path>	None	Triggers a file list backup. Currently, only Networker and Symantec NetBackup can pass environment variables to OneFS.
BACKUP_MODE	TIMESTAMP SNAPSHOT	TIMESTAM P	Enables or disables snapshot-based incremental backups. To enable snapshot-based incremental backups, specify SNAPSHOT.
BACKUP_OPTIONS	0x00000400 0x00000200 0x00000100 0x00000001 0x00000002 0x00000004	0	<p>This environment variable controls the behavior of the backup operations. The following settings are applicable only to datasets containing the CloudPools-driven SmartLink files:</p> <p>0x00000400 Backs up SmartLink files with full data. This is the combo copy backup option.</p> <p>0x00000200 Backs up all the cache data. This is the shallow copy backup option.</p> <p>0x00000100 Reads SmartLink file data from the cloud and backs up the SmartLink files as regular files. This is the deep copy option.</p> <p>0x00000001 Always adds DUMP_DATE into the list of environment variables at the end of a backup operation. The DUMP_DATE value is the time when the backup snapshot was taken. A DMA can</p>

Table 6 NDMP environment variables (continued)

Environment variable	Valid values	Default	Description
			<p>use the DUMP_DATE value to set BASE_DATE for the next backup operation.</p> <p>0x00000002</p> <p>Retains the backup snapshot of a token-based backup in the <code>dumpdates</code> file. Since a token-based backup has no LEVEL, its level is set to 10 by default. The snapshot allows a faster-incremental backup as the next incremental backup after the token-based backup is done.</p> <p>0x00000004</p> <p>Retains the previous snapshot. After a faster-incremental backup, the prior snapshot is saved at level 10. In order to avoid two snapshots at the same level, the prior snapshot is kept at a lower level in the <code>dumpdates</code> file. This allows the BASE_DATE and BACKUP_MODE=snaps hot settings to trigger a faster-incremental backup instead of a token-based backup. The environment variable settings prompt the NDMP server to compare the BASE_DATE value against the timestamp in the <code>dumpdates</code> file to find the prior backup. Even though the DMA fails the latest faster-</p>

Table 6 NDMP environment variables (continued)


Environment variable	Valid values	Default	Description
			incremental backup, OneFS retains the prior snapshot. The DMA can then retry the faster-incremental backup in the next backup cycle using the <code>BASE_DATE</code> value of the prior backup.
<code>BASE_DATE</code>			Enables a token-based incremental backup. The <code>dumpdates</code> file will not be updated in this case.
<code>DIRECT</code>	Y N	N	Enables or disables Direct Access Restore (DAR) and Directory DAR (DDAR). The following values are valid: Y Enables DAR and DDAR. N Disables DAR and DDAR.
<code>EXCLUDE</code>	<i><file-matching-pattern></i>	None	If you specify this option, OneFS does not back up files and directories that meet the specified pattern. Separate multiple patterns with a space.
<code>FILES</code>	<i><file-matching-pattern></i>	None	If you specify this option, OneFS backs up only files and directories that meet the specified pattern. Separate multiple patterns with a space.  Note: As a rule, files are matched first and then the <code>EXCLUDE</code> pattern is applied.
<code>HIST</code>	<i><file-history-format></i>	Y	Specifies the file history format. The following values are valid:

Table 6 NDMP environment variables (continued)

Environment variable	Valid values	Default	Description
			<p>D Specifies directory or node file history.</p> <p>F Specifies path-based file history.</p> <p>Y Specifies the default file history format determined by your NDMP backup settings.</p> <p>N Disables file history.</p>
LEVEL	<integer>	0	<p>Specifies the level of NDMP backup to perform. The following values are valid:</p> <p>0 Performs a full NDMP backup.</p> <p>1 - 9 Performs an incremental backup at the specified level.</p> <p>10 Performs Incremental Forever backups.</p>
MSB_RETENTION_PERIOD	Integer	300 sec	For a multi-stream backup session, specifies the backup context retention period.
MSR_RETENTION_PERIOD	0 through 60*60*24	600 sec	For a multi-stream restore session, specifies the recovery context retention period within which a recovery session can be retried.
RECURSIVE	Y N	Y	For restore sessions only. Specifies that the restore session should recover files or sub-directories under a directory automatically.

Table 6 NDMP environment variables (continued)


Environment variable	Valid values	Default	Description
RESTORE_BIRTHTIME	Y N	N	Specifies whether to recover the birth time for a recovery session.
RESTORE_HARDLINK _BY_TABLE	Y N	N	<p>For a single-threaded restore session, determines whether OneFS recovers hard links by building a hard-link table during recovery operations. Specify this option if hard links are incorrectly backed up and recovery operations are failing.</p> <p>If a recovery 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> <p> Note: This variable is not effective for a parallel restore operation.</p>
RESTORE_OPTIONS	0x00000001 0x00000002 0x00000004 0x00000100 0x00000200	0	<p>This environment variable controls the behavior of the restore operations.</p> <p>0x00000001 Performs a single-threaded restore operation.</p> <p>0x00000002 Restores attributes to the existing directories.</p> <p>0x00000004 Creates intermediate directories with default attributes. The default behavior is to get attributes from the first object under a given directory.</p> <p>The following settings are applicable only to datasets</p>

Table 6 NDMP environment variables (continued)

Environment variable	Valid values	Default	Description
			backed up with the combo copy backup option: 0x00000100 Forces deep copy restoration of the SmartLink files. That is, restores the backed up SmartLink file as a regular file on the target cluster. 0x00000200 Forces shallow copy restoration of the SmartLink files. That is, restores the backed up SmartLink file as a SmartLink file on the target cluster.
UPDATE	Y N	Y	Determines whether OneFS updates the <code>dumpdates</code> file. The default is to perform a combo copy restore. Y OneFS updates the <code>dumpdates</code> file. N OneFS does not update the <code>dumpdates</code> file.

Setting environment variables for backup and restore operations

You can set environment variables to support the backup and restore operations for your NDMP session.

You can set environment variables through a data management application (DMA) or the command-line interface. Alternatively, you can set global environment variables. The precedence to apply their settings for a backup or restore operation follows:


- The environment variables specified through a DMA have the highest precedence.
- Path-specific environment variables specified by the `isi ndmp settings variables` take the next precedence.
- Global environment variable settings of `"/BACKUP"` or `"/RESTORE"` take the lowest precedence.

You can set environment variables to support different types of backup operations as described in the following scenarios:

- If the `BASE_DATE` environment variable is set to any value and if you set the `BACKUP_MODE` environment variable to `SNAPSHOT`, the `LEVEL` environment variable is automatically set to 10 and an Incremental Forever backup is performed.
- If the `BASE_DATE` environment variable is set to 0, a full backup is performed.
- If the `BACKUP_MODE` environment variable is set to `snapshot` and the `BASE_DATE` environment variable is not set to 0, the entries in the `dumpdates` file are read and compared with the `BASE_DATE` environment variable. If an entry is found and a prior valid snapshot is found, a faster incremental backup is performed.
- If the `BACKUP_MODE` environment variable is set to `snapshot`, the `BASE_DATE` environment variable is not set to 0, and if no entries are found in the `dumpdates` file and no prior valid snapshots are found, a token-based backup is performed using the value of the `BASE_DATE` environment variable.
- If the `BASE_DATE` environment variable is set, the `BACKUP_OPTIONS` environment variable is set to `0x0001` by default.
- If the `BACKUP_MODE` environment variable is set to `snapshot`, the `BACKUP_OPTIONS` environment variable is set to `0x0002` by default.
- If the `BACKUP_OPTIONS` environment variable is set to `0x0004`, the snapshot is saved and maintained by the application used for the backup process.
- In order to run an Incremental Forever backup with faster incremental backups, you must set the following environment variables:
 - `BASE_DATE=<time>`
 - `BACKUP_MODE=snapshot`
 - `BACKUP_OPTIONS=7`

Managing NDMP contexts

Each NDMP backup, restore, restartable backup, and multi-stream backup process creates a context. The NDMP server stores the corresponding working files in the context. You can view or delete a context.


