

Quest® NetVault® Backup Plug-in *for NDMP*
11.4

Application Notes for Dell FluidFS



© 2017 Quest Software Inc.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, Join the Innovation, and NetVault are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Legend

- **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.

- ⚠ **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

- i **IMPORTANT NOTE, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Contents

Introduction	4
About Dell FluidFS	4
Key benefits	5
About this document	5
Target audience	5
Recommended additional reading	5
Installation and configuration prerequisites	6
Use cases	6
Installation prerequisites	8
Configuration prerequisites	8
Authentication	8
Local device topology	9
Adding the Dell FluidFS NAS	10
Add devices in NetVault Backup	10
Assumptions	10
Backing up data	11
About backup dump levels	11
Important notes	11
Backing up data using the plug-in	12
Backup content filtering	14
Backing up replica NAS containers	15
Backing up individual data paths	15
Restoring data	17
Restoring data using the plug-in	17
Restoring Incremental Backups	19
Renaming or relocating data	19
Restoring data to an alternate filer	20
Prerequisites	20
Restore procedure	20
Restoring data from or to a Replica NAS Container	20
Searching for files in savesets	21
Viewing media list	21
Troubleshooting	22
Common errors	22
About us	24
We are more than just a name	24
Our brand, our vision. Together.	24
Contacting Quest	24
Technical support resources	24

Introduction

- [About Dell FluidFS](#)
- [About this document](#)
- [Target audience](#)
- [Recommended additional reading](#)

About Dell FluidFS

Dell FluidFS is a proprietary, scaled-out, distributed file system that adds file services to the Dell EqualLogic, Dell Compellent, and Dell PowerVault storage product lines. (FluidFS running on Dell PowerVault NX3500/36x0, Dell EqualLogic FS7500/76x0, and Dell Compellent FS8600 NAS Appliances is also referred to as firmware.)

The FluidFS architecture is open-standards-based, supports industry-standard protocols, and provides innovative features that relate to high availability, performance, efficient data management, data integrity, and data protection. As a core component of the Dell Fluid Data architecture, FluidFS brings differentiated value to the various Dell storage offerings. It is a network-attached storage (NAS) file system that is accessed by using Common Internet File System (CIFS) and Network File System (NFS) protocols.

The high-availability and data-integrity capabilities of FluidFS provide a solid foundation for a strong data-protection strategy. Additional elements built on this foundation include snapshots, replication, and Network Data Management Protocol (NDMP)-based backups, which protect data against localized failures, outages, and disasters.

- **Snapshots:** Fast, point-in-time, read-only copies of NAS volumes that users can access to recover files or directories.
- **Replication:** Local or remote disk copy that administrators can use for disaster recovery or to recover data if a service interruption occurs.
- **NDMP-based backups:** Full data backups of NAS storage using NDMP-compliant Data Management Application (DMA) devices and backup devices such as tape drives and disks.
- **Two-way NDMP backup (Direct NDMP):** – Enables directly attaching FC tape to the FC switch used by FluidFS, and not through a DMA server. The backup data is transferred directly across Fibre Channel to the attached backup target without traveling across the LAN. Only the backup control data travels across the LAN from the NDMP client with backup software.

NOTE: iSCSI solutions and direct connection to a single controller are not supported.

The Quest® NetVault® Backup Plug-in *for NDMP* (Plug-in *for NDMP*) integrates with NetVault Backup so that you can maximize the benefits of NDMP. The Plug-in *for NDMP* enhances backup performance and uses data transfer over the network, while maintaining centralized backup administration, and providing fast, online backup and restore of data. Available for UNIX, Windows, and Linux, the plug-in, combined with the Dell NAS FluidFS architecture, provides a full-featured solution that includes replication, backup, recovery, automation control, direct access restores, and Dynamically Shared Devices. It provides central management and an easy-to-use interface, which simplifies deployment.

This document covers the scenarios and configuration options for using this integrated solution.

Key benefits

- Unified IP storage solution (FC, iSCSI, CIFS, NFS)
- Asynchronous, file-system-level, snapshot-based replication to peer NAS Series systems
- User-restorable snapshots
- NDMP-based backups

About this document

This document provides information about using the Plug-in *for NDMP* with Dell FluidFS. It is intended as a supplement to the *Quest NetVault Backup Plug-in for NDMP User's Guide* that describes the common procedures for installing and configuring the plug-in.

Target audience

This document is intended for system administrators and others responsible for installing, configuring, and using the Plug-in *for NDMP*. An understanding of filer administration and the host platform is assumed.

Recommended additional reading

- NetVault Backup documentation
 - *Quest NetVault Backup Installation Guide*: This guide provides details on installing the NetVault Backup Server and Client software.
 - *Quest NetVault Backup Administrator's Guide*: This guide explains how to use NetVault Backup and describes the functionality common to all plug-ins.
 - *Quest NetVault Backup CLI Reference Guide*: This guide provides a description of the command-line utilities.
 - *Quest NetVault Backup Plug-in for NDMP Application Notes*: These notes provide filer-specific information.

You can download these guides from <https://support.quest.com/technical-documents>.

