



Quest[®] NetVault[®] 13.0
Administrator's Guide



© 2020 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 Software, Quest, the Quest logo, QoreStor, 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.

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

NetVault Administrator's Guide
Updated - May 2020
Software Version - 13.0
NVG-101-13.0-EN-01

Contents

Introduction	15
About Quest NetVault	15
Key benefits	15
Feature summary	16
About this document	16
Target audience	17
Recommended additional reading	17
Getting started	18
About deploying NetVault	18
About NetVault components	19
NetVault Server	19
NetVault Clients	20
NetVault plug-ins	20
NetVault WebUI	21
NetVault Command Line Interface	21
Starting or stopping the NetVault Service	22
Enabling Web Service over HTTP or HTTPS	22
Logging in to NetVault	24
Overview of NetVault WebUI	24
Navigation pane	27
Quitting NetVault	31
Using the configuration wizard	31
Monitoring the NetVault Server	33
Sorting records in NetVault WebUI	34
Customizing tables in NetVault WebUI	35
Filter options	35
Deploying REST APIs	41
Configuring clients	43
About NetVault Clients	43
Installing software from WebUI (push installation)	43
About push installation	43
Prerequisites	44
Managing package stores	45
Performing push installations	47
Monitoring deployment tasks	52
Managing deployment tasks	53
Adding clients	54
Adding a client to the NetVault Server	55
Adding multiple clients	55
About firewall settings	58
Locating a client	60

Checking communication through a firewall	61
Adding a server as a client	61
Removing a client from the list of available clients	62
Managing clients	62
Viewing client details	62
Setting client description	63
Installing plug-ins	64
Checking for upgrades	65
Installing a product license file	66
Checking client access	67
Removing plug-ins	67
Removing a client from the server	67
Managing client groups	67
About client groups	68
Creating a client group	68
Viewing existing client groups	68
Modifying a client group	69
Removing a client group	69
Managing catalog search	70
About catalog search	70
Prerequisites	70
Configuring catalog search	71
Relocating the catalog search database directory	75
Changing the NetVault Catalog Database directory on a Windows-based machine	75
Migrating NetVault Catalog Server to different Windows-based client machine	75
Changing the NetVault Catalog Database directory on a Linux-based machine	76
Modifying catalog search	77
Disabling catalog search	78
Upgrading catalog search	78
Removing catalog search	78
Displaying Status of Catalog Search	78
Configuring storage devices	80
About storage devices	80
SAN considerations	80
Quest DR Series systems	81
About Quest DR Series systems	81
Quest DR Series system prerequisites	81
Adding a Quest DR Series system	82
Quest QoreStor	84
Secure Connect	84
Adding QoreStor	85
Adding a container as a Media for QoreStor	85
NetVault SmartDisk	86
About NetVault SmartDisk	86
Adding a NetVault SmartDisk	87

Migrating savesets and backup jobs from NetVault SmartDisk to QoreStor	88
EMC Data Domain Systems	90
About EMC Data Domain Systems	90
Data Domain System prerequisites	92
Adding a Data Domain System	92
DD Boost commands	94
Snapshot Array Manager	97
Supported Storage Array Manager	97
Prerequisites	97
Adding Snapshot Array Manager	97
Virtual Tape Libraries	98
About Virtual Tape Library	98
Virtual Tape Library considerations	99
Creating and adding a Virtual Tape Library	99
Re-adding a previously created VTL	101
Virtual standalone drives	101
About virtual standalone drive	101
Creating and adding a virtual standalone drive	101
Shared Virtual Tape Libraries	102
About Shared Virtual Tape Libraries	102
SVTL considerations	103
SVTL prerequisites	103
Creating and adding an SVTL	106
Re-adding a previously created SVTL	107
Physical tape devices	108
Prerequisite	108
Adding a standalone tape drive	108
Adding a tape library	109
Storage tiers	110
Creating a storage tier	111
Editing a storage tier	111
Removing a storage tier	112
Backing up data	113
About backing up data	113
Backup job definition	113
Backup indexes	114
Secondary Copy	114
Snapshot-based backups	115
Backup retirement	115
About NetVault Sets	117
Set types	117
Backup and recovery strategy	119
Creating backup jobs	119
Creating Schedule Sets	121
Creating Target Sets	124
Selecting the device	125