 **Note:** If you delete a restartable backup context, you cannot restart the corresponding backup session.

Managing NDMP restartable backups

An NDMP restartable backup also known as backup restartable extension (BRE) is a type of 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. There can be only one restartable backup context per restartable backup session. When your DMA restarts a failed backup session, the working files are reverted to the state corresponding to the NDMP restart backup request.

Restartable backups are supported for NetWorker 8.1 and later versions and CommVault Simpana DMAs.

 **Note:** NDMP multi-stream backup does not support restartable backups.

NDMP context settings

You can view the details of NDMP contexts and manage those contexts.

The following settings appear in the **Contexts** table:

Setting	Description
Type	The context type. It can be one of backup, restartable backup, or restore.
ID	An identifier for a backup or restore job. A backup or restore job consists of one or more streams all of which are identified by this identifier. This identifier is generated by the NDMP backup daemon.
Start Time	The time when the context started in <i>month date time year</i> format.
Actions	View or delete a selected context.
Status	Status of the context. The status shows up as <i>active</i> if a backup or restore job is initiated and continues to remain active until the backup stream has completed or errored out.
Path	The path where all the working files for the selected context are stored.
MultiStream	Specifies whether the multistream backup process is enabled.
Lead Session ID	The identifier of the first backup or restore session corresponding to a backup or restore operation.
Sessions	A table with a list of all the sessions that are associated with the selected context.

Configure NDMP restartable backup settings

You can specify the number of restartable backup contexts that OneFS can retain at a time, up to a maximum of 1024 contexts. The default number of restartable backup contexts is set to 64.

Procedure

- Run the `isi ndmp settings global modify` command.

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

```
isi ndmp settings global modify --bre_max_num_contexts=128
```

View NDMP contexts

You can view information about the NDMP backup, restartable backup, and recovery contexts.

Procedure

1. Click **Data Protection > NDMP > Contexts**.
2. In the **Contexts** table, click the check box corresponding to a context that you want to review and click **View Details**.
3. Review the information about the context in the **Display Backup Context** dialog box.

Delete an NDMP context

You can delete an NDMP context.

About this task

Backup and restore contexts have retention periods beyond which the contexts are deleted automatically. However, you can choose to delete a context before its retention period to free up resources. You cannot delete contexts with active sessions. Also, you cannot delete backup contexts with active BRE contexts. You can delete BRE contexts only if they are not a part of active sessions.

Procedure

1. Click **Data Protection > NDMP > Contexts**.
2. In the **Contexts** table, select a context and click **Delete**.
3. In the confirmation dialog box, click **Delete**.

Configure NDMP restartable backups for NetWorker

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

About this task

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** page, click the **Checkpoint enabled** checkbox.
 - b. In the **Checkpoint granularity** drop-down list, select **File**.
3. In the **Application information** field, type any NDMP variables that you want to specify.

The following variable setting 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. On the **Monitoring** page, locate the backup process in the **Groups** list.
 - b. Right-click the backup process and then, in the context menu, click **Restart**.

NetWorker automatically restarts the backup from the last checkpoint.


Managing NDMP sessions

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

NDMP session information

Data management applications (DMAs) establish sessions with the NDMP daemon running on the Fibre Attached Storage node. The communication from the DMA with the NDMP daemon is managed under the context of a session.

The following items appear in the **Sessions** table:

Item	Description
Session	Specifies the unique identification number that OneFS assigns to the session.
Elapsed	Specifies the time that has elapsed since the session started.
Transferred	Specifies the amount of data that was transferred during the session.
Throughput	Specifies the average throughput of the session over the past five minutes.
Client/Remote	Specifies the IP address of the backup server that the data management application (DMA) is running on. If a NDMP three-way backup or restore operation is currently running, the IP address of the remote tape media server also appears.
Mover/Data	<p>Specifies 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. 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> Note: When a session ID instead of a state appears, the session is automatically redirected.</p> <p>The following states might appear:</p> <p>Active</p> <p>The data mover or data server is currently sending or receiving data.</p> <p>Paused</p> <p>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</p> <p>The data mover or data server is not sending or receiving data.</p>

Item	Description
	<p>Listen</p> <p>The data mover or data server is waiting to connect to the data server or data mover.</p>
Operation	<p>Specifies the type of operation (backup or restore) that is currently in progress. If no operation is in progress, this field is blank.</p> <p>B ({M} {F} [L[0-10] T0 Ti S[0-10]] {r R})+</p> <p>Where:</p> <p>[a]—a is required</p> <p>{ a }—a is optional</p> <p>a b— a or b but not at the same time</p> <p>A(B)—The session is an agent session for a redirected backup operation</p> <p>M—Multi-stream backup</p> <p>F—File list</p> <p>L—Level-based</p> <p>T—Token-based</p> <p>S—Snapshot mode</p> <p>s—Snapshot mode and a full backup (when root dir is new)</p> <p>r—Restartable backup</p> <p>R—Restarted backup</p> <p>+—Backup is running with multiple state threads for better performance</p> <p>0-10—Dump Level</p> <p>R ({M s}[F D S]{h})</p> <p>Where:</p> <p>A(B)—The session is an agent session for a redirected restore operation</p> <p>M—Multi-stream restore</p> <p>s—Single-threaded restore (when RESTORE_OPTIONS=1)</p> <p>F—Full restore</p> <p>D—DAR</p> <p>S—Selective restore</p> <p>h—Restore hardlinks by table</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</p>

Item	Description
	are restoring data to the same location that you backed up your data from, the same path appears twice.
Device	Specifies the name of the tape or media changer device that is communicating with the Isilon cluster.
Mode	Specifies how OneFS is interacting with data on the backup media server through the following options: Read/Write OneFS is reading and writing data during a backup operation. Read OneFS is reading data during a restore operation. Raw The DMA has access to tape drives, but the drives do not contain writable tape media.
Actions	Allows you to probe or delete a session.

Example 1 NDMP backup and restore operations

Examples of active NDMP backup sessions indicated through the **Operation** field that is described in the previous table are as shown:

```

B(T0): Token based full backup
B(Ti): Token based incremental backup
B(L0): Level based full backup
B(L5): Level 5 incremental backup
B(S0): Snapshot based full backup
B(S3): Snapshot based level 3 backup
B(FT0): Token based full filelist backup
B(FL4): Level 4 incremental filelist backup
B(L0r): Restartable level based full backup
B(S4r): Restartable snapshot based level 4 incremental backup
B(L7R): Restarted level 7 backup
B(FT1R): Restarted token based incremental filelist backup
B(ML0): Multi-stream full backup

```

Examples of active NDMP restore sessions indicated through the **Operation** field that is described in the previous table are as shown:

```

R(F): Full restore
R(D): DAR
R(S): Selective restore
R(MF): Multi-stream full restore
R(sFh): single threaded full restore with restore hardlinks by
table option

```


View NDMP sessions

You can view information about active NDMP sessions.