- Dell FluidFS documentation: You can download the FluidFS documentation from <http://www.dellstorage.com/>.
- Dell FluidFS version 5.0 Administrator's Guide: You can download the Dell FluidFS version 5.0 Administrator's Guide from <http://www.dell.com/support/manuals/us/en/04/dell-compellent-fs8600/FluidFS-V5.0-Admin-Guide>.

Installation and configuration prerequisites

- [Use cases](#)
- [Installation prerequisites](#)
- [Configuration prerequisites](#)
- [Adding the Dell FluidFS NAS](#)
- [Add devices in NetVault Backup](#)
- [Assumptions](#)

Use cases

There are two primary groups of protection options for use with the NetVault Backup products and the Dell FluidFS NAS architecture:

- The first option is to use NetVault Backup as a configuration item providing services for data protection operations. This option lets you implement a simple method for data protection and recovery scenarios. You can store data on a disk-based target, such as the DR4000 Series appliance or a traditional tape library.
- The second option lets you use NetVault Backup as a configuration item that provides services for data-protection operations that complement NAS server replication. This option lets you implement a simple method for data protection and recovery scenarios, while providing extra availability options. As with the first option, you can store data on a disk-based target, such as the DR4000 Series appliance, or a traditional tape library.

The following list provides examples of these two protection options:

- Backing up one or more NAS Containers
- Backing up NAS Containers from multiple Dell NAS clusters
- Backing up Replica Containers
- Restoring backups to a different NAS Container
- Restoring or relocating (or both) backups to a different Dell NAS cluster
- Restoring backups from a Replica Container
- Include or exclude certain paths for NDMP backup and restore operations

Figure 1. Backing up multiple NAS Containers

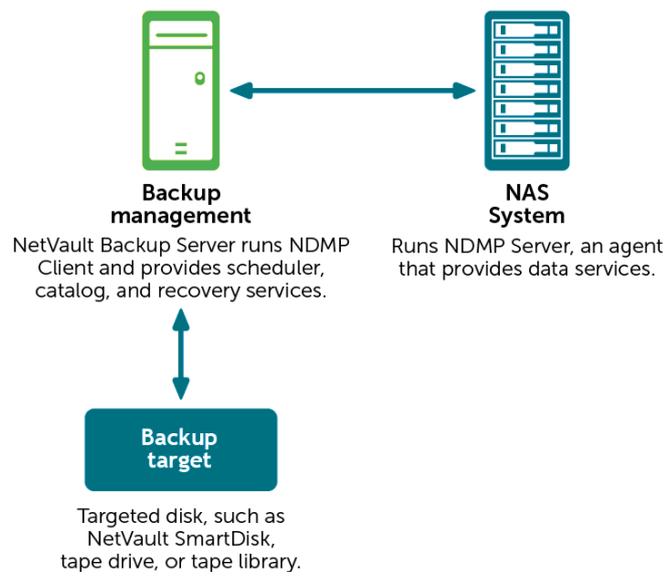


Figure 2. Backing up Replica NAS Containers

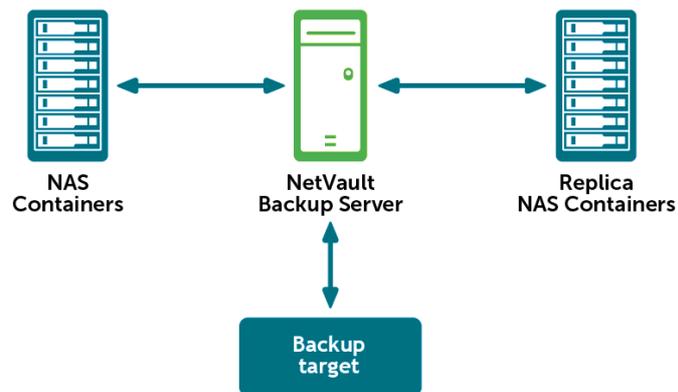
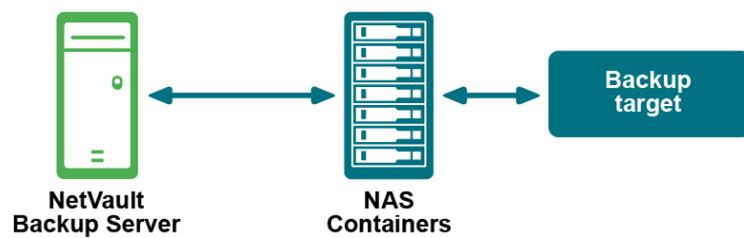


Figure 3. Backing up NAS Containers using Local Topology



Installation prerequisites

In addition to the installation prerequisites covered in the *Quest NetVault Backup Plug-in for NDMP User's Guide*, verify that the following requirements are met:

- **Hardware requirements:** For information about FluidFS Servers, see the Dell website at: <http://www.dellstorage.com/WorkArea/DownloadAsset.aspx?id=1578>
- **Software requirements:** For information about software requirements and supported filer operating systems, see the *Quest NetVault Backup Compatibility Guide* available at: <https://support.quest.com/technical-documents>

i | **NOTE:** The installation and configuration procedures are described in the *Quest NetVault Backup Plug-in for NDMP User's Guide*.

