

Dell PowerScale: NDMP Technical Overview and Design Considerations

August 2022

H18790

White Paper

Abstract

This white paper describes an overview of Dell PowerScale Network Data Management Protocol (NDMP) backup and recovery solutions. It also provides NDMP design considerations, best practices, and NDMP and OneFS feature integration details.

Dell Technologies

Copyright

The information in this publication is provided as is. Dell Inc. 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 software described in this publication requires an applicable software license.

Copyright © 2021-2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners.

Published in the USA August 2022 H18790.

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

Contents

Executive summary.....	4
Introduction	5
NDMP overview	5
NDMP configuration.....	8
NDMP considerations and best practices	15
NDMP and OneFS feature integration.....	17
Troubleshooting and logs	21
References.....	23

Executive summary

Overview

As unstructured data continues to grow exponentially, organizations face the daunting task of ensuring the protection of critical data on network attached storage (NAS) systems. In many cases, backup windows and budgets are shrinking as well.

Traditional backup solutions typically rely on full and incremental backups that are scheduled to run regular jobs. Unfortunately, the full backups often extend beyond the available backup window, which may result in unprotected data. If only a portion of a large file has changed, an incremental backup may take as long as a full backup because the whole file must be backed up. Traditional backup limitations can lead to data loss, declining user productivity, and constrained NAS system performance.

If the primary NAS system fails and data is not replicated to the secondary NAS system, there may be data loss on the primary NAS system. However, large amounts of data and limited WAN bandwidth can make data replication difficult. Some applications use snapshots, but snapshots are not a substitute for a comprehensive backup and recovery strategy.

Organizations require an appropriate balance of performance, reliability, management, and cost to meet their recovery time objectives (RTOs) and recovery point objectives (RPOs).

This document describes how Dell PowerScale NDMP backup solutions help organizations meet these requirements and understand NDMP best practices.

Audience

This document is intended for system engineers, storage administrators, backup administrators, and IT managers.

Revisions

Date	Description
June 2021	Initial release
December 2021	B100 backup accelerator release
August 2022	Updated best practices

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Jason He

Note: For links to other documentation for this topic, see the [PowerScale Info Hub](#).

Introduction

Overview

NDMP is an open network protocol that allows direct data transfers between storage and backup devices. The protocol allows backup administrators to back up critical data using any combination of compliant file servers, backup devices, and data management applications (DMA). NDMP separates the data path and control path and enables centralized control of enterprise-wide network data management while minimizing network traffic. NDMP simplifies the management and protection of business-critical data.

Dell PowerScale OneFS supports both NDMP versions 3 and 4 and facilitates the backup and restore functionality using NDMP. OneFS NDMP supports both two-way (referred to as local or direct NDMP backup) and three-way (referred to as remote NDMP) backup topologies.

The OneFS NDMP module includes the following functionality or features:

- Full (level 0) and incremental backups
- Token-based backups
- Snapshot-based incremental backups
- Multi-stream Backup (MSB)
- Backup Restartable Extension (BRE) and Multi-stream Backup Restartable Extension (MSBRE)
- Backup of file attributes
- Backup of Access Control Lists (ACLs)
- Backup of Alternate Data Streams (ADSS)
- NDMP tar and dump backup types
- Direct Access Restore (DAR) and Directory Direct Access Restore (DDAR)
- NDMP environment variables
- Including and excluding specific files and directories from backup
- NDMP Redirector
- NDMP Throttler
- Integration with OneFS features

NDMP overview

Overview

NDMP specifies a common architecture and data format for backups and restores of NAS. With NDMP, heterogeneous network file servers can directly communicate to tape devices for backup and restore operations. NDMP addresses the problems caused by the integrations of different backup software or DMA, file servers, and tape devices.

NDMP overview

The NDMP architecture is a client/server model with the following characteristics:

- The NDMP host is a file server that is being protected with an NDMP backup solution. The NDMP server is a virtual state machine on the NDMP host that is controlled using NDMP.
- The backup software is considered as a client to the NDMP server.

OneFS supports the following two types of NDMP backups:

- NDMP two-way backup
- NDMP three-way backup

In both the backup models, OneFS takes a snapshot of the backup directory to ensure consistency of data and the backup operates on the snapshot instead of the source directory. Also, this design allows users to continue normal read/write activities. OneFS makes entries in the file history that is transferred from the PowerScale cluster to the backup server during the backup.

NDMP two-way backup

The NDMP two-way backup is also known as the local or direct NDMP backup. The backup moves the backup data directly from the PowerScale cluster to the tape devices without moving to the backup server over the network. This backup model is the most efficient model and usually provides the best performance.

In this backup model, OneFS must detect the tape devices before you back up data to that tape devices. PowerScale cluster provides the option for NDMP two-way backups as shown in the following figure. You can connect the PowerScale cluster to a Backup Accelerator node and connect tape devices to that node. The Backup Accelerator node is synonymous with a Fibre Attached Storage node without adding primary storage and offloads NDMP workloads from the primary storage nodes. You can directly connect tape devices to the Fibre Channel ports on the PowerScale cluster or Backup Accelerator node using Fibre Channel. Alternatively, you can connect Fibre Channel switches to the Fibre Channel ports that connect tape devices to the PowerScale cluster or Backup Accelerator node.