Procedure

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

Abort an NDMP session

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

Procedure

1. Click **Data Protection > NDMP > Sessions**.
2. In the **Sessions** table, click the check box corresponding to the session you want to abort, and click **Delete**.
3. In the confirmation dialog box, click **Delete**.

Managing NDMP Fibre Channel ports

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

NDMP backup port settings

OneFS assigns default settings to each Fibre Channel port on the Fibre Attached Storage node attached to the Isilon cluster. These settings identify the port and determine how the port interacts with the NDMP backup devices.

The following settings appear in the **Ports** table:

Setting	Description
LNN	Specifies the logical node number of the Fibre Attached Storage node.
Port	Specifies the name and port number of the Fibre Attached Storage node.
Topology	<p>Specifies the type of Fibre Channel topology that is supported by the port. 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 for a switched-fabric topology.</p>

Setting	Description
WWNN	Specifies the world wide node name (WWNN) of the port. This name is the same for each port on a given node.
WWPN	Specifies the world wide port name (WWPN) of the port. This name is unique to the port.
Rate	Specifies 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, the Fibre Channel chip negotiates with connected Fibre Channel switch or Fibre Channel devices to determine the rate. Auto is the recommended setting.
State	Specifies whether a port is enabled or disabled.
Actions	Allows you to view and edit the port settings.

Enable or disable an NDMP backup port

You can enable or disable an NDMP backup port.

Procedure

1. Click **Data Protection > NDMP > Ports**.
2. In the row of a port, click **View/Edit**.
The **View Port** dialog box appears.
3. Click the **Edit Port** button.
The **Edit Port** dialog box appears.
4. From the **State** drop-down list, select `Enable` or `Disable`.
5. Click the **Save Changes** button.

View NDMP backup ports

You can view information about Fibre Channel ports of Fibre Attached Storage nodes attached to an Isilon cluster.

Procedure

1. Click **Data Protection > NDMP > Ports**.
2. In the **Ports** table, review information about NDMP backup ports. For more detailed information about a specific port, click the **View/Edit** button corresponding to that port.

Modify NDMP backup port settings

You can modify the settings of an NDMP backup port.

Procedure

1. Click **Data Protection > NDMP > Ports**.
2. Click the **View/Edit** button corresponding to the port you want to modify.
The **View Port** dialog box appears.
3. Click the **Edit Port** button.
The **Edit Port** dialog box appears.

4. Edit the settings in the **Edit Port** dialog box, and click **Save Changes** when finished.

Managing NDMP preferred IP settings

If you are performing NDMP three-way operations using Avamar in an environment with multiple network interfaces, you can create, modify, delete, list, and view cluster-wide or subnet-specific NDMP preferred IP settings.

You can manage NDMP preferred IP settings only through the OneFS command-line interface.

Create an NDMP preferred IP setting

If you are performing an NDMP three-way backup or restore operation using Avamar, you can create a cluster-wide or a subnet-specific NDMP preferred IP setting.

Procedure

- Create an NDMP preferred IP setting by running the `isi ndmp settings preferred-ips create` command.

For example, run the following command to apply a preferred IP setting for a cluster:

```
isi ndmp settings preferred-ips create cluster
groupnet0.subnet0,10gnet.subnet0
```

Run the command as shown in the following example to apply a preferred IP setting for a subnet group:

```
isi ndmp settings preferred-ips create 10gnet.subnet0
10gnet.subnet0,groupnet0.subnet0
```

Modify an NDMP preferred IP setting

If you are performing an NDMP three-way backup or restore operation using Avamar, you can modify an NDMP preferred IP setting by adding or deleting a subnet group.

Procedure

- Modify an NDMP preferred IP setting by running the `isi ndmp settings preferred-ips modify` command.

For example, run the following commands to modify the NDMP preferred IP setting for a cluster:

```
isi ndmp settings preferred-ips modify 10gnet.subnet0 --add-data-subnets
10gnet.subnet0,groupnet0.subnet0
```

Run the command as shown in the following example to modify the NDMP preferred IP setting for a subnet:

```
isi ndmp settings preferred-ips modify 10gnet.subnet0 --remove-data-
subnets groupnet0.subnet0
```

List NDMP preferred IP settings

If you are performing an NDMP three-way backup or restore operation using Avamar, you can list all the NDMP preferred IP settings.

Procedure

- List the NDMP preferred IP settings by running the `isi ndmp settings preferred-ips list` command.

For example, run the following command to list the NDMP preferred IP settings:

```
isi ndmp settings preferred-ips list
```

View NDMP preferred IP settings

If you are performing an NDMP three-way backup or restore operation using Avamar, you can view the NDMP preferred IP settings for a subnet or cluster.

Procedure

- View an NDMP preferred IP setting by running the `isi ndmp settings preferred-ips view` command.

For example, run the following command to view the NDMP preferred IP setting for a subnet:

```
isi ndmp settings preferred-ips view --scope=10gnet.subnet0
```

Delete NDMP preferred IP settings

If you are performing an NDMP three-way backup or restore operation using Avamar, you can delete an NDMP preferred IP setting for a subnet or cluster.

Procedure

- Delete NDMP preferred IP settings by running the `isi ndmp settings preferred-ips delete` command.

For example, run the following command to delete the preferred IP setting for a subnet:

```
isi ndmp settings preferred-ips delete --scope=10gnet.subnet0
```

Managing NDMP backup devices

After you attach a tape or media changer device to a Fibre Attached Storage 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.

In case the virtual tape library (VTL) device has multiple LUNs, you must configure LUN0 so that all the LUNs are detected properly.

NDMP backup device settings

OneFS creates a device entry for each device you attach to the cluster through a Fibre Attached Storage node.

The following table describes the settings you can review for a tape or media changer device in the **Devices** table and also through the **View Tape Devices** dialog box that appears when you select a device and click **View/Edit**:

Setting	Description
Name	Specifies 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, <code>Read/Write</code> appears. If the device is not in use, <code>Closed</code> appears.
WWNN	Specifies the world wide node name of the device.
Product (Vendor/Model/Revision)	Specifies the name of the device vendor and the model name or number of the device.
Serial Number	Specifies the serial number of the device.
Actions	Allows you to view, edit, or delete a device.
Path	Specifies the name of the Fibre Attached Storage node that is attached to the device and the port numbers to which the device is connected.
LUN	Specifies the logical unit number (LUN) of the device.
State	Specifies whether the device is active or inactive.
WWPN	Specifies the world wide port name (WWPN) of the port on the tape or media changer device.
Port ID	Specifies the port ID of the device that binds the logical device to the physical device.
Open Count	A counter of the active and open connections to the device.
Device Name	Specifies the regular device name that appears under the FreeBSD operating system.
Pass Name	Specifies the pass-thru device name that appears under the FreeBSD operating system.

Detect NDMP backup devices

If you connect a tape device or media changer to a Fibre Attached Storage 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 Isilon node, a specific port, or all ports on all nodes.

Procedure

1. Click **Data Protection > NDMP > Devices**.
2. Click the **Discover Devices** link.
The **Discover Devices** dialog appears.
3. (Optional) To scan only a specific node for NDMP devices, from the **Node** 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 > NDMP > Devices**.
2. In the **Tape Devices** and **Media Changer Devices** tables, review the 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 > NDMP > Devices**.
2. In the **Tape Devices** table, or the **Media Changer Devices** table, click the check box corresponding to the name of a backup device entry.
3. Click **View/Edit**.

The **View Tape Devices** or **View Media Changers** dialog box appears.

4. Click **Edit Tape Device**.

The **Edit Tape Devices** or **Edit Media Changers** dialog box appears.