Configuration prerequisites

In addition to the configuration prerequisites covered in the *Quest NetVault Backup Plug-in for NDMP User's Guide*, verify that the requirements in the following two sections are met.

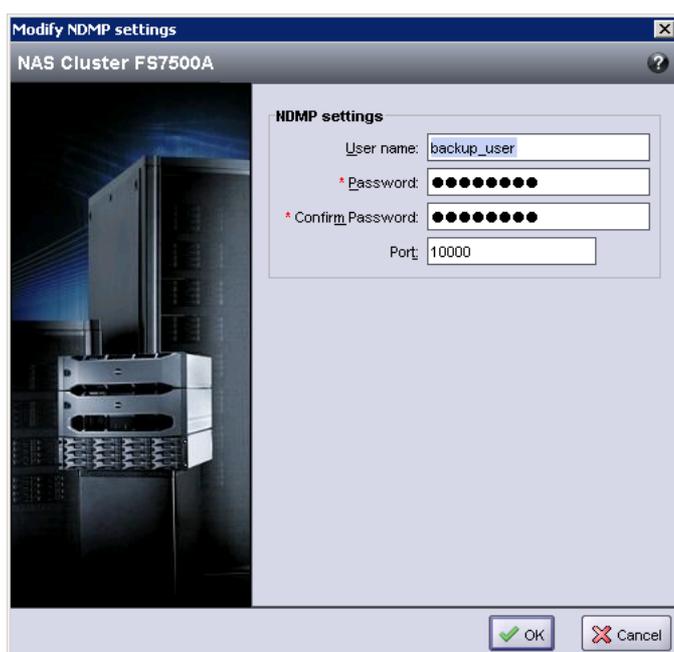
Authentication

Verify that the following authentication requirements are met:

- **NDMP Service:** On the FluidFS Server, ensure that the NDMP services are started.
- **User account for NDMP backups and restores:** Assign a password to the default NDMP user (backup_user) or create a user on the filer using either the FluidFS Group Manager or CLI. Consult the relevant filer documentation for instructions on creating a user account.

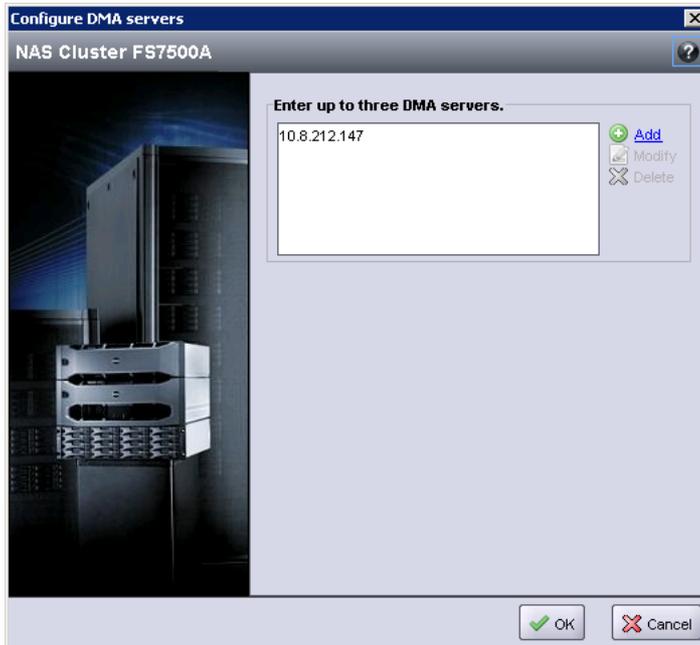
Log on with this user account when adding the NDMP Server or reconfiguring the plug-in authentication details.

Figure 4. Modify NDMP settings dialog box



- **Networking and infrastructure:** Verify that the NAS Cluster IP Addresses are preconfigured so that you can use them when you add the Dell FluidFS NAS to the Plug-in *for NDMP*.
 - **NOTE:** The Dell FluidFS NAS also assigns IP addresses to each controller in the cluster. Do not use these IP addresses when you add the Dell NAS to the plug-in.
- **Add the NetVault Backup Server to the Dell NAS:** In the **Configure DMA Servers** dialog box for Dell NAS, add the IP address of the NetVault Backup Server.

Figure 5. Configure DMA servers dialog box



Local device topology

This section describes local device topology configuration requirements and guidelines.

Before configuring or adding the NAS Filer to the NetVault Backup application, verify that the devices are attached and visible on all nodes in the cluster.

Example:

Verify that all aliases have been set and the operational status is “Optimal” by running the command “`system data-protection tape-devices list`”