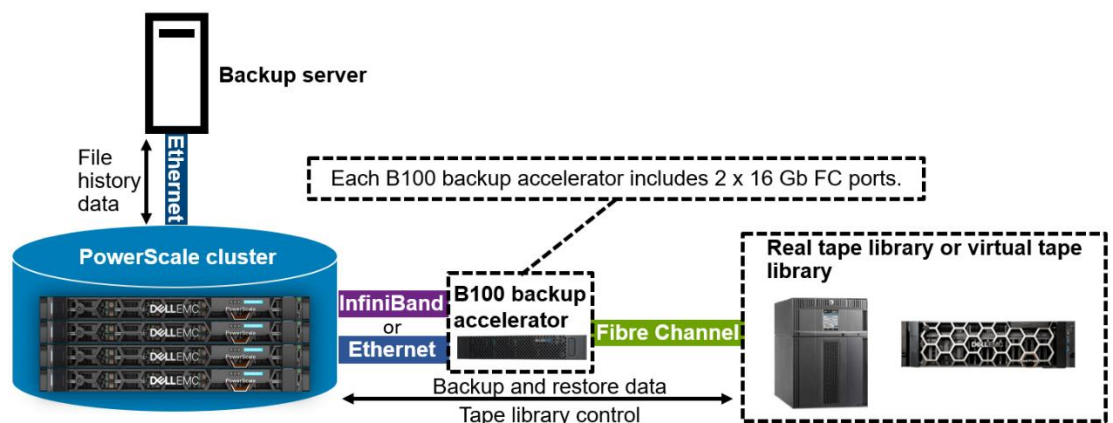


Figure 1. NDMP two-way backup with B100 backup accelerator connected to the PowerScale cluster

Note: See the section [NDMP considerations and best practices](#) for recommendations regarding PowerScale nodes with a Fibre Attached Storage node. NDMP can be an intensive workflow, so consult your Dell representative for appropriate sizing.

The following table shows the NDMP two-way backup supported by PowerScale.

Table 1. NDMP two-way backup supported by PowerScale

NDMP two-way backup options	Generation 5 PowerScale nodes with an InfiniBand back end	Generation 6+ PowerScale nodes with an InfiniBand back end	Generation 6+ PowerScale nodes with an Ethernet back end
B100 backup accelerator	Supported	Supported	Supported

Note: The B100 backup accelerator requires OneFS 9.3.0.0 or later.

NDMP three-way backup

The NDMP three-way backup, also known as the remote NDMP backup, is shown in the following figure.

In this backup mode, the tape devices are connected to the backup media server. OneFS does not detect tape devices on the PowerScale cluster and Fibre Channel ports are not required on the PowerScale cluster. The NDMP service runs on the NDMP server or the PowerScale cluster. The NDMP tape service runs on the backup media server. A DMA on the backup server instructs the PowerScale cluster to start backing up data from the PowerScale cluster to the backup media server over the network. The backup media server moves the backup data to tape devices. Both the servers are connected to each other across the network boundary. Sometimes, the backup server and backup media server reside on the same physical machine.

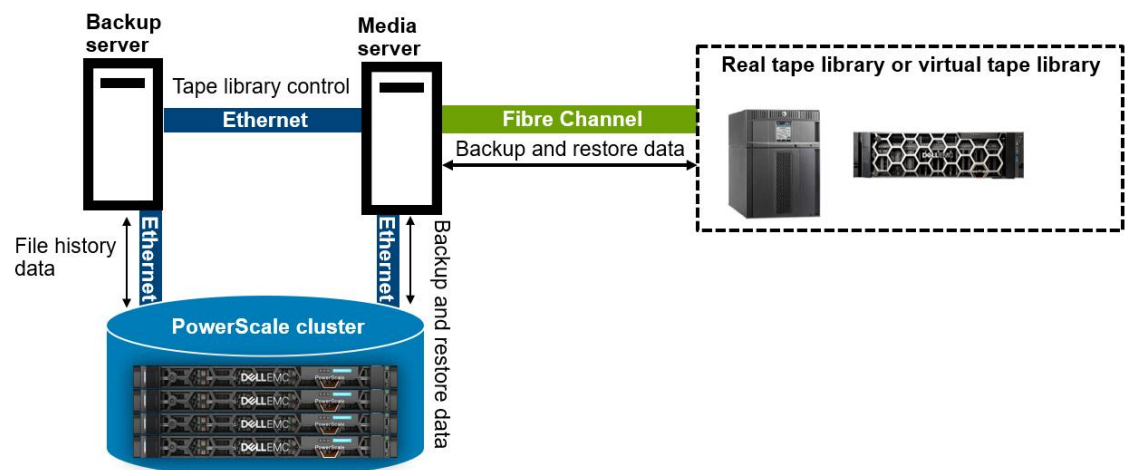


Figure 2. NDMP three-way backup

NDMP configuration

Backup types

This section describes the NDMP features supported by OneFS.