5. Edit the device name.
6. Click **Save Changes**.

Delete an entry for an NDMP backup device

If you physically remove an NDMP device from an Isilon 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.

About this task

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 > NDMP > Devices**.

2. In the **Tape Devices** table or the **Media Changer Devices** table, click the check box corresponding to the device that you want to remove.
3. Click **Delete**.
4. In the **Confirm Delete** dialog box, click **Delete**.

NDMP dumpdates file overview

When you set the `UPDATE` environment variable to `Y`, the NDMP daemon maintains a `dumpdates` file to record all but the token-based backup sessions. The timestamp within the `dumpdates` file helps identify the changed files for the next level-based backup. The entries within the `dumpdates` file also provide information about the last backup session at a given path and the type of backup session which can be a full, level-based incremental, or snapshot-based backup. This information determines the type of incremental backup you must run subsequently. The entries within the `dumpdates` file may be obsolete when the backup path is removed. In such a case, all the obsolete entries can be removed from the `dumpdates` file.

Managing the NDMP dumpdates file

You can view or delete entries in the NDMP `dumpdates` file.

NDMP dumpdates file settings

You can view details about the entries in the NDMP `dumpdates` file and delete them if required.

The following settings appear in the **Dumpdates** table:

Setting	Description
Date	Specifies the date when an entry was added to the <code>dumpdates</code> file.
ID	The identifier for an entry in the <code>dumpdates</code> file.
Level	Specifies the backup level.
Path	Specifies the path where the <code>dumpdates</code> file is saved.
Snapshot ID	Identifies changed files for the next level of backup. This ID is applicable only for snapshot-based backups. In all the other cases, the value is 0.
Actions	Deletes an entry from the <code>dumpdates</code> file.

View entries in the NDMP dumpdates file

You can view all the entries in the NDMP `dumpdates` file.

Procedure

1. Click **Data Protection > NDMP > Environment Settings**.
2. In the **Dumpdates** table, view information about the entries in the NDMP `dumpdates` file.

Delete entries from the NDMP dumpdates file

You can delete entries from the NDMP `dumpdates` file.

Procedure

1. Click **Data Protection > NDMP > Environment Settings**.
2. In the **Dumpdates** table, click **Delete** against the entry that you want to delete.
3. In the **Confirm Delete** dialog box, click **Delete**.

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

- 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. Click **Data Protection > NDMP > Environment Settings**.
2. In the **Dumpdates** table, view information about the snapshot-based incremental backups.

Delete snapshots for snapshot-based incremental backups

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

About this task

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. Click **Data Protection > NDMP > Environment Settings**.
2. In the **Dumpdates** table, click **Delete** against the entry that you want to delete.
3. In the **Confirm Delete** dialog box, click **Delete**.

NDMP restore operations

NDMP supports the following types of restore operations:

- NDMP parallel restore (multi-threaded process)
- NDMP serial restore (single-threaded process)

NDMP parallel restore operation

Parallel (multi-threaded) restore enables faster full or partial restore operations by writing data to the cluster as fast as the data can be read from the tape. Parallel restore is the default restore mechanism in OneFS.

The restore operation can restore multiple files concurrently through the parallel restore mechanism.

NDMP serial restore operation

For troubleshooting or for other purposes, you can run a serial restore operation which uses fewer system resources. The serial restore operation runs as a single-threaded process and restores one file at a time to the specified path.

Specify a NDMP serial restore operation

You can use the `RESTORE_OPTIONS` environment variable to specify a serial (single-threaded) restore operation.

Procedure

1. In your data management application, configure a restore operation as you normally would.
2. Make sure that the `RESTORE_OPTIONS` environment variable is set to `1` on your data management application.

If the `RESTORE_OPTIONS` environment variable is not already set to `1`, specify the `isi ndmp settings variables modify` command from the OneFS command line. The following command specifies serial restore for the `/ifs/data/projects` directory:

```
isi ndmp settings variables modify /ifs/data/projects RESTORE_OPTIONS 1
```

The value of the `path` option must match the `FILESYSTEM` environment variable that is set during the backup operation. The value that you specify for the `name` option is case sensitive.

3. Start the restore operation.

Managing file list backups

If your data management application (DMA) can pass environment variables to OneFS, you can control backups by specifying a file list.

When you specify a file list backup, only the listed files in a directory that are in a sorted order are backed up.

A backup level other than `0` triggers an incremental file list backup. In an incremental file list backup, only the listed files that were created or changed in the source directory since the last incremental backup of the same level are backed up.

To configure a file list backup, you must complete the following tasks:

- Create the file list in a sorted order and place it in OneFS.
- Specify the path of the source directory.
- Specify the file list location.

The file list is an ASCII text file that lists the pathnames of files to be backed up. The pathnames must be relative to the path specified in the `FILESYSTEM` environment variable. Absolute file paths in the file list are not supported. The pathnames of all files must be included, or they are not backed up.

To specify the full path of the source directory to be backed up, you must specify the `FILESYSTEM` environment variable in your DMA. For example:

```
FILESYSTEM=/ifs/data/projects
```

To specify the pathname of the file list, you must specify the environment variable, `BACKUP_FILE_LIST` in your DMA. The file list must be accessible from the node performing the backup. For example:

```
BACKUP_FILE_LIST=/ifs/data/proj_list.txt
```

Format of a backup file list

You must create a file list to enable a file list backup.

A file list backup requires an ASCII text file in a particular format to identify the pathnames of files to be backed up. Following is an example of a file list with pathnames relative to `/ifs/data/projects`:

```
proj001/plan/\001File
proj001/plan/\002File
proj001/plan/\003File
proj001/plan/\004File
proj001/plan/\005File
proj001/plan/\006File
proj001/plan/\aFile
proj001/plan/\bFile
proj001/plan/\tFile
proj001/plan/\nFile
proj002/plan/\vFile
proj002/plan/\fFile
proj002/plan/\rFile
proj002/plan/\016File
proj002/plan/\017File
proj002/plan/\020File
proj002/plan/\023File
proj002/plan/\024File
proj005/plan/\036File
proj005/plan/\037File
proj005/plan/ File
proj005/plan/!File
proj005/plan/"File
proj005/plan/#File
proj005/plan/$File
proj005/plan/%File
proj005/plan/&File
proj005/plan/'File
```

As shown in the example, the pathnames are relative to the full path of the source directory, which you specify in the `FILESYSTEM` environment variable. Absolute file paths are not supported in the file list.

Also as shown, the directories and files must be in sorted order for the backup to be successful. A `#` at the beginning of a line in the file list indicates to skip the line.

The pathnames of all files must be included in the file list, or they are not backed up. For example, if you only include the pathname of a subdirectory, the subdirectory is backed up, but not the files the subdirectory contains. The exception is ADS (alternate data streams). All ADS associated with a file to be backed up are automatically backed up.

Placement of the file list

Before you can perform a file list backup, you must place the file list in OneFS.

For example, suppose the `FILESYSTEM` environment variable specifies the full path of the directory to be backed up as `/ifs/data/projects`. You can place the text file containing the file list anywhere within the `/ifs` path.

Start a file list backup

You can configure and start a file list backup from your data management application (DMA).

Before you begin

You should have already specified and saved the list of files to be backed up in an ASCII text file.

About this task

Configure a file list backup from your DMA as you would any backup, but with a few additional steps as described in the following procedure.

Procedure

1. Copy the file list to the OneFS file system on the Isilon cluster containing the files to be backed up.

For example, if the directory that you specify in the `FILESYSTEM` environment variable is `/ifs/data/projects`, you could place your file list at `/ifs/data`.