Reference - Display all device alias:

```
CLI> system data-protection tape-devices list --
Name,Physical ID,Device Type,Status,Status per Controller
STK_L180,media:STK:L180:7UUDY0030W:0x2100000d7709a3dc,Library,Optimal
STK_L180D1,tape:HP:Ultrium_5-SCSI:7UUDY0030X:0x2100000d7709a3dc,Tape,Optimal
STK_L180D2,tape:HP:Ultrium_5-SCSI:7UUDY0030Y:0x2100000d7709a3dc,Tape,Optimal
STK_L180D3,tape:HP:Ultrium_5-SCSI:7UUDY0030Z:0x2100000d7709a3dc,Tape,Optimal
STK_L180D4,tape:HP:Ultrium_5-SCSI:7UUDY00310:0x2100000d7709a3dc,Tape,Optimal
STK_L180D5,tape:HP:Ultrium_5-SCSI:7UUDY0032G:0x2100000d7709a3dc,Tape,Optimal
STK_L180D6,tape:HP:Ultrium_5-SCSI:7UUDY0032H:0x2100000d7709a3dc,Tape,Optimal
STK_L180D7,tape:HP:Ultrium_5-SCSI:7UUDY0032I:0x2100000d7709a3dc,Tape,Optimal
STK_L180D8,tape:HP:Ultrium_5-SCSI:7UUDY0032J:0x2100000d7709a3dc,Tape,Optimal
STK_L180D9,tape:HP:Ultrium_5-SCSI:7UUDY0032K:0x2100000d7709a3dc,Tape,Optimal
STK_L180DA,tape:HP:Ultrium_5-SCSI:7UUDY0032L:0x2100000d7709a3dc,Tape,Optimal
```

If the aliases are not set or need modification, see the *Dell FluidFS version 5.0 Administrator's Guide*.

The examples provided here are for your reference; these examples are limited and noninclusive. For queries and issues, contact your FluidFS Support team.

Reference - Display all discovered devices:

```
CLI> system data-protection tape-devices discovered-list -
```

Reference - Add a device alias:

Changer:

```
CLI> system data-protection tape-devices add media:STK:L700::0x500065b152501600  
STK_L700P1
```

Tape:

```
CLI> system data-protection tape-devices add tape:IBM:ULT3580-  
TD4:9I0E1Z_01_:0x500065b152501600 STK_L700D1P1
```

Adding the Dell FluidFS NAS

You can configure the Dell FluidFS to use one or more NAS Cluster IP addresses. When you add the Dell NAS Server to the Plug-in for NDMP, use any one of the NAS Cluster IP addresses.

Add devices in NetVault Backup

On the NetVault Backup Server machine, add the library attached to the filer to the NetVault Backup Server to configure using the local topology. Ensure that the devices added are relative to the FluidFS.

To add a device, see the *Quest NetVault Backup Administrator's Guide*.

Assumptions

It is assumed that the users implementing the procedures described in this document have expertise in the following areas:

- NetVault Backup installation, configuration, and administration.
- FluidFS Server configuration and administration.
- Intermediate networking expertise, including routing, name resolution, and TCP/IP.
- Intermediate NFS and CIFS Protocol setup and administration expertise.
- Intermediate Windows Authentication and Active Directory administration expertise.

Backing up data

- [About backup dump levels](#)
- [Backing up data using the plug-in](#)
- [Backup content filtering](#)
- [Backing up replica NAS containers](#)
- [Backing up individual data paths](#)

About backup dump levels

The Plug-in *for NDMP* supports both Full and Incremental Backups. The backup type is determined by the Dump Level that can be set to Level 0 for a Full Backup or Levels 1 through 9 for Incremental Backups:

- **Level 0 or Full Backup:** A Full Backup provides a backup of all the data in the selected path.
- **Level 1 through 9 or Incremental Backup:** The Incremental Backups base themselves on the most recent lower-level Incremental Backup, and include any data that has changed or is new since the last Full or Incremental Backup.

Examples:

- Level 1 Incremental Backup includes any data that has changed or is new since the last Level 0 or Full Backup.
- Level 2 Incremental Backup includes any data that has changed or is new since the last Level 1 Incremental Backup, and so on, up to dump level 9.

Important notes

- If a directory is included in a dump level 0 backup, it is included in all future dumps even if no changes have occurred. In this instance, the directory is empty and does not contain content. However, it does not negatively affect the recovery of data.
- Only one dump level is maintained for each selected set of data. For example, after performing Level 1, Level 2, and Level 3 Incremental Backups, if you perform a Full Backup, the dump level is reset to 0. The next instance of the Incremental Backup is based on the Full Backup.
- If an Incremental Backup fails, the subsequent backup is based on the most recent backup of a lower level. For example, if dump level 1 is completed successfully, but dump level 2 fails, dump level 3 is based on dump level 1. The failure of an Incremental Backup is noted in the NetVault Backup logs.

Backing up data using the plug-in

- 1 In the Navigation pane, click **Create Backup Job**.

— or —

In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Create Backup Jobs**.

- 2 In **Job Name**, specify a name for the job.

Assign a descriptive name that lets you easily identify the job for monitoring its progress or restoring data. The job name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

- 3 In the **Selections** list, select an existing Backup Selection Set, or complete the following steps to create a set.

- a Next to the **Selections** list, click **Create New**.
- b On the **NetVault Backup Selections** page, open the NetVault Backup Server, and then open **NDMP Client**.
- c To display the **Volumes** node, open the applicable filer in the list of added servers.
- d Select the data that you want to back up:
 - **Back up all NAS Containers:** Select the **Volumes** node.
 - **Back up individual NAS Containers:** Open the **Volumes** node, and select the NAS Containers that you want to back up.
- e Click **Save**, and in the **Create New Set** dialog box, type a name for the set.

i **NOTE:** A set name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Linux OS, the names can have a maximum of 200 characters. On Windows OS, there is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

- f To close the dialog box, click **Save**.