NDMP backup has two types of backup. OneFS NDMP supports tar and dump backup formats for backup. An NDMP backup session can either be a full backup or an incremental backup. OneFS NDMP supports tar and dump backup formats for backup.

- **Full backup:** A full backup backs up all files and directories under a path.
- **Incremental backup:** An incremental backup backs up changed files under the same path of a full backup.

OneFS provides the following NDMP features to back up files. The features give users more flexibility for NDMP backups.

- **Level-based backup:** OneFS NDMP supports full backups (level 0) and incremental (levels 1 to 9) backups. The first backup of a path must be a full backup. Thereafter, incremental backup backs up the changed files since the last full or lower-level backup. OneFS detects the changed files or directories by ctime (change time) or mtime (modification time) of objects. All subsequent level 9 incremental backups will become cumulative incremental backups after performing a level 9 incremental backup. You can run a full backup before resuming the differential incremental backups. OneFS NDMP also supports incremental forever (level 10) backup. This backup level only backs up the changed files since the most recent incremental (levels 1 to 9) backup or the last incremental forever (level 10) backup. A DMA can pass the LEVEL environment variable to specify a file path.
- **Token-based backup:** Token-based backups use a timestamp to extend the level-based backups. With a token-based backup, the backup backs up the changed files after the timestamp. OneFS detects the changed files or directories by ctime (change time) or mtime (modification time) of objects. A timestamp 0 forces a full backup. A DMA can pass a timestamp value for a file path in the BASE_DATE environment variable.
- **Snapshot-based incremental backup:** A snapshot-based incremental backup (referred to as faster-incremental backup) leverages snapshots to identify files to be backed up. For each snapshot-based incremental backup, OneFS identifies the changed files between two snapshots. One snapshot is from the prior level backup, the other snapshot is from the current backup. OneFS retains the snapshot for backup that is used to identify the changed files when the next level backup is performed. During a snapshot-based incremental backup, OneFS does not walk through all whole dataset and check timestamp of each file to identify files to be backed up. OneFS can perform a snapshot-based incremental backup significantly faster if the change rate is low. You can set the BACKUP_MODE environment variable to SNAPSHOT to enable snapshot-based incremental backups. This feature is integrated with Avamar without setting the BACKUP_MODE environment variable. The following table shows DMA support for snapshot-based incremental backups.

Table 2. DMA support for snapshot-based incremental backups

DMA	DMA-integrated
Networker Enabled	Enabled only through an environment variable
Symantec NetBackup	Enabled only through an environment variable
Avamar	Yes
Commvault	Enabled only through an environment variable
Tivoli Storage Manager	Enabled only through an environment variable
Symantec Backup Exec	Enabled only through an environment variable
NetVault	Enabled only through an environment variable
ASG-Time Navigator	Enabled only through an environment variable

- **Multi-stream backup:** A multi-stream backup runs multiple backup sessions on the same dataset. OneFS supports the NDMP multi-stream backup in OneFS 8.0.0.0 and later releases. With multi-stream backup, you can use your DMA to specify multiple streams of data to back up concurrently.
- **File list backup:** File list backup allows users to back up files in the file list based on the BACKUP_FILE_LIST environment variable setting. Only Networker and Symantec NetBackup can pass the BACKUP_FILE_LIST environment variable to specify a file path. During a file list backup, OneFS does not walk through the entire backup directory for backing up files.
- **Backup Restartable Extension (BRE):** Starting from OneFS 7.1.0, OneFS supports the NDMP BRE feature that allows to resumption of a failed backup job from the last checkpoint taken before the failure. Restartable backups are supported for NetWorker 8.1 and later versions and Commvault DMAs. You can enable restartable backups in your DMA.
- **Multi-stream Backup Restartable Extension (MSBRE):** Starting from OneFS 9.1.0.0, OneFS supports the NDMP MSBRE that allows to resume a failed multi-stream backup job from the last checkpoint taken before the failure. Restartable backups are supported for NetWorker 8.1 and later versions and Commvault DMAs. You can enable restartable backups in your DMA.

Note: OneFS NDMP supports backing up file attributes, access control lists, and alternate data streams.

Run the following command to check the backup type (OP) of the running NDMP session.

```
# isi ndmp sessions list
```

For the completed NDMP session, check the backup type (BACKUP_METHOD) through the logs at `/ifs/.ifsvar/modules/ndmp/trace`.

The following are examples of the NDMP backup type as indicated through the above command or logs.

```
B(T0): Token based full backup
```

```
B(Ti): Token based incremental backup
```

NDMP configuration

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

Restore types

After successfully backing up files using NDMP, you can perform a restore of files from the NDMP backup. OneFS NDMP supports full restore, selective file restore, DAR, and DDAR.

