

Quest® NetVault® Backup Plug-in *for NDMP*  
12.2

**Application Notes for NetApp ONTAP  
(7.x-7G, 8.0 7-mode, 8.x C-mode, and  
9.5)**



© 2019 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, 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.

NetVault Backup Plug-in for NDMP Application Notes for NetApp ONTAP (7.x-7G, 8.0 7-mode, 8.x C-mode, and 9.5)  
Updated - April 2019  
Software Version - 12.2  
NCG-103-12.2-EN-01

# Contents

<b>Introduction</b> .....	<b>4</b>
About this document .....	4
Target audience .....	4
Recommended additional reading .....	4
<b>Installation and configuration prerequisites</b> .....	<b>5</b>
Installation prerequisites .....	5
Configuration prerequisites .....	5
Changing from NDMP node-scoped mode to NDMP SVM-scoped mode .....	7
Enabling the C-Mode browse extension .....	8
<b>Backing up data</b> .....	<b>9</b>
About Backup Dump levels .....	9
Important notes .....	9
Backing up data in 7.x-7G and 7-mode .....	10
Additional notes: BRE .....	14
Backing up data in C-mode .....	14
Backing up individual data paths in C-mode .....	16
<b>Restoring data</b> .....	<b>17</b>
Restoring data using the plug-in .....	17
Additional notes .....	20
Restoring Incremental Backups .....	20
Renaming or relocating data .....	20
Restoring data to an alternate filer .....	21
Prerequisites .....	21
Restore procedure .....	21
Searching for files in savesets .....	21
Viewing media list .....	22
<b>Troubleshooting</b> .....	<b>23</b>
Upgrading from 10.x or 11.1 to 11.4 .....	23
Common errors .....	23
<b>About us</b> .....	<b>25</b>
Technical support resources .....	25

---

# Introduction

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

## About this document

This document provides information about using the Quest® NetVault® Backup Plug-in *for NDMP* (Plug-in *for NDMP*) with a Network Appliance (NetApp) filer. It is intended as a supplement to the *Quest NetVault Backup Plug-in for NDMP User's Guide*, which 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 User's Guide*: This guide describes the procedures for installing and configuring the Plug-in *for NDMP*.

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

- Administrator's guide for your Network Data Management Protocol (NDMP)-based network-attached storage (NAS) appliance

# Installation and configuration prerequisites

- [Installation prerequisites](#)
- [Configuration prerequisites](#)
- [Enabling the C-Mode browse extension](#)

## 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 NetApp filers, see the NetApp website at: [www.netapp.com](http://www.netapp.com)
- **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>
- **Start the NDMP Daemon:** To start the NDMP Daemon:
  - At the prompt, type: `ndmpd on`
  - Alternatively, to start the daemon automatically at boot time, open the file `/etc/rc` on the filer, and add: `ndmpd on`

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

As described in the *Quest NetVault Backup Plug-in for NDMP User's Guide*, Quest recommends that you install the Plug-in for NDMP on the NetVault Backup Server instead of the client.

## Configuration prerequisites

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

- [Prerequisites for running NDMP backups and restores using the Cluster Aware Backup \(CAB\) extension](#)
- [Prerequisites for using Clustered Data ONTAP 8.2 with NetVault Backup 10.0.1 and earlier](#)
- [User authentication in Data ONTAP 7.x-7G and 7-mode](#)
- [Understanding supported NDMP data connection types](#)

## Prerequisites for running NDMP backups and restores using the Cluster Aware Backup (CAB) extension

NDMP Cluster Aware Backup (CAB) extension enables the NDMP server to provide affinity information about volumes and tape devices to the backup application that support the CAB extension. NetVault Backup uses this affinity information to determine the data transfer path used during backup and restore operations. If a target volume and tape device are on same cluster node, a local backup or restore operation is performed; if they are on different cluster nodes, a three-way backup or restore operation is performed.

To run NDMP backup using the CAB extension, verify that the following requirements are met:

- **Review the software requirements:** To use the CAB extension, you require NetVault Backup 10.0.5 or later, Plug-in for NDMP 10.0.5 or later, and NetApp Data ONTAP 8.2.2 C-Mode or later.