- 4 In the **Plugin Options** list, select an existing Backup Options Set, or complete the following steps to create a set.

- a Next to the **Plugin Options** list, click **Create New**.
- b Under **NDMP Backup Options**, configure the following settings.

Table 1. NDMP Backup Options

Option	Description
Dump Level	Select the type of backup that you want to perform. For a Full Backup, type or select 0. For an Incremental Backup, type or select any level from 1 to 9, based on the previous dump level.
Do Direct Backup	To use Direct Access Recovery (DAR) during restores, leave this check box selected. DAR enables precise tape positioning, which allows quick recovery of individual files. Without DAR, the backup saveset is read sequentially to restore files.
Snapshot	This variable controls whether a snap shot is taken for the backup.

Table 1. NDMP Backup Options

Option	Description
List of Include Paths	<p>This environment variable specifies a list of comma-separated paths to be included in the backup operation. The path specified could be a directory or a file. If the path is a directory, all child elements of that directory are included in the backup. Each path specified is relative to the backup root directory (as specified by the FILESYSTEM environment variable) and must start with the "/" character. Each path can be a maximum of 1 K bytes long. The last element of any path can contain a wildcard character ("*") at the beginning or at the end.</p> <p>A maximum of 32 paths can be specified.</p> <p>If INCLUDE_PATH is specified, SUBTREE_NAME must also be specified.</p> <p>If the specified INCLUDE_PATH does not exist, the backup fails.</p>
List of Exclude Paths	<p>This environment variable specifies a list of comma-separated paths to be excluded in the backup operation. The path specified could be a directory or a file. If the path is a directory, all child elements of that directory are excluded in the backup. Each path specified is relative to the backup root directory (as specified by the FILESYSTEM environment variable) and must start with the "/" character. Each path can be a maximum of 1 K bytes long. The last element of any path can contain a wildcard character ("*") at the beginning or at the end.</p> <p>A maximum of 32 paths can be specified.</p> <p>If EXCLUDE_PATH is specified, SUBTREE_NAME must also be specified.</p> <p>If the exclude path is not within the include path, it is ignored.</p>
Set Sub Tree Name for Include/Exclude Paths	<p>This environment variable must be specified if INCLUDE_PATH or EXCLUDE_PATH is specified. The specified name is used to track the time stamp of the backups at each backup level to handle incremental and differential backups. The environment variable must be a string of no more than 64 characters.</p>
File and Directory Exclusion	<p>Specifies a matching pattern for directory and filenames that are not to be backed up. The environment variable is a list of string separated by comma. Each entry is used to match to nodes encountered during backup. The string can contain an asterisk ("*") as the wildcard character but the asterisk must be the first or last character of the pattern. At most, 32 comma-separated strings are supported.</p>
Save File Information	<p>This option allows you to browse individual files and directories while selecting data during restore. You cannot restore individual files and directories if you clear this check box.</p> <p>Quest recommends keeping this check box selected.</p>

- c The Plug-in *for NDMP* includes the ability to generate multiple jobs for each container selected in a backup set. Quest recommends that you set the backups to run in parallel so that it is equal to the number of drives (virtual or physical) targeted for the data-protection backup operation.

To run parallel jobs, configure the following settings under **Sub Jobs**.

Table 2. Sub Jobs

Option	Description
Number of Backups to Run in Parallel	Type or select the number of jobs that can run concurrently. The default value is zero. With the default setting, all jobs in parallel, limited only by the availability of drives, network bandwidth, and other resources. To achieve the required result, set the same value for each job.
Delete Auto Generated Backup Jobs on Completion	The plug-in creates individual subjobs for each volume included in a backup. If this check box is not selected, all subjobs are listed on the Job Status page. This display can create confusion when you try to view the details of the parent job. To include only a single entry for the parent job, leave this check box selected.

- d Click **Save**, and in the **Create New Set** dialog box, type a name for the set.
 - e To close the dialog box, click **Save**.
- 5 Select or create the Schedule Set, Target Set, and Advanced Options Set.
For more information about these sets, see the *Quest NetVault Backup Administrator's Guide*.
- 6 To submit the job for scheduling, click **Save & Submit**. You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page.
To save the job definition without scheduling it, click **Save**. You can view, edit, or run this job from the **Manage Job Definitions** page. This job is not displayed on the **Job Status** page until you submit it.
For more information about **Job Status**, **View Logs**, and **Manage Job Definitions**, see the *Quest NetVault Backup Administrator's Guide*.

Backup content filtering

NDMP include and exclude paths add an ability to include or exclude certain paths for NDMP backup and restore operations. While defining a backup using DMA, you can select specific directories from the virtual NAS volume to include in, or exclude from, NDMP backup jobs.

During backup, you can filter the content using the following NDMP environment variables.

- **INCLUDE_PATH**: Specifies a list of paths to include. The paths are relative to the backup root.
- **EXCLUDE_PATH**: Specifies a list of paths to exclude. The paths are relative to the backup root.
- **EXCLUDE**: Specifies a list of filenames and patterns to exclude.