Specifying media options	126
Configuring media sharing options	128
Creating Source Sets for backup jobs	129
Creating Backup Advanced Options Sets	129
Setting backup retirement options	130
Specifying additional options	133
Creating a Secondary Copy	135
Configuring pre- and post script options for backup jobs	140
Configuring user-defined events for backup jobs	142
Managing sets	142
Modifying a set	142
Copying a set	143
Deleting a set	143
Managing policies	144
About policies	144
Creating a policy	145
Viewing existing backup policies	147
Viewing job status of policy jobs	148
Modifying a backup policy	148
Modifying jobs in a backup policy	148
Modifying clients in a backup policy	149
Modifying events in a backup policy	149
Quiescing a backup policy	149
Deleting a backup policy	150
Restoring data	151
About restoring data	151
Restore job definition	152
Creating restore jobs	152
Restoring data using an existing Restore Selection Set	154
Searching for files in savesets	155
Viewing media list	157
Creating Source Sets for restore jobs	157
Creating Restore Advanced Options Sets	158
Setting restore type	158
Specifying additional options	159
Configuring pre- and post -scripts for restore jobs	159
Configuring user-defined events for restore jobs	161
Managing online backup indexes	161
About managing online indexes	162
Manually deleting online indexes	162
Loading offline indexes	163
Manually compressing online indexes	163
Uncompressing online indexes	164
Managing Snapshots	164
Restoring snapshots	165

Setting expiry for Snapshots	165
Mounting Snapshots	165
Unmounting Snapshots	166
Managing NetVault dashboard	167
About NetVault dashboard	167
Viewing NetVault dashboard	167
Configuring Dashboard	169
Managing Dashboard	169
Adding Dashboard	169
Managing a widget on NetVault dashboard	171
Adding Widgets	172
Positioning Widgets	172
Widget filter options	172
Managing jobs	175
Viewing job activity	175
Viewing job calendar	177
Managing jobs	178
Running a job immediately	179
Aborting a job	179
Stopping a job	180
Restarting a job	180
Placing a job on hold	181
Resuming a job	181
Determining the reason for “waiting for media” status	181
Viewing log messages for a job	182
Viewing and managing a job	182
Monitoring job progress	184
Clearing job errors and warnings	185
Removing a job schedule	185
Managing job definitions	185
Viewing job definitions	185
Editing job definitions	187
Deleting job definitions	188
Adding a non-policy job to a policy	188
Viewing job history	188
Monitoring logs	190
About NetVault logs	190
Viewing log messages	190
Downloading logs	192
Exporting logs	193
Manually purging the log messages	194
Setting up a user-defined log event	194
Removing a user-defined log event	195
Searching the knowledge base	195

Managing storage devices	196
Monitoring device activity	196
Managing disk-based storage devices in list view	198
Viewing disk device details	198
Managing a Quest DR Series system, QoreStor, or Data Domain system	200
Checking the status of a disk-based storage device	208
Changing the status of a disk-based storage device	208
Scanning a disk-based storage device	209
Removing all savesets from a disk-based storage device	209
Removing a disk-based storage device	211
Managing disk-based storage devices in tree view	211
Viewing disk device details	211
Checking the status of a disk-based storage device	214
Changing the status of a disk-based storage device	214
Setting storage properties for a disk-based storage device	214
Scanning a disk-based storage device	215
Removing a disk-based storage device	215
Managing the Snapshot Array Manager	216
Modifying the Snapshot Array Manager settings	216
Changing the user credentials for Snapshot Array Manager	216
Changing the status of Snapshot Array Manager	217
Removing the Snapshot Array Manager	217
Managing tape libraries in list view	217
Viewing tape library details	218
Opening and closing library door	219
Opening and closing entry/exit ports	219
Unloading or importing tapes from entry/exit ports	219
Exporting tapes to entry/exit ports	220
Restarting ACSLS or NDMP Libraries	220
Importing shadow tapes (NetApp VTL)	220
Removing a tape library	221
Managing tape libraries in tree view	221
Viewing tape library details	222
Modifying a library	222
Changing the device view type	223
Opening and closing library door	223
Opening and closing entry/exit ports	224
Unloading or importing tapes from entry/exit ports	224
Exporting tapes to entry/exit ports	224
Restarting ACSLS or NDMP libraries	224
Importing shadow tapes (NetApp VTL)	225
Removing a tape library	225
Managing tape drives in list view	226
Viewing tape drive details	226
Configuring performance options for a tape drive	228
Checking the status of a tape drive	231
Changing the status of a tape drive	232