- **Full restore:** The full restore process reads all data across multiple tapes to restore all data to the chosen location. A full restore typically takes a painfully long time to complete. If only a subset of files in the backup needs to be restored, the full restore is inefficient.
- **Selective file restore:** This feature can restore independent files of the backup without having to restore all files in the backup. During a selective file or files restore, the restore process reads from the beginning of the backup tape until it locates the file or files. Then the process restores it to the chosen location.
- **DAR:** DAR provides DMA with positioning information of a file on tape. It allows the NDMP server to go directly to the location of that file on tape and restore that file. DAR can significantly reduce the restoration time for an individual file restoration.
- **DDAR:** DDAR is an extension of DAR. DAR provides DMA with positioning information of a directory on tape. It allows the NDMP server to go directly to the location of that directory on tape and restore that directory. DDAR can significantly reduce the restoration time for an individual directory restoration.

Run the following command to check the restore type (OP) of the running NDMP session.

```
# isi ndmp sessions list
```

For the completed NDMP session, you can check the restore type (DIRECT) through the logs at `/ifs/ifsvar/modules/ndmp/trace`.

The following are examples of NDMP restore type indicated through the above command or logs.

```
R(F): Full restore  
R(D): DAR  
R(S): Selective restore
```

NDMP environment variables

NDMP environment variables control the behavior of DMA for each backup and restore session. You can manage the NDMP environment variables for each backup path on the NDMP server or the PowerScale cluster. You can also specify NDMP environment variables through your DMA. NDMP environment variables can pass configuration

parameters to the NDMP server. The NDMP environment variables passed from DMA always override the NDMP environment variables that are set on the NDMP server.

Note: Symantec NetBackup and Dell NetWorker are the only two DMAs that allow you to directly set environment variables and pass them from DMA to the NDMP server.

The following table shows the NDMP environment variables supported by OneFS NDMP.

Table 3. NDMP environment variables

Environment variable	Supported values	Default	Description
BACKUP_FILE_LIST	<file-path>	None	This variable is used to specify the sorted file name list for file list backup.
BACKUP_MODE	TIMESTAMP or SNAPSHOT	TIMESTAMP	This variable is used to enable or disable snapshot-based incremental backups. To enable snapshot-based incremental backups, you can specify SNAPSHOT.
BACKUP_OPTIONS	0x00000001, 0x00000002, 0x00000004, 0x00000400, 0x00000200, or 0x00000100	0	This variable controls the behavior of the backup operations. 0x00000001: This option 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 use the DUMP_DATE value to set BASE_DATE for the next backup operation. 0x00000002: This option retains the backup snapshot of a token-based backup in the dumpdates 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. 0x00000004: This option retains the previous snapshot. After a snapshot-based incremental backup, the prior snapshot is saved at level 10. To avoid two snapshots at the same level, the prior snapshot is kept at a lower level in the dumpdates file. This option allows the BASE_DATE and BACKUP_MODE=snapshot settings to trigger a snapshot-based 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 dumpdates file to find the prior backup. Even though the DMA fails the latest snapshot-based incremental backup, OneFS retains the prior snapshot. The DMA can then retry the snapshot-based incremental backup in the next backup cycle using the BASE_DATE value of the prior backup. The following settings are applicable only to datasets containing the PowerScale OneFS CloudPools SmartLink files:

NDMP configuration

Environment variable	Supported values	Default	Description
			<p>0x0000400: This option backs up files as SmartLink files with file data. This option is the combo copy backup option. This backup option is the default setting.</p> <p>0x0000200: This option backs up files as SmartLink files without file data. This option is the shallow copy backup option.</p> <p>0x0000100: This option backs up files as regular files that read file data from the cloud. This option is the deep copy backup option.</p>
BASE_DATE	0 or DUMP_DATE value	0	This variable is used to enable a token-based incremental backup. The dumpdates file will not be updated in this case.
DIRECT	Y or N	N	<p>This variable is used to enable or disable DAR and DDAR.</p> <p>Y: This option enables DAR and DDAR.</p> <p>N: This option disables DAR and DDAR.</p>
EXCLUDE	<file-matching-pattern>	None	If you specify this option, OneFS NDMP does not back up files and directories that meet the specified pattern. You can separate multiple patterns with a space.
FILES	<file-matching-pattern>	None	<p>If you specify this option, OneFS NDMP backs up only files and directories that meet the specified pattern. You can separate multiple patterns with a space.</p> <p>Note: Files and directories are matched first and then the EXCLUDE pattern is applied.</p>
HIST	D, F, Y, or N	Y	<p>This variable is used to specify the file history format.</p> <p>D: This option specifies node-based file history.</p> <p>F: This option specifies path-based file history.</p> <p>Y: This option specifies the default file history format determined by your NDMP backup settings.</p> <p>N: This option disables file history.</p>
LEVEL	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, or 10	0	<p>This variable is used to specify the level of NDMP backup to perform.</p> <p>0: This option is to perform a full NDMP backup.</p> <p>1–9: These options are to perform an incremental backup at the specified level.</p> <p>10: This option is to perform incremental forever backups.</p>
MSB_RETENTION_PERIOD	<integer>	300	For a multi-stream backup session, specifies the backup context retention period in seconds.
MSR_RETENTION_PERIOD	0 to 60*60*24	600	For a multi-stream restore session, specifies the recovery context retention period in seconds within which a recovery session can be retried.