For information about software requirements, see the *Quest NetVault Backup Compatibility Guide* available at: <https://support.quest.com/technical-documents>

- **Set the NDMP Service to Storage Virtual Machine (SVM)-scoped mode on the filer:** Verify that the NDMP service is set to SVM-scoped mode. In Data ONTAP 8.1 and earlier versions, you must explicitly disable the node-scoped mode for NDMP.

Commands:

- Verify the NDMP mode status: `system services ndmp node-scope-mode status`
- Disable node-scoped mode: `system services ndmp node-scope-mode off`

- **Configure an NDMP service user account:** In the SVM-scoped NDMP mode, you must use NDMP Service user account to access a storage system to perform tape backup and restore operations. You must also generate a password for the user account:

Command:

```
vserver services ndmp generate-password vserver <vserverName> -user <userName>
```

**Example:**

```
vserver services ndmp generate-password -vserver vs1 -user admin
```

- **Add the Logical Interface (LIF) to the plug-in:** To add the NDMP Client to the plug-in, use the IP address of the Logical Interface (LIF) and the NDMP service user account.

You can use the command `network interface show` to obtain information about LIFs. The availability of volumes and tape devices depends on the LIF type that you use.

**Table 1. LIF type and availability of volumes and tape devices**

NDMP connection LIF type	Volumes available for backup or restore	Tape Devices available for backup or restore
Cluster Management LIF	All volumes	All tape devices
Node Management LIF	All volumes hosted by the node	Tape devices connected to the node hosting the node management LIF
Data LIF	All volumes that belong to the Vserver	None

For more information about the availability of volumes and tape devices, see the *Clustered Data ONTAP Data Protection Tape Backup and Recovery Guide*.

## Prerequisites for using Clustered Data ONTAP 8.2 with NetVault Backup 10.0.1 and earlier

NetVault Backup 10.0.1 and earlier do not support NDMP in SVM-scoped mode. If you are using NetVault Backup 10.0.1 or an earlier version with Data ONTAP 8.2 C-mode, verify that the following requirements are met:

- **Set the NDMP Service to node-scoped mode:** Verify that the NDMP service is set to node-scoped mode. In a new clustered Data ONTAP 8.2 installation, NDMP is in SVM-scoped mode by default.

Commands:

- Verify the NDMP mode status: `system services ndmp node-scope-mode status`
- Enable node-scoped mode: `system services ndmp node-scope-mode on`
- **Configure NDMP root password on cluster nodes:** In the node-scoped NDMP mode, you must use an NDMP-specific user account to access a storage system to perform tape backup and restore operations. User authentication must be configured on each node. The default user ID is root. Ensure that you change the default NDMP password associated with the root user account.

Commands:

- Set user password: `system services ndmp modify -node * -clear-text false -user-id root -enable true`
- Verify the configuration using the following command: `system services ndmp show`
- **Create the LIF and add the LIF to the plug-in:** To add the NDMP Client to the plug-in, use the IP address of the LIF and the NDMP root user account.

Example commands:

- Create the LIF: `network interface create -vserver vnetapp -lif For_Backup -role cluster-mgmt -home-node vserver-01 -home-port e0d -address 192.168.0.22 -netmask 255.255.255.0`
- Verify the configuration of the LIF: `net int show -fields role,data,data-protocol,vserver,address,curr-node,curr-port`

**i** | **NOTE:** If a volume affinity changes to a different node, you must manually migrate the LIF to the node that is now hosting the volume: `net int migrate -vserver vserver1 -lif For_Backup -dest-node vserver1-02 -dest-port e0d`

## User authentication in Data ONTAP 7.x-7G and 7-mode

To use the NDMP **root** user account for accessing the filer, set a valid password for this account. Password cannot be left blank for this user account. To use a nonroot user, verify that the account has necessary privileges to perform NDMP backups and restores.

## Understanding supported NDMP data connection types

In Vserver-aware NDMP mode, the supported data connection types depend on the NDMP control connection LIF type and the status of the CAB extension. The type of data connection type indicates whether you can perform a local or a three-way NDMP backup or restore operation. For more information, see <https://library.netapp.com/ecmdocs/ECMP1196874/html/GUID-070D947D-F73F-48E2-904E-74CB1CF71E78.html>.