Configuring cleaning slots	232
Configuring the cleaning lives option	232
Configuring automatic cleaning options for a tape drive	233
Manually submitting a drive cleaning request	233
Unloading a tape	234
Loading a tape	234
Removing a tape drive	234
Managing tape drives in tree view	235
Viewing tape drive details	235
Configuring physical tape drives	237
Checking the status of a tape drive	238
Changing the status of a tape drive	238
Configuring cleaning slots	238
Configuring the cleaning lives option	238
Configuring automatic cleaning options for a tape drive	239
Manually submitting a drive cleaning request	239
Unloading a tape	239
Loading a tape	240
Marking media for re-use	240
Removing a tape drive	240
Adding shared devices	241
Adding shared drives to non- shared library using the semi- automatic method	241
Adding shared drives to non-shared library manually	241
Sharing standalone drives	242
Managing storage media	244
Viewing storage summary	244
Viewing disk storage details	245
Viewing tape storage details	247
Managing tape storage media	248
Labeling tape media	248
Blanking tape media	252
Scanning tape media	254
Marking a tape as unusable	256
Marking a tape as read-only	256
Marking a tape for reuse	257
Removing an offline tape	257
Managing savesets	257
Viewing saveset details	258
Configuring saveset expiry options	258
Expiring all savesets on disk-based storage	260
Deleting savesets from disk-based storage devices	260
Deleting savesets from tape-based storage devices	261
Managing user and group accounts	263
About user accounts	263
About user and group privileges and presets	264
Creating a user account	264

Setting user password	265
Configuring user details	266
Configuring memberships for a user account or user group	267
Granting privileges and quota to a user account	267
Granting add and remove permissions	268
Setting up user notification profile	269
Modifying a user account	269
Deleting a user account	269
Setting a user policy	270
Using Secure Mode	270
Enabling or Disabling Secure Mode	270
Using Presets	271
Creating a Preset	271
Modifying a Preset	271
Deleting a Preset	271
User privileges	272
Predefined Presets	274
Integrating Active Directory with NetVault	277
Considerations for Linux-based NetVault	278
Managing AD users	278
Managing AD groups	280
Using Azure Active Directory as an identity provider	281
Setting up an Azure AD login with NetVault	282
Managing an Azure AD user	284
Managing Azure AD groups	285
Monitoring events and configuring notifications	287
About NetVault events	287
About notification methods	287
Event classes	288
Events types	288
Viewing event logs	292
Using global notification methods	293
About global notification methods	293
Setting up a global notification method	293
Using custom notification methods	295
About custom notification methods	295
Creating a custom notification method	296
Setting up a custom notification method	296
Removing a user-defined job event	297
Removing a user-defined report event	297
Reporting in NetVault	298
Reporting system overview	298
Working with reports	298
Generating reports	299
Setting a favorite report	300

Setting filters for report	300
Editing report job definitions	301
Using notification methods to send reports by email or print reports	302
Customizing table views for reports	302
Adding charts to reports	304
Exporting reports to PDF files	305
Available reports	306
Working with client clusters	323
About client cluster support	323
Virtual clients	324
Device configuration in cluster environment	324
Installing and upgrading cluster-aware plug-ins	325
Prerequisites	325
Installing a cluster-aware plug-in	325
Upgrading a cluster-aware plug-in	326
Configuring a cluster-aware plug-in	326
Configuring preferred network address for cluster nodes	327
Configuring default settings for a cluster-aware plug-in	327
Managing virtual clients	328
Modifying a virtual client	328
Checking access to a virtual client	328
Determining the current real client	329
Removing a virtual client	329
Backups using cluster-aware plug-ins	329
Cluster failover during backups	330
Restores using cluster-aware plug-ins	330
Configuring default settings for NetVault	331
About configuring default settings	331
Other configuration utilities	333
Configuring encryption settings	333
Configuring plug-in options	334
Configuring default settings for Disk Devices Plug-in	334
Configuring default settings for nvjobstart	335
Configuring default settings for post-scripts	336
Configuring default settings for Verify Plug-in	336
Configuring Deployment Manager Settings	337
About Deployment Manager	337
Configuring default settings for Deployment Manager	338
Configuring Job Manager settings	338
About Job Manager	338
Configuring default settings for Job Manager	338
Configuring Logging Daemon settings	339
About Logging Daemon	339
Configuring disk space alert thresholds	340
Modifying the purge policy for log messages	341