Environment variable	Supported values	Default	Description
RECURSIVE	Y or N	Y	This variable is only for restore sessions and is used to specify that the restore session should recover files or subdirectories under a directory automatically.
RESTORE_BIRTHTIME	Y or N	N	This variable is used to specify whether to restore the birth time for a restore session.
RESTORE_HARDLINK_BY_TABLE	Y or N	N	For a single-threaded restore session, determines whether OneFS NDMP restores hard links by building a hard-link table during restore operations. This variable is used to specify if hard links are incorrectly backed up and restore operations are failing. If a restore operation fails because hard links were incorrectly backed up, the following message appears in the NDMP backup logs: Bad hardlink path for <path> Note: This variable is not effective for a parallel restore operation.
RESTORE_OPTIONS	0x00000001, 0x00000002, 0x00000004, 0x00000100, or 0x00000200	0	This environment variable controls the behavior of the restore operations. 0x00000001: This option performs a single threaded restore operation. 0x00000002: This option restores directory attributes to the existing directories. 0x00000004: This option creates intermediate directories with default attributes. The default behavior is to get attributes from the first object under a given directory. The following settings are applicable only to datasets backed up with the combo copy backup option: 0x00000100: This option forces deep copy restoration of the SmartLink files. This option is to restore the backed-up SmartLink files as regular files on the target cluster. 0x00000200: This option forces shallow copy restoration of the SmartLink files. This option is to restore the backed-up SmartLink files as SmartLink files on the target cluster. This restore option is the default setting.
UPDATE	Y or N	Y	This variable determines whether OneFS updates the dumpdates file. Y: OneFS updates the dumpdates file. N: OneFS does not update the dumpdates file.

Preferred IP addresses for NDMP three-way sessions

OneFS NDMP provides a feature to select the most desirable network interfaces and preferred network for NDMP three-way sessions. You can configure one preferred IP address to a cluster or network subnet.

Note: Avamar and Commvault DMAs support selecting the most desirable network interface for NDMP three-way session by the preferred IP addresses on the NDMP server or the PowerScale cluster. Other DMAs do not support the preferred IP addresses feature.

A preferred IP addresses setting is a list of prioritized IP addresses on the NDMP server. The NDMP server connects to a data server or media server using the preferred IP addresses during an NDMP three-way backup or restore.

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

Run the following command through the command-line interface to configure a preferred IP to a subnet defined in OneFS:

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

Include and exclude files and directories from backup

OneFS NDMP supports features to include or exclude files and directories from backup.

The **include** feature allows for flexible selection of which files or directories that are included when backing up data. You can include files and directories from NDMP backup operations by specifying NDMP the environment variable **FILES** through a DMA. If you include a file or directory for a backup and do not exclude that file or directory, that file or directory will only be backed up. In other words, all other files and directories are automatically excluded from backups.

The **exclude** feature allows for flexible selection of which files or directories that are excluded when backing up data. You can exclude files and directories from NDMP backup operations by specifying NDMP the environment variable **EXCLUDE** through a DMA. If you exclude a file or directory for a backup, all files and directories except the excluded file or directory will be backed up.

If you set both the include and exclude patterns for a backup, the include pattern is first processed followed by the exclude pattern. If you set both the include and exclude patterns for a backup, any excluded files or directories under the included directories will not be backed up. If the excluded directories are not found in any of the included directories, the exclude specification will have no effect for the backup.

NDMP redirector

NDMP Redirector is introduced in OneFS 8.2.0 and is disabled by default. NDMP Redirector distributes NDMP two-way sessions automatically over PowerScale nodes. You only can enable NDMP Redirector by command-line interface to automatically distribute NDMP two-way sessions to other nodes with lesser loads. Before redirecting the

NDMP sessions, NDMP Redirector checks CPU usage, the number of running NDMP sessions, and the availability of tape devices.

Note: Virtual Tape Library (VTL) does not know NDMP Redirector due to LONG FORM READ POSITIONING even though NDMP redirector is enabled. The NDMP two-way sessions still work without NDMP Redirector. OneFS will not distribute NDMP two-way sessions to other nodes if VTL is only configured for the NDMP two-way solution.

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

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

NDMP throttler

NDMP Throttler is introduced in OneFS 8.2.0 and is disabled by default. NDMP Throttler is global to all nodes in the PowerScale cluster and manages CPU usages of both NDMP two-way and three-way sessions. NDMP throttler ensures that NDMP operations do not compete with things that are performance sensitive like SMB and NFS user activity. You only can enable NDMP Throttler by command-line interface.