# Changing from NDMP node-scoped mode to NDMP SVM-scoped mode

After changing from NDMP node-scoped mode to NDMP SVM-scoped mode, complete the following tasks:

- Delete the previous NDMP sessions on the filer.
- Add the tape libraries and drives to the NetVault Backup Server again. For more information, see the *Quest NetVault Backup Administrator's Guide*.
- Create the Target Sets again. For more information, see the *Quest NetVault Backup Administrator's Guide*.
- Modify the IP address for the NDMP Client. In the SVM-scoped mode, you must use the IP address of the LIF. For more information about editing the IP address, see the *Quest NetVault Backup Plug-in for NDMP User's Guide*.

- Create the Backup Selection Sets again if there is any change in the NDMP Client name. For more information, see the *Quest NetVault Backup Administrator's Guide*.

## Enabling the C-Mode browse extension

In support of the NDMP protocol extension for Cluster Aware Backup (CAB), you can enable support for the NetApp Data ONTAP Cluster Mode (C-Mode) browse extension. To enable the extension, run the following commands from the filer console:

```
set diag
services ndmp extensions modify -is-extension-0x2050-enabled true
```



# Backing up data

- [About Backup Dump levels](#)
- [Backing up data in 7.x-7G and 7-mode](#)
- [Backing up data in C-mode](#)
- [Backing up individual data paths in C-mode](#)

## 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 Level 1 through Level 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.

**i** | **NOTE:** Clustered Data ONTAP 8.3 supports 32 (0 to 31) levels of Incremental Backups (Level 0 is a Full Backup and Level 1 through Level 31 are Incremental Backups). For more information, see the *Clustered Data ONTAP 8.3 Data Protection Tape Backup and Recovery Guide*.

## 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 in 7.x-7G and 7-mode

In Data ONTAP 7.x-7G and 7-mode, the Plug-in *for NDMP* can be used to back up files, directories, qtrees, and volumes as well as snapshots stored on volumes.

- 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 In the list of added servers, open the applicable filer to display the following nodes:
  - **Volumes**
  - **Snapshots**
- d Select the data that you want to back up.
  - **Back up all volumes:** Select the **Volumes** node.
  - **Back up individual volumes:** Open the **Volumes** node, and select the volumes that you want to back up.

**i** **NOTE:** Although LUNs are displayed as valid backup points, if you select a LUN for a backup, the backup succeeds; however, when you try to restore the LUN, Data ONTAP is unable to identify it because the metadata required to identify the LUN is not backed up when you back up a LUN directly.

The workaround is to back up a snapshot that includes the LUN. Snapshots contain the metadata required to identify a LUN. When the snapshot is restored, Data ONTAP is able to identify the restored LUN.

- **Back up individual files and directories:** Open the **Volumes** node, and then open the target volume. Select the files and directories that you want to back up.

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

- **Back up snapshots:** To list all volumes, open the **Snapshots** node, and then open the applicable volume to list directories that have any snapshots available for a backup. Open the applicable directory, and select the snapshot that you want to back up. You can only select a single snapshot.

**i** **IMPORTANT:** If you do not want to use Direct Access Restore (DAR), use the NetApp NODAR build.

- 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 2. NDMP Backup Options**

<b>Option</b>	<b>Description</b>
<b>Dump Level</b>	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.
<b>Update /etc/dumpdates</b>	The file “/etc/dumpdates” keeps a record of the backup, including the backup name, backup level, and snapshot time.  Leave this check box selected if your backup strategy includes Incremental Backups. The information in this file is used to determine what data must be included in an Incremental Backup.
<b>Backup ACLs</b>	Leave this check box selected to back up ACLs associated with the selected files.
<b>Do Direct Backup if Possible</b>	To use Direct Access Restore (DAR) during restores, leave this check box selected. DAR enables precise tape positioning, which allows faster recovery of an individual file or a small set of files.  Without DAR, the backup saveset is read sequentially to restore files. DAR requires Data ONTAP 6.x or later.
<b>Save File Information</b>	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.  Quest recommends that you leave this check box selected.
<b>Request Unicode File Information</b>	To convert the names of directories and files, and display them correctly according to the locale configured on the NetVault Backup WebUI, select this check box. This feature can only be used if the filer supports UTF-8 encoding:  To use this feature, do the following: <ul style="list-style-type: none"> <li>• Enable UTF-8 conversion on the filer.</li> <li>• Set the appropriate language type on each target volume.</li> <li>• Set the <b>Use UTF-8 Conversions</b> option when adding the NDMP Server or using the <b>Edit Server</b> method. For more information, see the <i>Quest NetVault Backup Plug-in for NDMP User’s Guide</i>.</li> <li>• Set the configuration option <b>Codepage for UTF-8 Conversions</b>. For more information, see the <i>Quest NetVault Backup Plug-in for NDMP User’s Guide</i>.</li> </ul>