Configuring additional settings for Logging Daemon	341
Configuring Media Manager settings	342
About Media Manager	343
Configuring general settings for Media Manager	343
Configuring retirement rules for dependent backups	345
Configuring Media Manager settings for tape devices	346
Configuring Media Manager settings for backup indexes	347
Configuring Media Manager settings for RAS devices	349
Configuring transfer update settings for foreign RAS devices	350
Configuring data transfer stall timeout period for NetVault SmartDisk	350
Configuring media request weightings	350
Configuring default interval for backup retirement scans	351
Configuring an alternate index read block size for Quest DR Series systems	352
Configuring Network Manager settings	352
About Network Manager	353
Configuring timeout settings for Network Manager	353
Enabling or disabling availability broadcasts	354
Configuring network addresses for multihomed machines	355
Reducing startup delay	356
Configuring default port for Network Manager	357
Configuring default port for Communications Manager	358
Configuring Process Manager settings	358
About Process Manager	358
Configuring shared memory settings	358
Configuring RAS device settings	359
Configuring connection settings for NetVault SmartDisk	359
Configuring Schedule Manager settings	360
About Schedule Manager	360
Configuring default settings for Schedule Manager	360
Configuring default job priority settings	362
Configuring Web Service settings	362
Configuring Web Service to disable ciphers	362
Configuring Web Service to disable protocols	363
Enabling Web Service auditing	363
Configuring timeout period for client operations	363
Configuring timeout period for saveset removal request	364
Configuring Auditor Daemon settings	364
About Auditor Daemon	365
Configuring Auditor to log only failed requests	365
Modifying the purge policy for audit logs	365
Configuring firewall settings	366
Configuring general settings	367
Relocating default directories	367
Relocating the NetVault Database directory	369
Configuring TCP/IP buffer sizes	370
Changing language and locale settings	371
Disabling pre-installation package compatibility check	372

Configuring license expiration warning period	372
Configuring security settings	373
Disabling password protection for a client	373
Changing NetVault password	373
Synchronizing NetVault Time	374
About NetVault Time	374
Configuring an alternate NetVault Time Server	374
Configuring default settings for global notification methods	375
Configuring email server settings for notifications	375
Configuring SysOp email ID for notifications	376
Configuring a default printer for notifications	376
Configuring network manager host settings for notifications	376
Configuring the reporting utility	377
About reporting utility	377
Customizing CSV report templates	378
Customizing HTML report templates	378
Customizing miscellaneous report templates	379
Customizing plain text report templates	380
Configuring timeout setting for report generation	380
Configuring default mail format type for report	381
Configuring default settings for Statistics Manager	381
Creating a global purge policy for the Reports Database	382
Creating table-specific purge policy	383
Configuring NetVault WebUI default settings	384
Configuring NetVault to use a specific VSS provider	385
Configuring default settings using Txtconfig	386
Diagnostic tracing	388
About diagnostic tracing	388
Managing trace filters	388
Enabling tracing	389
Downloading trace files	391
Changing the trace directory location	392
Enabling tracing using Txtconfig	392
Disabling tracing	393
Deleting trace session directories	393
Managing diagnostic data	395
About support diagnostics	395
Downloading diagnostic data	395
Uploading diagnostic data	396
Using the deviceconfig utility	398
About deviceconfig	398
Configuring default settings for tape libraries	398
General settings for tape libraries	399
Drive cleaning settings	400

Mixed media settings	401
Configuring default settings for tape drives	402
NDMP settings	403
General settings for tape drives	403
Software compression settings	405
Drive performance settings	406
Statistics collection settings	407
Generic cleaning settings	407
NetVault processes	408
About NetVault processes	408
Process description	408
Environment variables	414
Environment variables in NetVault	414
Network ports used by NetVault	416
Ports used or required	416
Troubleshooting	418
Common errors	418
Safe Mode in NetVault	430
About us	432
Technical support resources	432

Introduction

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

About Quest NetVault

Quest® NetVault® (NetVault) offers the most advanced, cross-platform data protection capabilities as well as unsurpassed ease of use, out-of-the-box deployment, and pain-free scalability

NetVault allows you to safeguard your data and applications in both physical and virtual environments from one intuitive user interface and to protect a massive number of servers that contain many petabytes of data. NetVault also features heterogeneous support, so you can safeguard data on a wide range of operating systems, applications, databases, processor architectures, and networked storage devices. Such cross-platform versatility makes it easy for you to tailor NetVault to match the ever-changing and growing landscape of your IT infrastructure.