Run the following commands through the command-line interface to enable NDMP Throttler, and set the throttler CPU threshold value in percent:

```
# isi ndmp settings global modify --enable-throttler true
# isi ndmp settings global modify --throttler-cpu-
threshold=<value-in-percent>
```

NDMP considerations and best practices

Overview

Most customers back up their data after business hours because the backup operation does not compete with the user activity from performance sensitive protocols like SMB and NFS. NDMP backup and restore performance depends on many factors, such as the NDMP backup models, network bandwidth, DMA configuration, available system resources, directory depth, number of files per directory, and file size.

The following considerations and best practices can help you optimize OneFS NDMP backup and restore operations.

- Install the latest roll up patches (RUP) for OneFS and DMA. We strongly recommend installing OneFS 8.2.2 GA RUP 2021-09 or later Patch and OneFS 9.1.0.14+ when using OneFS NDMP.
- Run a maximum of eight NDMP concurrent sessions for each Fibre attached storage node. Run four NDMP concurrent sessions for each PowerScale backup accelerator node to obtain optimal throughput in each session.
- Isolate NDMP backup sessions from other workloads such as SMB or NFS user activity.
- For NDMP two-way backups, the B100 backup accelerator requires OneFS 9.3.0.0 or later.

- Assign static IP addresses to Fibre Attached Storage nodes. Attach more Fibre Attached Storage nodes to larger PowerScale clusters or if you are backing up to more tape devices. The following table shows the recommend number of PowerScale nodes for each Fibre attached storage node.

Table 4. Recommended PowerScale nodes for each Fibre attached storage node

PowerScale node type	Recommended number of PowerScale nodes for each Fibre Attached Storage node
F-Series	3
H-Series	3
A-Series	3
X-Series	3
NL-Series	3
S-Series	3
HD-Series	3

- NDMP backups result in high Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs). You can attach one or more Fibre Attached Storage nodes to the PowerScale cluster before running NDMP two-way sessions and reduce the RPO and RTO.
- NDMP backup and restore performance depends on the file size. The performance 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 restore files to another PowerScale cluster using NDMP, we recommend having the same OneFS and patch version installed on the source and target clusters.
- If you are performing NDMP three-way backups, run multiple NDMP sessions on multiple nodes in the PowerScale cluster.
- For NDMP three-way backups, configure using dedicated switch ports for all NDMP PowerScale nodes.
- Use Dell PowerScale SmartPools and file pool policies to move data to a node pool for NDMP backups.
- With NDMP, compressed chunks are decompressed. Each file is then reassembled and sent over Fibre Channel (2-way NDMP) or Ethernet (3-way NDMP) in uncompressed form to the backup devices.
- With NDMP, deduplicated data is rehydrated when it leaves the cluster, and shadow stores and shadow links are not preserved on the backup. The NDMP tape device or VTL will need to have sufficient space to house the full size of the dataset. When NDMP backs up deduplicated data, the NDMP performance will be impacted.
- Recover files through DDAR if you recover large numbers of files frequently.
- Use the largest tape block size available for OneFS to increase NDMP backup and restore throughput.

- If possible, do not include or exclude files from NDMP backups. Including or excluding files can affect backup performance, due to filtering overhead.
- Limit the depth of nested subdirectories in the OneFS file system for NDMP backups.
- Enable parallelism on DMA to enable multi-stream backup if DMA supports this option, and ensure that there are enough tape devices to support the multi-stream backup. This feature allows OneFS NDMP to back up data to multiple tape devices simultaneously and improves the backup performance. We recommend starting with two streams for the multi-stream backup. If there are too many attempts, multiple NDMP jobs lock onto the same snapshot. This causes locking contention.
- All PowerScale nodes support NDMP. Considering that the archive nodes have fewer CPU resources, we recommend against running NDMP jobs on archive nodes which can impact other workloads such as SMB or NFS user activity.
- Backing up **/ifs** is not a supported configuration.
- The maximum number of client sessions is 3,000 per node. If the number of client sessions exceeds 3000, it can result in degraded performance.
- NDMP MSBRE is not supported in releases prior to OneFS 9.1.0.0.
- Consider the following NDMP limitations when using NDMP backups.
 - NDMP supports block sizes up to 512 KB. We recommend setting it to a maximum of 256 KB.
 - NDMP does not support more than 4 KB file path length.
 - NDMP operates best with a directory depth at 15 or fewer directories.
 - NDMP does not back up OneFS configuration data, such as file protection level policies and quotas.
 - NDMP does not support recovering data from a file system other than OneFS.
 - Fibre Attached Storage nodes cannot interact with more than 4096 tape paths.
 - The maximum length of the **FILESYSTEM** NDMP environment variable supported for a backup operations is 1024.

NDMP and OneFS feature integration

Overview

NDMP seamlessly integrates with OneFS core capabilities and features. These features include snapshots, CloudPools, SmartConnect, and SmartLock which are described in this section.

SnapshotIQ

Dell PowerScale SnapshotIQ can take read-only, point-in-time copies (snapshots) of any directory or subdirectory within OneFS. It also supports up to 1,024 snapshots in any single directory.