The NDMP server first examines INCLUDE_PATH and only traverses to the list of included paths. Then EXCLUDE_PATH and EXCLUDE are evaluated to further filter the content. The difference between EXCLUDE_PATH and EXCLUDE is that EXCLUDE_PATH is applied only at the specified path while EXCLUDE is applied to all elements. For example,

```
EXCLUDE: a*  
EXCLUDE_PATH: /home/david/b*
```

In this case, all directories and files starting with “a” are excluded. Directories and files under /home/david starting with “b” are excluded.

The EXCLUDE environment variable specifies a name for filtering and a wildcard character at the beginning or end of the pattern. For example, you can specify “*.o” to filter out all object files during backup. The same rule is applicable to the last path element of INCLUDE_PATH and EXCLUDE_PATH. For example, if INCLUDE_PATH is “/home/a*”, then all directories and files under /home that starts with “a” are backed up.

The NDMP server follows a general rule that prioritizes INCLUDE_PATH. So, if something is specified by INCLUDE_PATH, it gets backed up even if EXCLUDE_PATH or EXCLUDE excludes it. The following are a few examples:

Example 1

```
INLCUDE_PATH: /home/david/a*
EXCLUDE_PATH: /home/david
```

In this case, /home/david and all child nodes starting with “a” are backed up even though EXCLUDE_PATH has excluded /home/david.

Example 2

```
INLCUDE_PATH: /home/david/a*
EXCLUDE: david
```

In this case, /home/david and all child nodes starting with “a” are backed up even though EXCLUDE has excluded all elements name “david”.

Example 3

```
INLCUDE_PATH: /home/david/a*
EXCLUDE: *c
```

In this case, /home/david and all child nodes starting with “a” are backed up except for child nodes ending with a “c”. So “/home/david/aaa” are backed up while “/home/david/abc” are excluded.

Exception:

There is an exception to the preceding rules. If the last element of a INCLUDE_PATH does not contain a wildcard character and the named element is a directory, it is backed up no matter what EXCLUDE_PATH and EXCLUDE is set to. For example,

```
INLCUDE_PATH: /home/david/abc
EXCLUDE: *c
```

In this case, if “/home/david/abc” is a directory, it is backed up even though it matches EXCLUDE’s “*c”. However, if “/home/david/abc” is not a directory, it is excluded.

Backing up replica NAS containers

- 1 Complete [Step 1](#) through [Step 3](#) in [Backing up data using the plug-in](#).

i | **NOTE:** The Replicated Container must either be promoted to a Local Container or it must be a Read-Only Replica Container.

- 2 Next to the **Plugin Options** list, click **Create New**.
- 3 Configure the **Sub Job** options.
For more information, see [Sub Jobs](#).
- 4 Complete [Step 5](#) and [Step 6](#) in [Backing up data using the plug-in](#).

Backing up individual data paths

- 1 In the Navigation pane, click **Create Backup Job**.

— or —

In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Create Backup Jobs**.

- 2 In **Job Name**, specify a name for the job.

Assign a descriptive name that lets you easily identify the job for monitoring its progress or restoring data. The job name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin

characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

- 3 Next to the **Selections** list, click **Create New**.
- 4 On the **NetVault Backup Selections** page, open the NetVault Backup Server, and then open **NDMP Client**.
- 5 Select the applicable filer, and in the **Actions** list, click **Enter Backup Path**.
- 6 In the **NDMP Backup Request** dialog box, configure the **Directory** option.

Specify the volume or the complete path to the directory that you want to back up. Use the following format to specify the backup path:

```
/<NAScontainer>/<parentDirectory>/<subdirectory>
```

Example:

```
/Vol1/testdata/small/data
```

- 7 Set the **Dump Level** and the remaining backup options.
For more information, see [NDMP Backup Options](#).
- 8 To save the settings, click **OK**.
- 9 Complete [Step 5](#) and [Step 6](#) in [Backing up data using the plug-in](#).

Restoring data

- Restoring data using the plug-in
- Restoring Incremental Backups
- Renaming or relocating data
- Restoring data to an alternate filer
- Restoring data from or to a Replica NAS Container
- Searching for files in savesets
- Viewing media list

Restoring data using the plug-in

- 1 In the Navigation pane, click **Create Restore Job**.

On the **Create Restore Job — Choose Saveset** page, the saveset table provides a list of available savesets. The table shows the saveset name (Job Title and Saveset ID), creation date and time, saveset size, and saveset status.

The saveset status is indicated using the following icons.

Icon	Description
	Saveset is online (all segments are online).
	Saveset is partially online (some segments are online).
	Saveset is offline (all segments are offline).

The saveset list is sorted by creation date (newest to oldest). You can sort the list by one or more columns, as required. The arrowhead next to the column name indicates the sort order.

- 2 To filter the saveset list, use the following filter options.

Filter	Description
Client	Displays savesets created for particular clients. The default selection is Any . <ol style="list-style-type: none"> 1 To use this filter, click the Client box. 2 In the Choose Client dialog box, select the clients. 3 To close the dialog box, click OK.
Plugin Type	Displays savesets created using a particular plug-in. The default selection is Any . <ol style="list-style-type: none"> 1 To use this filter, click the Plugin Type box. 2 In the list, select the plug-in.