Key benefits

- Simple, out-of-the-box deployment for fast time to value
- Fully automated installation manager that simplifies deployment of Windows clients
- Protection for both physical and virtual environments for cost savings
- Heterogeneous server support for unparalleled coverage in diverse IT environments
- Broad application support
- Disk-based backup and deduplication to significantly improve storage efficiency
- Seamless integration with the Quest DR Series systems for source-side deduplication and WAN-optimized replication
- Comprehensive Network Attached Storage (NAS) protection to safeguard critical data
- Feature rich protection of virtual environments
- Bare metal recovery to drastically reduce the time it takes to recover a failed server
- Extensive storage attachment options to allow for distributed backup targets and workload
- Dynamic device sharing to optimize backup data transfers and reduce points of failure
- Fast file-system backups with high performance multi-streaming

Feature summary

- **Application protection:** Ensure the availability of business-critical applications such as Oracle, SQL Server, Exchange, SharePoint, MySQL, PostgreSQL, Domino, DB2, Informix, SAP, and Sybase with application plug-ins. These plug-ins complement native solutions to save you time on integration. No scripting is required to run backup and recovery jobs.
- **NAS protection:** Get advanced data protection for information stored on NAS appliances, including the devices made by Dell, EMC, Hitachi, IBM, NetApp, and Sun. You can reduce traffic over the LAN and maximize performance by backing up data using Network Data Management Protocol (NDMP). NetVault supports many different storage topologies and configurations, allowing you to perform backups directly to a locally attached SCSI device, a SAN-attached device, or a storage device elsewhere on the network.
- **Enterprise-wide control:** Use the web-based interface to configure, manage, and monitor your backup and recovery operations. Convenient wizards guide you through the common tasks, such as creating backup jobs, assigning policies, configuring storage devices, and generating reports.
- **Back up to disk and tape:** Leverage disk- and tape-based backups to a wide range of storage targets, including NAS devices and third-party deduplication appliances. NetVault also allows you to move data from one storage target to another for off-site storage and disaster recovery purposes.
- **Data deduplication:** Reduce your data storage footprint with powerful deduplication technologies. NetVault integrates seamlessly with the Quest DR Series system of deduplication appliances and the Rapid Data Access (RDA) client-side deduplication technology, enabling you to take full advantage of the appliance's powerful deduplication, compression, and replication capabilities. Additionally, NetVault supports software-defined Quest QoreStor™ storage devices (Linux only), which also provide the advantage of RDA client-side deduplication. NetVault also supports EMC Data Domain backup and recovery platform, as well as its DD Boost technology for deduplication.
- **Virtualization support:** Extend advanced data protection to VMware and Hyper-V environments. NetVault gives you consistent, reliable, point-and-click backup and restore for virtual environments without requiring you to be an expert.
- **High-performance multi-streaming:** Save time and ease management by backing up multiple workloads simultaneously.
- **Strong security:** Meet regulatory requirements without sacrificing backup windows or deduplication performance with encryption plug-ins for CAST-128, CAST-256, and AES-256 algorithm support. Flexible job-level encryption lets you easily select which data to encrypt.
- **Simple, straight-forward licensing:** License NetVault based on capacity or by component. This option gives you incredible flexibility to choose the model that best meets the organization's needs. With component-based licensing, you can choose which modules you need. You can also license NetVault by capacity and deploy an unlimited number of clients or application plug-ins. Quest offers two capacity-based licensing editions for NetVault. One offers great value for growing businesses; the other offers comprehensive protection for large enterprises.

About this document

This guide describes how to configure and use NetVault to protect your data. It provides comprehensive information about all NetVault features and functionality.



IMPORTANT:

- NetVault stores the system data in a PostgreSQL database, which resides on the NetVault Server. You should not attempt to modify the NetVault Database directly using any PostgreSQL tools unless directed by Quest Technical Support personnel. Improper changes to the database can cause irrecoverable data corruption. Before modifying the NetVault Database, make sure that you create a backup copy of the database. For more information about backing up the NetVault Database, see the *Quest NetVault Built-in Plug-ins User's Guide*.
- NetVault stores the system configuration settings in “.cfg” files, which reside in the **config** directory under the NetVault installation directory. The settings in these files should only be modified under the guidance of Quest Technical Support personnel. Improper changes to these files can cause errors and other unexpected behavior. Before modifying a configuration file, make sure that you create a backup copy of the file.