For NDMP backups, OneFS takes a snapshot of the backup directory to ensure consistency of data and the NDMP backup operates on the snapshot instead of the

source directory. The snapshot-based incremental backup uses snapshots to identify files to be backed up.

You are not required to activate a SnapshotIQ license on the PowerScale cluster to perform NDMP backups. If you have activated a SnapshotIQ license on the PowerScale cluster, you can generate a snapshot and back up that snapshot using NDMP. If you back up a SnapshotIQ snapshot, OneFS does not create another snapshot for the NDMP backup.

CloudPools

PowerScale CloudPools allows tiering cold or infrequently accessed data to lower-cost cloud storage. After file data has been archived to the cloud storage, the file is truncated to an 8 KB file. The 8 KB file is called a SmartLink file. CloudPools allows applications and users to seamlessly retain access to data through the same network path and protocols, regardless of where the file data physically resides.

OneFS NDMP is CloudPools-aware and supports three backup methods for CloudPools:

- **Combo copy backup:** Combo copy backup is introduced in OneFS 8.2.0. This option is used to back up files as SmartLink files with file data. Files can be restored as regular files or SmartLink files. This option is the default setting in OneFS 8.2.0 and higher.
- **Deep copy backup:** This option is used to back up files as regular files or unarchived files. Files can only be restored as regular files.
- **Shallow copy backup:** This option is used to back up files as SmartLink files without file data. Files can only be restored as SmartLink files. This option is the default setting before OneFS 8.2.0.

Note: Deep copy and combo copy backups retrieve file data from the cloud. The file data is not cached or stored on disks during the backups. Retrieving file data from the cloud may incur charges from cloud vendors, and the backup performance may degrade. Shallow copy backup backs up SmartLink files without file data and improves the backup performance.

OneFS NDMP provides the following capabilities with CloudPools:

- Restoring SmartLink files or regular files to the same PowerScale cluster.
- Restoring SmartLink files or regular files to another PowerScale cluster. After SmartLink files are restored and detected on the target cluster, OneFS automatically configures the cloud storage account, CloudPools, and the file pool policy matching the source cluster configuration.
- Backing up version information with each SmartLink file and restoring the SmartLink file after verifying the version compatibility on the target PowerScale cluster.

Note: NDMP backs up the relevant information along with SmartLink files, such as CloudPools version information, account information, encryption keys that are associated with the SmartLink file. NDMP backs up the cache of SmartLink files before OneFS 8.2.0. NDMP does not back up the cache of SmartLink files in OneFS 8.2.0 and later.

The following table shows the NDMP supported restore use cases when running a different version of CloudPools on the source and target PowerScale clusters.

Table 5. NDMP supported restore use cases with a different version of CloudPools

Source cluster	Target cluster	Shallow copy NDMP	Deep copy NDMP
CloudPools 1.0 in OneFS 8.0.x/OneFS 8.1.x	CloudPools 2.0 in OneFS 8.2.0 or higher	Supported	Supported
CloudPools 2.0 in OneFS 8.2.0 or higher	CloudPools 1.0 in OneFS 8.0.x/OneFS 8.1.x	Not Supported	Supported

You can specify how SmartLink files are backed up and restored by setting the NDMP environment variables `BACKUP_OPTIONS` and `RESTORE_OPTIONS`. For details about how to set and manage NDMP environment variables, see the *PowerScale OneFS CLI Administrator Guide*.

CloudPools 2.0 (SmartLink files version) is backward compatible with CloudPools 1.0. However, CloudPools 1.0 is incompatible with CloudPools 2.0. The following table shows the NDMP supported backup options and restore options with CloudPools.

Table 6. NDMP supported backup options and restore options with CloudPools

Backup option	Restore option	SmartLink files version compatible	SmartLink files version incompatible
Combo copy backup	Shallow copy restore	SmartLink files are restored on the target cluster.	Not supported
	Deep copy restore	Regular files are restored on the target cluster.	Regular files are restored on the target cluster.
Shallow copy backup	Shallow copy restore	SmartLink files are restored on the target cluster.	Not supported
Deep copy backup	Deep copy restore	Regular files are restored on the target cluster.	Regular files are restored on the target cluster.

We recommend using the combo copy backup option. When the combo copy backup option is used for backup, you can use the shallow copy, or deep copy restore options to recover SmartLink files. For the NDMP backup and restore options, the considerations include the following:

- If you do not specify any backup option, NDMP backs up SmartLink files using the combo copy backup option by default.
- If you do not specify any restore operation, NDMP restores SmartLink files from combo copy backups using the combo copy restore option by default. For the combo copy restore, you can use shallow copy restore option or deep copy restore option by setting the `RESTORE_OPTIONS` environment variable.
- If the version check operation on the target cluster is successful, the shallow copy restore operation restores SmartLink files as SmartLink files on the target cluster.