Table 2. NDMP Backup Options

Option	Description
<b>Exclude</b>	<p>Use this option to specify the files and directories that you want to exclude from the backup. Type the full file path, starting with a forward slash (“/”). If the full path is not provided, the plug-in excludes any item that matches the entry. To separate multiple items, use a comma.</p> <p>You can use up to two wildcard characters per file. Wildcard characters can only be used at the first or the last place in the filename. For example, if you enter “/dir1/tmp*, *.o” in <b>Exclude</b>, the plug-in excludes the following items:</p> <ul style="list-style-type: none"> <li>Files and directories starting with the letters “<b>tmp</b>” in the directory /dir1.</li> <li>All files and directories with the extension “.o”.</li> </ul> <p>Filenames are case-sensitive for NFS shares and case-insensitive for CIFS shares.</p>
<b>Incremental forever</b>	<p>The number of dump levels for clustered environments ranges from 0 to 31. If your environment contains a large number of files that are regularly added and rarely deleted, select this option if you want to avoid generating a Full Backup after dump level 31.</p> <p>To use this feature, do the following after selecting the check box:</p> <ul style="list-style-type: none"> <li>In the <b>Number of Incremental backups</b> list, select 0 to generate only Incremental Backups after the initial Full Backup. If you want to generate a Full Backup after reaching a dump level that exceeds 31, select that number.</li> <li>From the <b>Next Backup</b> list, select the type of backup to perform.</li> </ul> <p>The Dump Date for Differential field indicates the date on which a change was identified.</p>

- c Under **Volume Backup Options**, configure the following settings.

Option	Description
<b>Backup Subvolumes Individually</b>	<p>If you are backing up more than one qtree from the same volume, select this check box. It allows you to run separate jobs for each qtree. For example, if two qtrees are selected in addition to other non-qtree files and directories, the plug-in runs three subjobs, one for each of the qtrees and one for the non-qtree data</p> <p>Note the following:</p> <ul style="list-style-type: none"> <li>The Plug-in <i>for NDMP</i> automatically creates individual jobs if the selected qtrees reside on different volumes, even if you do not select this option.</li> </ul> <p>The plug-in also creates individual jobs for each volume included in the backup.</p> <ul style="list-style-type: none"> <li>Each child job creates a separate saveset. In the preceding example, three savesets are created. This step facilitates quick restore.</li> <li>When used with other NetVault Backup features, this option can cause unwanted results. For example, when used with the option <b>Ensure this Backup is the First on the Media</b>, each subjob is set to exist first on the media and requires a separate media item.</li> </ul>

Option	Description
<b>Backup Subvolume Data Only</b>	<p>You can use this check box to back up only the qtree data while excluding the non-qtrees files and directories.</p> <p>It eliminates the need to select individual qtrees. You can select the volume node, and then select this check box to back up qtree data and exclude non-qtrees data. You must use this option with the <b>Backup Subvolumes Individually</b> option.</p> <p><b>NOTE:</b> The backup fails if there are no qtrees in the selected volume.</p>
<b>Backup Non-Subvolume Data ONLY</b>	<p>You can use this option to back up only the non-qtrees data while excluding the qtree data.</p> <p>It eliminates the need to select individual files and directories. You can select the volume node, and then select this check box to back up all non-qtrees data and exclude the qtree data. This option cannot be used if the <b>Backup Subvolumes Individually</b> check box is selected.</p>

- d Under **Sub Jobs**, configure the following settings.