2. In your DMA, specify the `BACKUP_FILE_LIST` environment variable to be the full pathname of the file list that resides on the Isilon cluster.

For example, if the file list was named `proj_backup.txt`, and you placed it at `/ifs/data`, specify `/ifs/data/proj_backup.txt` as the full pathname of the file list.

3. Start your backup as you normally would.

Results

The files in your file list are backed up as specified.

Managing cluster performance for NDMP sessions

NDMP Redirector distributes NDMP loads automatically over nodes with the Sheba network interface card (NIC). You can enable NDMP Redirector to automatically distribute NDMP two-way sessions to nodes with lesser loads. NDMP Redirector checks the CPU usage, number of NDMP operations already running, and the availability of tape devices for the operation on each node before redirecting the NDMP operation. The load-distribution capability results in improved cluster performance when multiple NDMP operations are initiated.

Enable NDMP Redirector to manage cluster performance

You must enable NDMP Redirector in order to automatically distribute NDMP two-way sessions to nodes with lesser loads.

Before you begin

Make sure that the cluster is committed before enabling NDMP Redirector.

Procedure

1. Run the following command through the command line interface to enable NDMP Redirector:

```
isi ndmp settings global modify --enable-redirector true
```

2. View the setting change by running the following command:

```
isi ndmp settings global modify
```

A sample output of the previous command is shown:

```
Service: False
      Port: 10000
      DMA: generic
      Bre Max Num Contexts: 64
      Context Retention Duration: 300
      Smartlink File Open Timeout: 10
      Enable Redirector: True
```

NDMP Redirector limitations

NDMP Redirector has the following limitations.

- NDMP Redirector does not work on NDMP three-way sessions. In other words, redirection can happen only between nodes that share the same Fibre Channel back-end network.
- SCSI 2-only tape drives are not supported for redirection.
- A redirected session runs on the redirected node until the session is closed by a DMA. A redirected session fails if the connection to the DMA is broken or if the DMA changes the session to run an NDMP three-way operation.
- If there is no suitable node for redirection, the node receiving the request will run the requested operation.

Managing CPU usage for NDMP sessions

NDMP Throttler manages the CPU usage during NDMP two-way sessions on 6th Generation nodes. The nodes are then available to adequately support other system activities.

Enable NDMP Throttler

You must enable NDMP Throttler in order to manage CPU usage of NDMP sessions on 6th Generation nodes.

Procedure

1. Run the following command through the command line interface to enable NDMP Throttler:

```
isi ndmp settings global modify --enable-throttler true
```

2. View the setting change by running the following command:

```
isi ndmp settings global modify
```

A sample output of the previous command is shown:

```
Service: False
      Port: 10000
      DMA: generic
      Bre Max Num Contexts: 64
Context Retention Duration: 600
Smartlink File Open Timeout: 10
      Enable Throttler: True
Throttler CPU Threshold: 50
```

3. If required, change the throttler CPU threshold as shown in the following example:

```
isi ndmp settings global modify -throttler-cpu-threshold 80
```

NDMP Throttler limitations

NDMP Throttler has the following limitations.

- Throttler settings are global to all the nodes in a cluster.
- In addition to controlling NDMP two-way sessions, Throttler also controls three-way data server operations.

Configuring NDMP backups with NetWorker

You can configure OneFS and NetWorker to backup data stored on an Isilon cluster. The following procedures explain how to configure NDMP backup with NetWorker 9.0.

Note: The steps described in the procedures are general guidelines only. They might change for different versions of NetWorker. Consult your DMA vendor documentation for the configuration information for a specific version of NetWorker.

Create a group

With 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. On the **Enterprise** tab, select a NetWorker server name and double-click.
3. Click **Launch NetWorker Administration**.
4. In the administration window, click **Protection**.
5. Right-click **Groups** and then click **New**.
6. In the **Name** field, type a name for the group.
7. Click **OK**.

Create a policy

With NetWorker 9.0 and later versions, you can create data protection policies and assign those policies to clients that specify the data to be backed up from an Isilon cluster.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. On the **Enterprise** tab, select a NetWorker server name and double-click.
3. Click **Launch NetWorker Administration**.
4. In the administration window, click **Protection**.
5. Right-click **Policies** and then click **New**.
6. In the **Name** field, type a name for the policy.
7. Click **OK**.

Create a workflow

With NetWorker 9.0 and later versions, you can schedule a backup operation and define workflows related to that operation. You can schedule the time interval to start the backup operation and the type of backup operation to run which can be a level-based or an incremental backup.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. On the **Enterprise** tab, select a NetWorker server name and double-click.
3. In the NetWorker administration window, click **Protection**.
4. In the left navigation pane, expand **Policies** and select the policy that you created.
5. With the policy selected, click **File > New** to start a new workflow.

6. In the **New Workflow** window, type a name for the workflow.
7. In the **Groups** area, click **+**, select the protection group to which the workflow applies using the **Create Group** window, and click **Add**.
8. The **Actions** table displays a list of actions in the workflow. To edit or delete an action in the workflow, select the action and click **Edit** or **Delete**. To create one or more actions for the workflow, click **Add**.
9. Click **OK**.

Scan for tape devices

With 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. On the **Enterprise** tab, select a NetWorker server name and double-click.
3. Click **Launch NetWorker Administration**.
4. In the NetWorker administration window, click **Devices**.
5. Right-click **Libraries**, and then click **Scan for Devices**.
6. Make sure that no existing storage nodes are selected.
7. Click **Create a new Storage Node**.
8. Configure the following settings:

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

9. Click **Start Scan**.

Configure a library

With 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. On the **Enterprise** tab, select a NetWorker server name and double-click.
3. Click **Launch NetWorker Administration**.
4. In the NetWorker administration window, click **Devices**.
5. Right-click **Libraries** and then click **Refresh**.

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

6. Right-click the name of the tape library you want to configure and then click **Configure Library**.

7. In the **Configure Library** window, click **Check All**.
8. Click **Start Configuration**.

Create a data media pool

With 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. On the **Enterprise** tab, select a NetWorker server name and double-click.
3. Click **Launch NetWorker Administration**.
4. In the NetWorker administration window, click **Media**.
5. Click **Media Pools**.
6. In the **Media Pools** area, right-click and then click **New**.
7. Configure the following settings:

Tab	Setting name	Setting value
General	Name	A name for the media pool
	Enabled	Selected
	Pool type	The type of backups that you perform which could be one of Archive , Archive Clone , Backup , or Backup Clone .
Selection Criteria	Devices	Each tape device that you want to use.
Configuration	Max parallelism	The maximum number of tape drives to use for concurrent backups.
Legacy	Levels	Select 1 , 9 , full , or incremental .
	Groups	Specify the group that you created for the Isilon cluster.

Label tape devices

With 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. On the **Enterprise** tab, select a NetWorker server name and double-click.
3. Click **Launch NetWorker Administration**.
4. In the NetWorker administration window, click **Devices**.
5. Click the name of the library that you configured.

6. In the device list, highlight all the tape drives that you want to label.
7. Right-click the highlighted list, and then click **Label**.
8. In the **Label Library Media** window, from the **Target Media Pool** list, select the name of the media pool you created.
9. Make sure that the **Prompt to Overwrite Existing Label** box is cleared.
10. Click **OK**.