Target audience

This guide is intended for backup administrators and other technical personnel who are responsible for designing and implementing a backup strategy for the organization. A good understanding of the operating systems on which the NetVault Server and Clients are running is assumed.

Recommended additional reading

- *Quest NetVault Installation Guide*: This guide provides information about installing the NetVault Server and Client software.
- *Quest NetVault CLI Reference Guide*: This guide provides information about using the NetVault command-line utilities.
- *Quest NetVault Plug-in for FileSystem User's Guide*: This guide provides information about installing, configuring, and using NetVault Plug-in for *FileSystem*.
- *Quest NetVault Built-in Plug-ins User's Guide*: This guide provides information about configuring and using the following plug-ins:
 - NetVault Plug-in for *Consolidation*
 - NetVault Plug-in for *Data Copy*
 - NetVault Plug-in for *Databases*
 - NetVault Plug-in for *Raw Devices*
- *Quest NetVault SmartDisk Installation/Upgrade Guide*: This guide provides information about installing the NetVault SmartDisk software.
- *Quest NetVault SmartDisk Administrator's Guide*: This guide provides information about administering a NetVault SmartDisk instance.
- *Quest NetVault Compatibility Guide*: This guide provides information about the platforms, operating system versions, and application versions that are supported by NetVault.

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

Getting started

- [About deploying NetVault](#)
- [About NetVault components](#)
- [Starting or stopping the NetVault Service](#)
- [Enabling Web Service over HTTP or HTTPS](#)
- [Logging in to NetVault](#)
- [Overview of NetVault WebUI](#)
- [Navigation pane](#)
- [Quitting NetVault](#)
- [Using the configuration wizard](#)
- [Monitoring the NetVault Server](#)
- [Sorting records in NetVault WebUI](#)
- [Customizing tables in NetVault WebUI](#)
- [Deploying REST APIs](#)

About deploying NetVault

In a NetVault setup, one machine is configured as the NetVault Server and various other machines that you want to protect function as NetVault Clients. A single server and its associated clients form a **NetVault Domain**.

The following diagram depicts a NetVault deployment.

NetVault Clients

The NetVault Client is installed on machines that you want to protect using the NetVault solution. These machines can be file servers, database servers, email servers, application servers, and workstations.

The NetVault Clients are assigned to a NetVault Server, which manages all data protection operations for the clients. A single server and its associated clients form a NetVault Domain. A NetVault SmartClient license is required to attach physical or virtual storage devices locally to a NetVault Client.

The NetVault Client can run on AIX, FreeBSD, HP-UX, Linux, Mac OS X, Solaris, and Windows operating systems.

NetVault plug-ins

The NetVault Plug-ins are used to protect various applications and data stored on the server and client machines. There are two categories of NetVault plug-ins: built-in plug-ins and licensed plug-ins.

Built-in plug-ins

The built-in plug-ins are packaged with the NetVault software, and automatically installed on the respective machines when you install the NetVault Server and Client software.

NetVault offers the following types of built-in plug-ins:

- **NetVault Plug-in for File System (Plug-in for File System):** The Plug-in for File System protects critical file system data, and minimizes downtime by allowing you to restore full volumes, individual partitions, or individual directories and files quickly and reliably with minimal interaction.
- **NetVault Plug-in for Consolidation (Plug-in for Consolidation):** The Plug-in for Consolidation lets you create a composite saveset by combining a Full Backup and its associated Incremental Backups. You can use this consolidated saveset as the base for subsequent Incremental Backups. The Plug-in for Consolidation does not back up data from a client; it just creates a composite set from existing savesets.
- **NetVault Plug-in for Data Copy (Plug-in for Data Copy):** The Plug-in for Data Copy lets you create one or more copies of backups for off-site storage and disaster recovery purposes. The Plug-in for Data Copy does not back up data from a client; it just creates a copy of an existing backup.
- **NetVault Plug-in for Databases (Plug-in for Databases):** The Plug-in for Databases protects system data (such as configuration files, system settings, backup indexes, backup media information, job schedule, licenses, and other data) stored in the NetVault Database. You can use this backup to recover a functional NetVault Server after a failure.
- **NetVault Plug-in for Encryption (Plug-in for Encryption):** These plug-ins provide support for CAST-128, AES-256, and CAST-256 algorithms to meet regulatory backup security requirements.
- **NetVault Plug-in for Raw Devices (Plug-in for Raw Devices):** The Plug-in for Raw Devices protects data stored on physical disks. The plug-in lets you recover the Master Boot Record (MBR), system partitions, and individual user partitions from a point-and-click user interface without writing any complex scripts.
- **NetVault Plug-in for Rapid Data Access (Plug-in for RDA):** The Plug-in for RDA lets you use the client-side and inline deduplication capabilities that are available in other products, such as Quest DR Series system of deduplication appliances and software-defined Quest QoreStor storage devices (Linux only).