Table 3. Sub Jobs

Option	Description
<b>Number of backups to run in parallel</b>	<p>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.</p> <p>To achieve the required result, set the same value for each job.</p>
<b>Delete auto generated backup jobs on completion</b>	<p>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 <b>Job Status</b> 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.</p>

- e The Backup Restart Extension (BRE) options allows you to resume a failed backup job from the last checkpoint before failure. Without BRE, any failure encountered during a backup renders the entire saveset unusable even if a significant portion of the data was successfully transferred to the backup media.

To take advantage of BRE, configure the following settings under **Restart**.

Option	Description
<b>Enable restart, if possible</b>	To enable checkpoint restarts, select this check box.
<b>Number of Gigabytes between checkpoints</b>	This option specifies the amount of data written to the media between two checkpoints. You can set it to any value from 1 through 30000 GB. The default value is zero (0). With the default setting, the checkpoints are set to media boundaries.
<b>Number of attempts to restart backup</b>	Type or select the number of times the plug-in attempts to restart the backup. You can set it to any value from 1 through 100. The default value is zero (0). With the default setting, the plug-in keeps trying to restart the backup infinitely.
<b>Delay, in minutes, between retries</b>	Specify the interval between two restart attempts. You can set it to any value from 1 through 1440 minutes. The default value is zero (0). With the default setting, the plug-in tries to resume the backup immediately upon failure.

- f Click **Save**, and in the **Create New Set** dialog box, type a name for the set.

- g 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*.

## Additional notes: BRE

- A backup job can be resumed only if suitable checkpoints exist for the job. The checkpoints are set during phase 4 of an NDMP backup and are logged as Background messages in the NetVault Backup logs.
- BRE may fail if you restart the filer. Certain circumstances may prevent the filer from correctly continuing a backup restart. When this error occurs, the message “Node Data Missing for File <Filename>” is written to the NetVault Backup logs during the index building stage. If this message is displayed in the log, it implies that the restart has failed and the job cannot be resumed.
- If you restart the NetVault Backup Service, all information related to backup checkpoints is lost. BRE does not work over NetVault Backup Service restarts.
- If network connectivity is lost during backup and the filer is unable to detect this failure, the session becomes unresponsive. When this error occurs, you must manually kill the session. To kill a session, type the following command at the command prompt on the filer:

```
ndmpd kill <sessionID>
```

- When you kill a session, the drives or changer in use go offline. To bring them back online, use the **Check** method.

Another method to bring offline drives back online is to restart the NetVault Backup Service. However, BRE does not work if you restart the service.

## Backing up data in C-mode

Starting with 10.0.5, NetVault Backup uses the CAB extension to perform backup and restore operations for cluster nodes running Data OnTAP 8.2.2 and later. The CAB extension enables the NDMP server to provide affinity information about volumes and tape devices to the backup application that support the CAB extension. NetVault Backup uses this affinity information to determine the data transfer path used during backup and restore operations. If a target volume and tape device are on same cluster node, a local backup or restore operation is performed; if they are on different cluster nodes, a three-way backup or restore operation is performed.

**i** | **NOTE:** During an NDMP backup or restore using the CAB extension, if a fail-over occurs, the job continues without any interruption. After the fail-over, if the volume is moved to a different node, the data transfer path may change depending on the affinity information provided by NetApp.

### To back up one or more volumes in C-mode:

- 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 volumes:** Select the **Volumes** node.
    - **Back up individual Vserver:** Open the **Volumes** node, and select the Vserver that you want to back up.
    - **Back up individual volumes:** Open the **Volumes** node, open the Vserver, and select the volumes that you want to back up.
    - **Back up individual files and directories:** Open the **Volumes** node, open the Vserver, and then open the target volume. Select the files and directories that you want to back up.
- i** | **NOTE:** If you omit child-level items after selecting the parent node, the job fails. To perform selective backups, open the parent node, and select each item that you want to back up.
- **Back up snapshots:** Open the **Snapshots** node, open the Vserver, and then open the applicable volume to list the snapshots available for a backup. Select the snapshot that you want to back up. You can only select a single snapshot.
- i** | **IMPORTANT:** If you do not want to use Direct Access Restore (DAR), use the NetApp NODAR build.

**i** | **NOTE:** NetVault Backup supports volume-level and file-level browsing.

In Data ONTAP 8.x C-Mode, you can only define jobs at the volume level; you cannot define jobs at the directory or file level. However, you can also define a single path that includes a directory or file, but you must define this path by manually entering the absolute path to the applicable content.