NDMP and OneFS feature integration

- If the version check operation on the target cluster fails, the deep copy restore operation forces the recovery of the SmartLink files as regular files on the target cluster.

OneFS allows users to update the file data through SmartLink files and writes back the updated data to the cloud storage. Multiple versions of SmartLink files can be backed up to tapes using NDMP, and multiple versions of cloud data objects are protected on the cloud under the data retention setting. You can restore a specific version of a SmartLink file from tapes to a PowerScale cluster and continue to access to the restored SmartLink files.

Data retention is a concept used to determine how long to keep cloud objects on the cloud storage. There are the following three different retention periods in CloudPools setting:

- **Cloud data retention period:** This option is used to specify the length of time cloud objects are retained after the files have been fully recalled or deleted. The default value is one week.
- **Incremental backup retention period for NDMP incremental backup and SyncIQ:** This option is used to specify the length of time that CloudPools retains cloud objects referenced by a SmartLink file. Also, SyncIQ replicates the SmartLink file or NDMP backs up the SmartLink file using an incremental NDMP backup. The default value is five years.
- **Full backup retention period for NDMP only:** This option is used to specify the length of time that OneFS retains cloud data referenced by a SmartLink file. Also, NDMP backs up the SmartLink file using a full NDMP backup. The default value is five years.

If more than one period applies to a file, the longest period is applied. CloudPools checks the retention periods of SmartLink files and ensure that SmartLink files are valid. CloudPools disallows restoring invalid SmartLink files. When restoring SmartLink files, the retention period must be set appropriately to ensure that SmartLink files are valid. If the retention time of a SmartLink file is past the restore time, CloudPools prevents NDMP from restoring that SmartLink file.

Note: Do not delete a cloud storage account that is in use by archived files. NDMP backup and restore of SmartLink files will fail when the cloud storage account has been deleted.

SmartConnect

PowerScale SmartConnect acts as a DNS delegation server to return IP addresses for SmartConnect zones, generally for load-balancing connections to the PowerScale cluster. SmartConnect can continue to provide data availability without any user or application interruption if there is a PowerScale node failure or preventative maintenance.

Note: Consult your DMA vendor documentation check if DMA supports the PowerScale SmartConnect feature for NDMP backup and restore operations.

SmartConnect allows the filer to distribute the NDMP sessions across the PowerScale nodes or the backup accelerator nodes. The filer configures and controls the distribution algorithm. SmartConnect can look up for both single and multi-stream NDMP sessions.

Note the following considerations regarding NDMP and SmartConnect integration:

- The NDMP two-way backup solution with SmartConnect requires Fibre Attached Storage node for backup and restore operations.
- The NDMP three-way backup solution with SmartConnect does not require Fibre Attached Storage nodes for backup and restore operations.
- For the NDMP two-way backup solution with SmartConnect, you can initiate the backup session on DMA through a dedicated SmartConnect zone. This zone consists of a pool of NICs on the Fibre attached storage nodes.
- For NDMP two-way backup without SmartConnect, you can initiate the backup session on DMA through a static IP address or FQDN of the Fibre attached storage node.
- For the NDMP three-way backup solution, the front-end Ethernet network or the network interfaces of the PowerScale nodes are used to serve the backup and restore traffic. We recommend initiating NDMP sessions on DMA through the PowerScale nodes that are only used for NDMP operations. This configuration ensures that it does not compete with things that are performance sensitive like SMB and NFS user activity.
- For the NDMP three-way backup solution with or without SmartConnect, you can initiate the NDMP backup session on DMA using the IP addresses of the PowerScale nodes. The IP addresses can be identified and connected for running the NDMP sessions.

SmartLock

PowerScale SmartLock is used to protect critical data from unauthorized alteration. SmartLock allows you to commit files to a write once, read many (WORM) state, which prevents users from erasing or rewriting those files.

You can back up WORM files through NDMP. SmartLock retention settings are retained through the backup and restore processes, so you are not required to recommit files after a restore.

Troubleshooting and logs

Overview

An NDMP environment includes backup applications, file servers, and tape devices. These three components are connected over networks. Failure of each part will cause the NDMP session to fail.

This section describes various NDMP troubleshooting methodologies, which include:

- DMA logs
- OneFS NDMP logs

DMA logs

Consult your DMA vendor documentation for the logs of NDMP backup and restore operations.

OneFS NDMP logs

Check the OneFS NDMP log if needed. The location of NDMP logs is as follows:

- OneFS NDMP error log is at `/var/log/isi_ndmp_d` and `/var/log/messages`

Troubleshooting and logs

- The directory of NDMP trace logs is at **`/ifs.ifsvar/modules/ndmp/trace`**
- The directory of NDMP contexts is at **`/ifs.ifsvar/modules/ndmp/contexts`**
- The directory of NDMP file history is at **`/ifs.ifsvar/modules/ndmp/filehistory`**

References

Dell Technologies documentation

The following Dell Technologies documentation provides other information related to this document. Access to these documents depends on your login credentials.

- [PowerScale Info Hub](#)