Filter	Description
Date	<p>Displays savesets created during a specified period. The default selection is Any.</p> <ol style="list-style-type: none"> To use this filter, click the Date box. In the list, select the option that you want to use. The available options are Last 24 hours, Last Week, Last Month, Last 6 Months, Last Year, and Any.
Job	<p>Displays savesets created for particular jobs. The default selection is Any.</p> <ol style="list-style-type: none"> To use this filter, click the Job box. In the Choose Job dialog box, select the jobs. To close the dialog box, click OK.

- 3 Select the saveset that you want to use, and click **Next**.

When you select a saveset, the following details are displayed in the **Saveset Information** area: job ID number, job title, tag, server name, client name, plug-in name, saveset date and time, retirement setting, Incremental Backup or not, archive or not, and saveset size.

- 4 On the **Create Selection Set** page, select the data that you want to restore:

- **Restore entire saveset:** Select the root node.
- **Restore individual files and directories:** Select the target files and directories in the selections tree. The selections tree can only be browsed if the **Save File Information** option was selected during backup.

i | **NOTE:** If you omit child-level items after selecting the parent node, the job fails. To perform selective restores, open the parent node, and select each item that you want to restore.

- 5 Click **Edit Plugin Options**, and configure the following settings.

Table 3. Restore options

Option	Description
NDMP Server	This option specifies the target filer name. It is only required when you want to restore data to an alternate filer. By default, it is set to the name of the original filer from which the data was backed up. Do not change the NDMP Server while restoring data to the same filer.
Direct Access Restore Mode	<p>Select the appropriate DAR mode from the following:</p> <ul style="list-style-type: none"> • Use Direct Access Restore if Possible: To perform a direct-access restore if the backup was direct and the mover and data server can complete a direct-access restore, select this option. If these conditions are not met, a nondirect access restore is performed. • Use Direct Access Restore by Recovering Individual Files: To restore a directory and its file contents, select this option. When you use this option, empty directories in the saveset are ignored. You must select this option if any item was omitted during data selection; otherwise, the restore job fails. • Only Do Direct Access Restore: To force a direct-access restore, select this option. It can be only used if the backup was direct and the mover and data server can complete a direct-access restore; otherwise, the restore job fails. • Never Do Direct Access Restore: To perform a standard nondirect-access restore, select this option.
Overwrite Existing Files	If you do not want to overwrite the existing files, clear this check box.

i | IMPORTANT:

- When individual files are selected for a direct-access restore, their parent-level directory permissions are not restored.
- When running a direct-access restore, NetVault Backup organizes the restore requests into groups of 1024 files at a time. When a restore job consists of more than 1024 total files, multiple restore requests are issued. So, multiple log entries are displayed for the job.

6 To save the settings, click **OK**, and then click **Next**.

7 In **Job Name**, specify a name for the job.

Assign a descriptive name that lets you easily identify the job for monitoring its progress. The job name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms

In the **Target Client** list, the client from which data was backed up is selected by default. Do *not* change this setting.

8 Select or create the Schedule Set, Restore Source Set, and Advanced Options Set.

For more information about these sets, see the *Quest NetVault Backup Administrator's Guide*.

9 To submit the job for scheduling, click **Submit**.

You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page. For more information about these functions, see the *Quest NetVault Backup Administrator's Guide*.

Restoring Incremental Backups

The incremental restores build on each other. To restore an Incremental Backup, you must have all the backups from Level-0 through the last backup in the backup sequence that you want to restore.

1 Restore the Full or Level 0 Backup.

For more information, see [Restoring data using the plug-in](#).

2 Restore each Incremental Backup in the backup sequence, starting with the lowest-level backup and going to the last backup that you want to restore (that is, in the same order in which they were created).

For each dump level, follow the steps outlined in [Restoring data using the plug-in](#).

i | IMPORTANT: All NDMP backups are effectively snapshots of the volume. When you restore a dump level in an increment series, it restores all the data in the saveset. At the same time, it deletes the files present on the filer, but not available in that saveset. Thus, it reconstructs the subtree as it was at the time of backup.

Renaming or relocating data

When restoring a backup, you can rename the directories to create a copy of the data instead of overwriting of the existing versions. You can also relocate the data to a different volume. This procedure can be useful if the original volume is down or if you are dealing with limited bandwidth or capacity.

i | NOTE: This feature is only available to savesets that were created with the **Save File Information** option.

To rename or relocate data during restore:

1 Complete [Step 1](#) through [Step 4](#) in [Restoring data using the plug-in](#)

2 Select the volume that you want to rename, and in the **Actions** list, click **Rename**.

- 3 In the **Restore Rename** dialog box, provide the following information:
 - To rename a file or directory, type the new name.
 - To relocate the item to a different volume, type the volume name. You can also rename while relocating by specifying a new path. The volume to which you are relocating must exist.

Click **OK** to close the dialog box.

After you close the dialog box, the plug-in updates the corresponding node in the selections tree to display the new name and location for the file or directory.

- 4 Complete [Step 5](#) through [Step 9](#) in [Restoring data using the plug-in](#).