In Data ONTAP 9.x C-Mode, you can define Selection Sets at the SVM, volume, snapshot, directory, or file level.

For more information about defining and using single paths, see [Backing up individual data paths in C-mode](#).

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

**i** | **NOTE:** This step is not applicable if a file or directory is the root selection item of a backup. The backup options are only available if you have selected a volume or subvolume as the root item for a backup.

- a Next to the **Plugin Options** list, click **Create New**.
- b Configure the following settings:
  - **NDMP Backup Options:** For more information about these options, see [NDMP Backup Options](#)
  - **Sub Jobs:** For more information about these options, see [Sub Jobs](#).
- c Click **Save**, and in the **Create New Set** dialog box, type a name for the set.
- d To close the dialog box, click **Save**.

5 Complete steps 5 and 6 in [Backing up data in 7.x-7G and 7-mode](#).

# Backing up individual data paths in C-mode

In Data ONTAP 8.x C-Mode, you cannot open the **Volumes** node to browse or select the contents. However, you can specify a single data path for a backup.

- 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 Click the applicable filer, and select **Enter Backup Path** from the context menu.

- 6 In the **NDMP Backup Request** dialog box, configure the **Directory** option.

Specify the full path to the item that you want to back up. You can include a single volume, qtree, directory, or snapshot. This option does not support individual files. If you specify a filename, backup fails.

Use the following format to specify the backup path:

```
/<vserverName>/<volumeName>/ .snapshot/<snapshotName>
```

**Example:**

```
/vs1/vol2/.snapshot/daily.2014-06-16_0010
```

For more information about snapshot names, see the relevant NetApp Data ONTAP documentation.

- 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 steps 5 and 6 in [Backing up data in 7.x-7G and 7-mode](#).



# Restoring data




- Restoring data using the plug-in
- Restoring Incremental Backups
- Renaming or relocating data
- Restoring data to an alternate filer
- 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
<b>Client</b>	Displays savesets created for particular clients. The default selection is <b>Any</b> . <ol style="list-style-type: none"> <li>1 To use this filter, click the <b>Client</b> box.</li> <li>2 In the <b>Choose Client</b> dialog box, select the clients.</li> <li>3 To close the dialog box, click <b>OK</b>.</li> </ol>
<b>Plugin Type</b>	Displays savesets created using a particular plug-in. The default selection is <b>Any</b> . <ol style="list-style-type: none"> <li>1 To use this filter, click the <b>Plugin Type</b> box.</li> <li>2 In the list, select the plug-in.</li> </ol>

Filter	Description
<b>Date</b>	<p>Displays savesets created during a specified period. The default selection is <b>Any</b>.</p> <ol style="list-style-type: none"> <li>1 To use this filter, click the <b>Date</b> box.</li> <li>2 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.</li> </ol>
<b>Job</b>	<p>Displays savesets created for particular jobs. The default selection is <b>Any</b>.</p> <ol style="list-style-type: none"> <li>1 To use this filter, click the <b>Job</b> box.</li> <li>2 In the <b>Choose Job</b> dialog box, select the jobs.</li> <li>3 To close the dialog box, click <b>OK</b>.</li> </ol>

- 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.

If the backup contains a single file, the filename is added to the directory path. In the selection tree, it is displayed as **/path/to/filename/filename** and it is restored in that manner. To restore an individual file, the parent directory must have been selected for the backup job; otherwise, the restore job fails.

**i** **NOTE:** While restoring data from a snapshot backup, you must use the **Relocate** option. If you do not use this option, the restore job fails. The **Relocate** option allows you to relocate data to the same or a different volume. For more information about relocating data, see [Renaming or relocating data](#).