Create a client

With NetWorker, you must create a client that specifies the data to be backed up from an Isilon cluster. If you are creating a client using NetWorker 9.0 and later, you must enable the diagnostics mode to set the storage node. Furthermore, to perform a two-way backup, the storage node must be set as the name of the Backup Accelerator node or Fibre Attached Storage node.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. On the **Enterprise** tab, select a NetWorker server name and double-click.
3. Click **Launch NetWorker Administration**.
4. In the NetWorker administration window, click **Protection**.
5. To configure a new client, right-click **Clients** and select **New** to open the **Create Client** window.
6. Configure a client through the **General** tab.
 - a. In the **Name** field, type a name or IP address for the client.
 - b. In the **Save set** field, type the full path of the directory that you want to back up.
 - c. From the **Pool** list, select the name of the data media pool you created.
 - d. From the **Protection group list**, select a group that you created previously.
7. Configure a remote user through the **Apps & Modules** tab.
 - a. In the **Remote user** field, type the name of an NDMP user you configured on the cluster.
 - b. In the **Password** field, type the password of the NDMP user.
 - c. Select **NDMP** and then, in the **Backup command** field, type the following backup command:

```
nsrndmp_save -T tar
```

- d. In the **Application information** field, type the 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.

Back up data through NetWorker

After configuring NetWorker 9.0 for managing NDMP backups on an Isilon cluster, you can use this DMA to back up data stored on an Isilon cluster.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. On the **Enterprise** tab, select a NetWorker server name and double-click.
3. Click **Launch NetWorker Administration**.
4. In the administration window, click **Protection**.
5. In the left pane, expand **Policies** and select the workflow that you have created previously with the backup schedule.
6. Right-click the workflow and click **Start** to start the backup process.

Recover backed up data through NetWorker

You can use NetWorker to recover data that you backed up previously.

About this task

This procedure describes the recovery operation that you can perform through NetWorker 9.0. If you are running a previous version of NetWorker, see the vendor documentation for instructions.

Procedure

1. Connect to the NetWorker server from the NetWorker Management Console Server.
2. On the **Enterprise** tab, select an enterprise—for example, `nw-win-1`—and double-click.
3. Click **Launch NetWorker Administration**.
4. In the NetWorker administration window, click the **Recover** tab.
5. Click **Recover > New Recover**.
6. In the **Recover Configuration** window, specify the following settings:
 - a. In the **Select the Recovery Hosts** screen, specify the following settings and click **Next**.
 - In the **Source Host** area, type the name or IP address of the source host with the backed data.
 - In the **Destination Host** area, specify whether you want to recover data to the same host or specify a new destination host name or IP address.
 - b. In the **Select the Data to Recover** screen, specify the following settings and click **Next**.
 - Click the **Browse** tab and select a date and time through which you want to recover data.
 - Browse and select the data to recover.
 - c. In the **Select the Recovery Options** screen, specify the file path for recovery and click **Next**.
 - d. In the **Obtain the Volume Recovery** screen, select **Allow NetWorker to select the required volumes for recovery (Recommended)** and click **Next**.
 - e. In the **Perform the Recovery** screen, specify a recovery name, recovery start time, and click **Run Recovery**.

- f. Check the results of the recovery operation through the **Check the Recovery Results** screen.

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.

Note: The steps described in the procedures are general guidelines only. They might change for different versions of Symantec NetBackup. Consult your DMA vendor documentation for the configuration information for a specific version of Symantec NetBackup.

Add an NDMP host

You must add an Isilon cluster as an NDMP host before you can back up 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 that you want to back up data from.
 - If you have a single Backup Accelerator or a Fibre Attached Storage node in the cluster, type the fully qualified domain name, host name, and the IPv4 or IPv6 address of the Backup Accelerator node or Fibre Attached Storage node.
 - If you have multiple Backup Accelerator nodes or Fibre Attached Storage nodes in the cluster, type the name of a SmartConnect zone that contains only the Backup Accelerator nodes or Fibre Attached Storage nodes.
 - If you are performing a three-way NDMP backup, type the fully qualified domain name (FQDN), host name, SmartConnect zone, and the IPv4 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 user name 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 or Fibre Attached Storage nodes, you must configure Symantec NetBackup to recognize those storage devices.

About this task

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 anisilon 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:


 **Note:** Refer to the [NetBackup documentation](#) to understand how to enable and use multistreaming.

Setting name	Setting value	Notes
Policy Type	NDMP	Required
Policy volume pool	The name of the volume pool you created	Required

Setting name	Setting value	Notes
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
	<code>set DIRECT=Y</code>	Optional. Enables direct access restore (DAR) for the directory. It is recommended that you enable DAR for all backups.
	<code>set HIST=F</code>	Optional. Specifies path-based file history. It is recommended that you specify path-based file history for all NetBackup backups.
	<code>set NEW_STREAM</code> 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 back up data stored on an Isilon cluster, including multi-stream backups. The following procedures explain how to configure NDMP backup with CommVault Simpana.

 **Note:** The steps described in the procedures are general guidelines only. They might change for different versions of CommVault Simpana. Consult your DMA vendor documentation for the configuration information for a specific version of CommVault Simpana.

Add a NAS client

With CommVault Simpana, you must add a NAS client for an Isilon cluster before you can back up 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 that you want to back up data from.</p> <ul style="list-style-type: none"> • If you have a single Backup Accelerator node or Fibre Attached Storage node in the cluster, type the fully qualified domain name, or the host name, or the IPv4 or IPv6 address of the Backup Accelerator node or Fibre Attached Storage node. • If you have multiple Backup Accelerator nodes or Fibre Attached Storage nodes in the cluster, type the name of a SmartConnect zone that contains only the Backup Accelerator nodes or Fibre Attached Storage nodes. • If you are performing a three-way NDMP backup, type the fully qualified domain name, host name, SmartConnect zone, and the IPv4 or IPv6 address of any node in the cluster.
NDMP Login	The name of the NDMP user account that you configured on the Isilon cluster.
NDMP Password	The password of the NDMP user account that 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.
Detect MediaAgent	If your Simpana server is located on a different host machine, select the media agent client name from this drop-down menu.

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 **Next**.
 - 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 for Primary Copy** list, select the name of the NDMP library you configured previously and click **Next**.
 - b. From the **MediaAgent** list, select the Isilon cluster you want to back up data from. If necessary, also select the appropriate drive pool.
 - c. Click **Next**.
3. From the **Scratch Pool** list, select **Default Scratch**.
4. (Optional) To enable multi-streaming, specify the **Number of Device Streams** setting as a number greater than one and up to the number of tape devices in the library. Although OneFS supports up to 64 streams, we recommend that you do not specify more than eight device streams for performance reasons.

 **Note:** We recommend that you enable multi-streaming to increase the speed of backup operations.

5. Click **Next**.
6. Select **Hardware Compression**, and then click **Next**.
7. Review your storage policy selections and click **Finish**.

You can monitor the progress of the backup through the CommVault Simpana user interface, or by running `isi ndmp sessions list` at the OneFS command-line interface.

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. To specify a new subclient, right-click in the Subclients window and select **All Tasks > New Subclient**.

The **Subclient Properties** window appears.

3. Enter a name for the subclient.
4. Select the **Content** tab and enter the full path of the directory you are backing up in the **Backup Content Path** field.
5. Click **Add**.
6. Select the **Advanced Options** tab and enter the number of data readers you are using. If you are doing a multi-stream backup, this number cannot exceed the number of tape devices you specified when creating the NDMP library.
7. Select the **Storage Policy** tab and select the storage policy you created from the **Storage Policy** drop-down list.
8. Click **OK**.

The **Backup Schedule** window appears.

9. Select **Do Not Schedule** to perform the backup manually at any time. Select the other options to associate the backup with an existing schedule policy, or to specify a date and time to perform the backup.
10. Click **OK**.