For more information about built-in plug-ins, see the *Quest Plug-in for File System User's Guide* and *Quest NetVault Built-in Plug-ins User's Guide*.

Licensed plug-ins

The licensed plug-ins are available as separate products, and installed on the NetVault Server and Client machines to protect specific applications and appliances.

NetVault offers the following types of licensed plug-ins:

- **Plug-ins for applications:** These plug-ins provide data protection to business-critical applications, such as Oracle, SQL Server, Exchange, SharePoint, MySQL, PostgreSQL, Domino, DB2, Informix, SAP, and Sybase.
- **Plug-ins for NDMP-based NAS appliances:** These plug-ins enable NDMP-based backups and restores for appliances that support this protocol. NetVault also offers specialized plug-ins that integrate the NetApp SnapMirror, SnapVault, and Snapshot technologies to provide advanced data protection to NetApp appliances.
- **Plug-ins for virtual environments:** These plug-ins provide data protection to virtual machines in VMware and Hyper-V environments.
- **Plug-ins for device integration:** These plug-ins enable configuration of specialized tape libraries for use in a NetVault environment.
- **Plug-ins for bare metal recovery:** These plug-ins let you recover an entire system, including the operating system, applications, system settings, partition information, and data on supported Windows and Linux operating systems.

NetVault also offers cluster-aware versions of various plug-ins that enable data protection for distributed data.

For more information about the licensed plug-ins, see the respective plug-in user's guide.

NetVault WebUI

NetVault offers a web-based user interface, called the NetVault WebUI, to configure, manage, and monitor your NetVault system. You can use the NetVault WebUI to perform various tasks such as the following:

- Configure performance, security, and other options
- Manage clients, storage devices, and storage media
- Perform backups and restores
- Monitor jobs, device activities, and logs
- Set up notifications
- Generate and view reports

You can access the NetVault WebUI from any standard browser. You can use the WebUI to remotely administer a NetVault Server from any system on which you can run a supported web browser and connect to the server over HTTP or HTTPS.

NetVault Command Line Interface

NetVault also provides a Command Line Interface that allows you to configure and manage a NetVault system from a command prompt window or terminal window.

You can use the NetVault CLI utilities to perform various tasks such as the following:

- Start and stop the NetVault Service
- Configure performance, security, and other options
- Manage clients, storage devices, and storage media
- Perform backups and restores
- Generate and view reports

The NetVault command-line utilities can be accessed from the NetVault Server and Client machines. For more information about the command-line interface, see the *Quest NetVault Command Line Interface Reference Guide*.

Starting or stopping the NetVault Service

The NetVault Service is configured to start automatically on the server and client machines. You can use the Txtconfig utility or CLI to manually start or stop the NetVault Service.

i | **NOTE:** You must be logged-in with Administrator privileges on Windows-based systems and root user privileges on Linux- and UNIX-based systems to use these utilities.

To manually start or stop the NetVault Service:

1 Using Txtconfig:

- a In a terminal or command prompt window, type, `txtconfig`, and press **Enter** or **Return**.
- b On the **Machine** page, press **p** to view the main menu, and then press the option number for the **Services** page.
- c Depending on the current state of the NetVault Service, press the option number to stop or start the service. Press **s**, and then press **q**.

2 Using CLI:

- a On Linux-based systems:
To start the service, type: `$NV_HOME/etc/startup.sh start`
To stop the service, type: `$NV_HOME/etc/startup.sh stop`
- b On Windows-based systems:
To start the service, type: `net start "NetVault Process Manager"`
To stop the service, type: `net stop "NetVault Process Manager"`

Enabling Web Service over HTTP or HTTPS

By default, the Web Service is enabled over HTTPS. You can modify the Web Service settings to change the port or protocol. These settings also let you specify the security certificate file and private key file for HTTPS communications.