Restoring data to an alternate filer

To relocate a backup to an alternate filer, use the following procedure. This procedure can be useful during a server migration or disaster recovery operation.

Prerequisites

Before you start the restore procedure, add the target filer to the plug-in. For more information, see the *Quest NetVault Backup Plug-in for NDMP User's Guide*.

Restore procedure

- 1 Complete [Step 1](#) through [Step 4](#) in [Restoring data using the plug-in](#).
- 2 To rename or relocate data during restore, complete [Step 2](#) and [Step 3](#) in [Renaming or relocating data](#).
- 3 Click **Edit Plugin Options**, and in the **NDMP Server box**, type the name of the target filer.
- 4 Select the appropriate DAR method. For more information about the available options, see [Restore options](#).
- 5 Complete [Step 6](#) through [Step 9](#) in [Restoring data using the plug-in](#).

Restoring data from or to a Replica NAS Container

You can use the Plug-in *for NDMP* to restore data from or to a Replica NAS Container.

- To restore data from a Replica NAS Container, restore the data to a regular NAS Container, but verify that you select the Rename option.

For more information, see [Renaming or relocating data](#).

i | **NOTE:** If you do not select the Rename option, the restore job fails because the Replica NAS containers are read-only.

- To restore data to a Replica NAS Container, promote it to a Local Container, and then submit the restore job.

Searching for files in savesets

The **Search** option on the **Create Restore Job — Choose Saveset** page allows you to find specific files or data items without opening any savesets or browsing through their contents. You can use filenames or regular expressions to find the data items that you want to restore.

- 1 On the **Create Restore Job — Choose Saveset** page, click **Search**.
- 2 In the **Search for files in savesets** dialog box, configure the following options:
 - **Search String:** Type the search string.
 - **Regular expression search:** To use POSIX (Portable Operating System Interface for Unix) regular expressions in the **Search String** box, select this check box.
 - **Case sensitive:** To perform a case-sensitive search, select this check box.
- 3 Click **Search**. On the **Search Results** page, you can view the savesets that contain the specified files or data items.
- 4 Select the items you want to restore. You can only restore items from one saveset. Click **Next**.
- 5 Complete [Step 5](#) through [Step 9](#) in [Restoring data using the plug-in](#).

Viewing media list

The **Media List** option on the **Create Restore Job — Choose Saveset** page lets you view information about the media items used to store a backup. You can view details about the data segments and index segments for a backup.

- 1 On the **Create Restore Job — Choose Saveset** page, select the applicable saveset.
- 2 In the **Saveset Information** area, click **Media List**.

In the dialog box that is displayed, you can view the following details:

 - **Backup size:** This area shows the total size of the saveset in number of bytes
 - **Data segment table:** This table shows information about the media items that contain the data segments. You can view the following details: media label, media group label, stream ID, starting byte number, ending byte number, and media location
 - **Index segment table:** This table shows information about the media items that contain the index segments. You can view the media label and media location.
- 3 To close the dialog box, click **Close**.

Troubleshooting

- [Common errors](#)

Common errors

The following is a list of some common errors and their solution.

Table 4. Troubleshooting

Description	Symptom	Solution
The NetVault Backup Service fails to start on a Windows-based NetVault Backup Server.	Check the Windows Event Viewer for the following message: PDT FATAL: lock file "postmaster.pid" already exists	NetVault Backup cannot start if the PostgreSQL database that is used to store the system data does not start. To correct this issue, delete the " postmaster.pid " file from the location referenced in the log and restart the NetVault Backup Server.
After restarting the machine, the NetVault Backup Service sometimes fails to start on a Windows-based NetVault Backup Server.	Check the Windows Event Viewer for the following message: FATAL: could not create any TCP/IP sockets " for a PostgreSQL source	NetVault Backup cannot start if the PostgreSQL database that is used to store the system data does not start. To correct this issue, start the Task Manager, and click Show processes from all users . You can see multiple instances of postgres32.exe running on the system. To remove all instances of postgres32.exe , select any instance of this process, and click End Process . Start the NetVault Backup Service.

Table 4. Troubleshooting

Description	Symptom	Solution
The NetVault Backup Service starts, and then stops immediately on a Linux-based machine.	No error messages are displayed.	This issue can occur if the Postgres service cannot resolve the host name localhost , and fails to start. Check the <code>/etc/hosts</code> file, and if the file does not contain an entry for localhost , add the entry.
While restoring a Full Backup for a volume that maintains an existing <code>.clusterConfig</code> folder and for which the Overwrite Existing Files option is not selected (the default), an error occurs.	Check the NetVault Backup logs for the following message: Error 2017/06/01 23:38:06 493 Data Plugin AUTOINT9 NDMP: ERROR 0: ERROR: /mnt/backup/v6vol2/.clusterConfig/backup/c413ec4b-13e7-4ab8-87e6-28c53b478b65@3688372308_0/.BackupIndex0/.BackupIndex0 already exists, not overwriting.	To correct this issue, do either of the following: <ul style="list-style-type: none"> • Delete the <code>.clusterConfig</code> folder, and then run the restore job again. • Select the Overwrite Existing Files option for the restore job to ensure that the folder and its files are overwritten.

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece—you—to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit www.quest.com.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.