Configure backup options

To enable or disable faster incremental (using snapshots for backup), you can configure the backup options for the subclient you created previously.

Procedure

1. From the CommCell browser, select *<client computers>* > *<your client>* > **NAS > defaultBackupSet**.
2. Right-click on the name of the subclient you created previously.
3. Select **Backup** from the right-click menu.

The **Backup Options** dialog box appears.

4. In the **Select Backup Type** area, select **Full**.
5. (Optional) If you do not want to use snapshots in your backup:

- a. Click **Advanced**.
 - b. In the **Advanced Backup Options** dialog, select the **NAS Options** tab.
 - c. De-select the **Fast Incremental (Use snapshot for backup)** check box.
 - d. Click **OK**.
6. Click **OK**.
 7. (Optional) To verify that your backup is running, select the **Job Controller** tab. You can also use the `isi ndmp sessions list` command to verify.

Restore backed up data

You can use CommVault Simpana to restore data you previously backed up using this data management application (DMA).

About this task

Regardless of whether you performed a single-stream or multi-stream backup using CommVault Simpana, this DMA only supports single-stream restoration.

Procedure

1. From the CommCell browser, select **<client computers> > <your client> > NAS > defaultBackupSet**.
2. Right-click on the name of the subclient you created previously.
3. Select **Browse and Restore** from the right-click menu.
The **Browse and Restore Options** dialog box appears.
4. In the **Time Range** tab, select either **Latest Backup** or specify a time range from when the desired backup was performed.
5. (Optional) In the **Advanced Options** tab, specify criteria to refine the scope of the data restoration.
6. Click **View Content**.
A new tab appears specific to the subclient you created previously.
7. In the left pane, expand the NAS client to display the full path of the backed up data.
8. Right-click on the subdirectory you want to restore. For example, if you previously backed up a data set located in `/ifs/data/mydata`, right-click on `mydata`.
9. Select **Restore Current Selected** or **Restore All Selected**.
10. (Optional) Click **Advanced**.
11. (Optional) Select the **NAS Options** tab. The Use Direct Access Restore feature is enabled by default; select OFF if you do not want to use it.
12. In the **General** tab of the **Restore Options for Current/All Selected Items** window:
 - a. Verify that the correct destination client is displayed.
 - b. De-select the **Restore to same folder** check box.
 - c. Specify the correct file path that you want to use to restore data, such as `/ifs/data/myrestore`.
 - d. Click **OK**.
13. (Optional) To verify that your restoration is running, select the **Job Controller** tab. You can also use the `isi ndmp sessions list` command to verify.

Configuring NDMP backup with IBM Tivoli Storage Manager

You can configure OneFS and IBM Tivoli Storage Manager (TSM) to backup data stored on an Isilon cluster. The following procedures explain how to configure NDMP backup with IBM TSM.

Note: The steps described in the procedures are general guidelines only. They might change for different versions of IBM TSM. Consult your DMA vendor documentation for the configuration information for a specific version of IBM TSM.

Initialize an IBM Tivoli Storage Manager server for an Isilon cluster

You must initialize an IBM Tivoli Storage Manager server to manage NDMP backups on an Isilon cluster.

Procedure

1. Open the Tivoli Storage Manager Management Console.
2. Expand the **Tivoli Storage Manager** folder.
3. Right-click the host name of your local machine and then select **Add a New Tivoli Storage Manager Server**.
4. In the **Initial Configuration Task List** window, click **Minimal configuration**.
5. Click **Start**.
6. Follow the prompts to configure a Tivoli Storage Manager server.

Configure an IBM Tivoli Storage Manager server for an Isilon cluster

You can configure an IBM Tivoli Storage Manager (TSM) server to manage NDMP backups on an Isilon cluster.

About this task

Configure the TSM server by following these steps:

1. Configure the tape library.
2. Configure your system for the backup and restore operations.
3. Define a virtual filespace mapping and perform the backup operation.
4. Define a virtual filespace mapping if you are restoring data to a location other than the location from where you backed up the data, and then perform the restore operation. Otherwise, perform the restore operation without defining a virtual filespace mapping.

Configure a tape library

With IBM Tivoli Storage Manager (TSM), you must configure a tape library that contains the tape devices for backup and restore operations.

Procedure

1. Create a TSM node by running the following command:

```
register node <node-name> <admin-password> userid=admin domain=STANDARD
type=NAS
```

The following command creates a TSM node called node001:

```
register node node001 password123 userid=admin domain=STANDARD type=NAS
```

2. Define a data mover for the node you want to back up data from by running the following command:

```
define datamover <datamover_name> type=NAS
hladdress=<backup_accelerator_ip_address> lladdress=<ndmp_port>
userid=<ndmp_username> password=<ndmp_password> dataformat=ndmpdump
```

The following command defines a data mover for node001:

```
define datamover node001 type=NAS hladdress=10.13.17.117
lladdress=10000 userid=ndmp password=ndmppw dataformat=ndmpdump
```

3. Define a tape library by running the following command:

```
def libr <library-name> libtype=scsi
```

The following command defines a tape library called ISILIB:

```
def libr ISILIB libtype=scsi autolabel=overwrite shared=no
serial=autodetect
```

4. Define the path for the data mover by running the following command:

```
define path <data_mover_name> <tape_library> srctype=datamover
desttype=library device=<backup_accelerator_library_name>
```

Specify `device` as the name of the device entry for the tape library on the Isilon cluster.

The following command defines a path for the data mover created in step 2.

```
define path node001 ISILIB srctype=datamover desttype=library
device=mc005
```

5. Define drives for the tape library by running the following command:

```
define drive <tape_library> <backup_accelerator_tape_drive>
serial=autodetect online=yes element=<tape_element_address>
```

The following commands defines four tape drives and configures TSM to automatically detect the addresses of the tape drives.

```
define drive ISILIB tape015 serial=autodetect online=yes element=256
define drive ISILIB tape016 serial=autodetect online=yes element=257
define drive ISILIB tape017 serial=autodetect online=yes element=258
define drive ISILIB tape018 serial=autodetect online=yes element=259
```

6. Define paths for tape drives by running the following command:

```
define path <data_mover_name> <path_name> srctype=datamover
desttype=drive libr=<tape_library>
device=<backup_accelerator_tape_drive>
```

The following commands define paths for the tape drives defined in the previous step:

```
define path node001 tape015 srctype=datamover desttype=drive
libr=ISILIB device=tape015
define path node001 tape016 srctype=datamover desttype=drive
libr=ISILIB device=tape016
define path node001 tape017 srctype=datamover desttype=drive
libr=ISILIB device=tape017
define path node001 tape018 srctype=datamover desttype=drive
libr=ISILIB device=tape018
```

7. Label the tape media by running the following command:

```
label libvol <tape_library> voll=<tape_media_label>
```

The following commands create labels for the tape media in the tape library:

```
label libvol ISILIB voll=C90000LA search=yes overwrite=yes
labelsource=barcode checkin=scratch
label libvol ISILIB voll=C90001LA search=yes overwrite=yes
labelsource=barcode checkin=scratch
label libvol ISILIB voll=C90002LA search=yes overwrite=yes
labelsource=barcode checkin=scratch
label libvol ISILIB voll=C90003LA search=yes overwrite=yes
labelsource=barcode checkin=scratch
```

8. Verify that the tape library has been configured accurately by performing the following steps:
- Verify that the tapes are online by running the following command:

```
query tape libr=<tape_library>
```

- Query the tape media by running the following command:

```
query libvol libr=<tape_library>
```

- Query the path configured to ensure that it is accurate:

```
query path <data_mover_name>
```

- Audit the tape library by running the following command:

```
audit library <lib_name>
```

Configure your system for backup and restore operations

With IBM Tivoli Storage Manager (TSM), you must configure your system for performing backup and restore operations.