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 4. Restore options

Option	Description
<b>NDMP Server</b>	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.
<b>Direct Access Restore Mode</b>	<p>If you are using Data ONTAP 9.5 or later, DAR is selected by default. Otherwise, select the appropriate DAR mode from the following:</p> <ul style="list-style-type: none"> <li>• <b>Use Direct Access Restore if Possible:</b> 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.</li> <li>• <b>Use Direct Access Restore by Recovering Individual Files:</b> 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.</li> <li>• <b>Only Do Direct Access Restore:</b> 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.</li> <li>• <b>Never Do Direct Access Restore:</b> To perform a standard nondirect-access restore, select this option.</li> </ul>
<b>Incremental Restore</b>	<p>To recursively extract data and reconstruct the selected subtree to the state it was in at the time of backup, select this check box.</p> <p>This option is equivalent to setting the environment variable <code>EXTRACT=N</code>. The plug-in overwrites the existing files and deletes the files that were not present at the time of backup.</p>
<b>Restore ACLs</b>	To restore ACLs from the backed-up file, select this option. This option is equivalent to setting the environment variable <code>EXTRACT_ACL=Y</code> . You can use this option only if the <b>Backup ACLs</b> option was selected during backup.

**i | IMPORTANT:**

- When individual files are selected for DAR, their parent-level directory permissions are not restored.
- When running DAR, 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*.

## Additional notes

When restoring data, you may notice one or more of the following entries in the NetVault Backup logs. These messages are routine and do not require any intervention:

- **Mover has paused due to seek outside the supplied window:** This entry is displayed as a Background message. BRE backups create a segment on tape for each checkpoint (as configured in **Number of Gigabytes Between Checkpoints** in the Backup Options Set). This message is displayed after each individual segment is restored and it just indicates the underlying mechanism behind the NDMP operation.
- **Received media request for UID:** This entry is displayed as an Information message. It is displayed after each segment is restored as a request is sent to the Media Manager for the media that contains the next segment. This message is displayed even if the next segment resides on the same media.
- **No suitable drive or media for the job — Request held in queue:** This entry is displayed as a Warning. It is displayed when the required piece of media is not available or it is in use. If the piece of media being requested is the current one, this message indicates a temporary condition and the restore resumes shortly.

## 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.
- 2 For more information, see [Restoring data using the plug-in](#).
- 3 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 backup 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 Click the volume that you want to rename, and select **Rename** from the context menu.
- 3 In the **Restore Rename** dialog box, configure the following option:
  - 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 on the filer.

In Data ONTAP 8.1 C-mode, use the following format to specify the new path:

```
/<vserverName>/<volumeName>/<folderName>
```

**Example:**

```
/vs1/vol4/data/restore
```

- 4 To close the dialog box, click **OK**.

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.

- 5 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. It 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 configure the following option:  
**NDMP Server:** This option specifies the target filer name. By default, it is set to the name of the original filer from which the data was backed up. To relocate the data to an alternate server, 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](#).

## 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

- [Upgrading from 10.x or 11.1 to 11.4](#)
- [Common errors](#)

## Upgrading from 10.x or 11.1 to 11.4

If you are using a Network Appliance (NetApp) filer and Data ONTAP and you upgrade from version 10.x or 11.1 of the plug-in to version 11.4, the selection set created using the earlier version is not displayed correctly in the selection tree on the **NetVault Backup Selections** page. You must manually clear and select the correct Volumes to include in backup jobs. Use the **Manage Sets** option to update and save the selections.

## Common errors

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

Table 5. 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 <b>"postmaster.pid"</b> file from the location referenced in the log and restart the NetVault Backup Server.

**Table 5. Troubleshooting**

Description	Symptom	Solution
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 <b>Show processes from all users</b> . You can see multiple instances of <b>postgres32.exe</b> running on the system. To remove all instances of <b>postgres32.exe</b> , select any instance of this process, and click <b>End Process</b> . Start the NetVault Backup Service.
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 <b>localhost</b> , and fails to start. Check the <code>/etc/hosts</code> file, and if the file does not contain an entry for <b>localhost</b> , add the entry.



Quest provides software solutions for the rapidly-changing world of enterprise IT. We help simplify the challenges caused by data explosion, cloud expansion, hybrid datacenters, security threats, and regulatory requirements. We are a global provider to 130,000 companies across 100 countries, including 95% of the Fortune 500 and 90% of the Global 1000. Since 1987, we have built a portfolio of solutions that now includes database management, data protection, identity and access management, Microsoft platform management, and unified endpoint management. With Quest, organizations spend less time on IT administration and more time on business innovation. For more information, visit [www.quest.com](http://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.