Procedure

1. Define a device class by running the following command:

```
define devclass <class-name> devtype=<dev-type> library=<library-name>
mountretention=0 estcapacity=120g
```

The following command defines a device class called ISICLASS:

```
define devclass ISICLASS devtype=NAS library=ISILIB mountretention=0
estcapacity=120g
```

2. Define an NDMP storage pool by running the following command:

```
define stgpool <ndmp-pool-name> <class-name> maxscratch=10
dataformat=ndmpdump
```

The following command defines an NDMP storage pool called NDMPPPOOL:

```
define stgpool NDMPPPOOL ISICLASS maxscratch=10 dataformat=ndmpdump
```

3. Define a device class for the table of contents (TOC) of the files to be backed up by running the following command:

```
def devclass TOC devtype=<dev-type>
```

The following command defines a device class for the TOC for a device type called file:

```
def devclass TOC devtype=file
```

4. Define a storage TOC by running the following command:

```
define stgpool <TOC_POOL_NAME> DISK
```

The following command defines a storage TOC called TOC:

```
define stgpool TOC DISK
```

5. Define a volume for the storage pool by running the following command:

```
define volume <TOC_NAME> <PATH> formatsize=<SIZE_IN_MB> wait=no
```

The following command defines a storage pool volume:

```
define volume TOC'e:\\Program Files\\Tivoli\\TSM\\server1\
\tsm_6.toc.dsm'formatsize=1024 wait=no
```

6. Define a domain by running the following command:

```
def domain domain_name
```

The following command defines a domain called NASDOMAIN:

```
def domain NASDOMAIN
```

7. Define a policy by running the following command:

```
def pol domain_name policy_name
```

The following command defines a policy called NASPOLICY:

```
def pol NASDOMAIN NASPOLICY
```

8. Define a management class by running the following command:

```
define mgmtclass <domain_name> <policy_set_name> <mgmt_class_name>
```

The following command defines a management class called NASMGMT:

```
define mgmtclass STANDARD STANDARD NASMGMT
```

9. Define a copy group by running the following command:

```
define copygroup <domain_name> <policy_set_name> <mgmt_class_name>
type=backup destination=<ndmp_pool_name> [serialization=static]
tocdestination=<toc_pool_name>
```

The following command defines a copy group:

```
define copygroup STANDARD STANDARD ISIMC type=backup
destination=NDMPPOOL serialization=static tocdestination=TOCPOOL
```

10. Assign a default management class by running the following command::

```
assign defmgmtclass <domain_name> <policy_set_name> <mgmt_class_name>
```

In order to allow a TOC to be backed up to disk, a specific management class must be defined because the standard management class cannot be modified. The following

command assigns the ISIMC management class as the default for the standard policy set and the standard domain:

```
assign defmgmtclass STANDARD STANDARD ISIMC
```

11. Validate the policy set by running the following command:

```
validate policyset <domain_name> <policy_set_name>
```

The following command validates the standard policy set.

```
validate policyset STANDARD STANDARD
```

12. Activate the policy set by running the following command:

```
activate policyset <domain_name> <policy_set_name>
```

The following command activates the standard policy set:

```
activate policyset STANDARD STANDARD
```

13. Update the TSM node to the domain that you created in step 6 by running the following command:

```
update node <ip_addr_or_node_name> <domain_name>
```

The following command updates the node to the domain that you created in step 6:

```
activate policyset 10.27.49.39 NASDOMAIN
```

14. Update the path to the NAS library by running the following command:

```
update path <ip_addr_or_node_name> <tape_library> srctype=datamover  
desttype=libr
```

The following command updates the path to the NAS library for node001:

```
update path node001 ISILIB srctype=datamover desttype=libr online=yes
```


Define a virtual filesystem mapping and perform a backup operation

As a part of configuring IBM Tivoli Storage Manager for managing NDMP backups on an Isilon cluster, you must define a virtual filesystem mapping before performing the backup operation.

Procedure

1. Define virtual filesystem mapping by running the following command:

```
def virtualfs <ip_addr_or_node_name> <virtual_filespace_name>
<file_system_name> <path>
```

The following command creates a virtual filesystem mapping for /ifs/data.

```
def virtualfs node001 /data-backup /ifs/data
```

2. Perform the backup operation by running the following command:

```
backup node <ip_addr_or_node_name> <virtual_filespace_name>
mode=backup_mode
```

The following command performs a backup operation on node001.

```
backup node node001 /data-backup mode=full toc=yes wait=yes
```

Define a virtual filesystem mapping for the restore operation

If you are restoring data to a location that is different from the location that you backed up the data from, define a virtual filesystem mapping for the restore operation. You must perform this process as a part of configuring IBM Tivoli Storage Manager for managing NDMP backups on an Isilon cluster.

Procedure

1. Define a virtual filesystem mapping by running the following command:

```
def virtualfs <ip_addr_or_node_name> <virtual_filespace_name>
<file_system_name> <path>
```

The following example creates a filesystem mapping for /ifs/data.

```
def virtualfs node001 /data-restore /ifs /data
```

2. Restore data by running the following command:

```
restore node <ip_addr_or_node_name> <virtual_fs_backup_location>
<restore_file_system_name>
```

Note: If you are restoring data to the same location that you backed up the data from, you do not need to define a virtual filesystem mapping.

The following example restores data on node001:

```
restore node node001 /data-backup /data-restore wait=yes
```

CHAPTER 8

Sparse file support

This chapter contains the following topics:

- [Commvault sparse file support: data aging and space reclamation](#)..... 124
- [Enable or disable SMB sparse file support](#)..... 124

Commvault sparse file support: data aging and space reclamation

You can add support for file sparsing through SMB, specific to Commvault backup workflows.

The Commvault backup application interfaces with OneFS over SMB protocol. The Commvault application stores all data needing to be backed up in large, indexed files. When data ages past a retention policy, if Commvault cannot delete the entire file, it attempts to free up sections of the tarball by converting them to sparse blocks. These tarball files can be large (more than 3 GB), and the range of hole punch arguments can also be large. To accomplish this, Commvault will send an FSCTL SET ZERO DATA call over SMB. This request comes with a range, and the expectation is that the server must convert the specified range to sparse blocks, freeing up space. This feature is provided in addition to the existing support for space savings from the delete and truncate operations on backup container files.

Capacity utilization with sparse file savings is best viewed with the `du -sh` command on the Isilon cluster. The file and folder properties from Windows represents the savings from delete and truncate operations but may not represent the sparse file savings.

Enable or disable SMB sparse file support

You can enable or disable SMB sparse file support from the OneFS command line.

Before you begin

This procedure is available only through the OneFS command-line interface (CLI).

Procedure

1. To view if the sparse file support is enabled or not, run the following command:

```
isi smb settings share view | grep Sparse
```

The system displays output similar to the following example:

```
Sparse File: No
```

or,

```
Sparse File: Yes
```

2. If sparse file support is enabled and you want to disable it, run the following command:


```
isi smb settings share modify --sparse-file=False
```
3. If sparse file support is disabled and you want to enable it, run the following command:


```
isi smb settings share modify --sparse-file=True
```
4. After you adjust settings, confirm that the setting is correct. Run the following command:


```
isi smb settings share view | grep Sparse
```