You can modify the Web Service settings through the following interfaces: NetVault WebUI (through the **Change Settings** page), or Txtconfig.

i | **NOTE:** NetVault uses port 8443 as the default HTTPS listen port for incoming Web Service connections. When installing the server, if the default port (8443) is in use by any other application, NetVault assigns the first available port in the port range 50486 to 65535.

You must be logged-in with Administrator privileges on Windows-based systems and root user privileges on Linux- and UNIX-based systems to use this procedure.

To configure the Web Service settings for NetVault:

- 1 Access the Web Service settings by using any of the following methods:
 - **NetVault WebUI:**
 - a Start the NetVault WebUI, and in the Navigation pane, click **Change Settings**.
 - b On the **NetVault Server Settings** page, under **Services**, click **Web Service**.

- **Txtconfig:**
 - a In a terminal or command prompt window, type, `txtconfig`, and press **Enter** or **Return**.
 - b On the **Machine** page, press **p** to view the main menu, and then press the option number for the **Web Service** page.

2 Configure the following settings.

Table 1. Web Service settings

Setting	Description
HTTP settings	<p>To access the Web Service through HTTP, configure the following settings:</p> <ul style="list-style-type: none"> • Enable Web Service over HTTP: Select this check box. If you are using Txtconfig, press the option number for this setting to change the value to ON. • HTTP Listen port for incoming Web Service connections: The default HTTP port is 80. If this port is in use by any other server or application, configure an alternate port.
HTTPS settings	<p>To access the Web Service through HTTPS, configure the following settings:</p> <ul style="list-style-type: none"> • Enable Web Service over HTTPS: This protocol is selected by default. HTTPS is the preferred protocol. This protocol provides encrypted communication between the client and server. It protects sensitive data such as NetVault passwords passed between the browser and NetVault Web Service. • HTTPS Listen port for incoming Web Service connections: By default, NetVault uses port 8443 for HTTPS. If this port is in use by any other server or application, configure an alternate port. • WebService security certificate file: To use HTTPS, provide a SSL certificate. NetVault provides a self-signed certificate (server.crt), which resides in the etc directory under the NetVault installation directory. This certificate generates warnings in most browsers. For the browser to accept a certificate without warnings, provide a valid certificate file signed by a trusted certificate authority. • Ciphers to disable for incoming Web Service connections: You can modify the Web Service configuration to disable ciphers for incoming web service connections settings to prevent from allowing one or more ciphers. These settings can be configured from the NetVault Server Settings page. By default, this field is blank and all the ciphers are allowed. • Protocols to disable for incoming Web Service connections: You can modify the Web Service configuration to disable protocol for incoming web service connections settings to prevent from allowing one or more protocols. These settings can be configured from the NetVault Server Settings page. By default, this field is blank and all the protocols are allowed. • WebService private key file: Provide the private key file required for HTTPS communications. The default key file is server.key, which resides in the etc directory under the NetVault installation directory.

3 To save the settings, do the following:

- **NetVault WebUI:** Click **Apply** to apply the settings and close the WebUI dialog box.
- **Txtconfig:** Press **s** to save the settings, and then press **q** to quit Txtconfig.

Logging in to NetVault

The following options are available for logging in to the NetVault Core:

- NetVault local user
- Active Directory Domain (AD) user
- OpenLDAP Directory Services user
- Samba Active Directory user]
- Microsoft Azure Active Directory (Azure AD) user

i **NOTE:** You must be logged-in with Administrator privileges on Windows-based systems and root user privileges on Linux- and UNIX-based systems to run NetVault.

For more information about logging in using AD, see [Integrating Active Directory with NetVault](#).

Before logging in with the Azure AD option, register NetVault in the Azure portal and then configure Azure AD in NetVault. For more information, see [Using Azure Active Directory as an identity provider](#).

To log in to NetVault:

- 1 Open a browser window. In the address bar, type:

```
https://<machine-name>:8443
```

Press **Enter**.

- 2 In the Login dialog box, complete one of the following options:

- To log in using NetVault local or Domain user credentials, type your user name and password, and then click **Sign In**.

i **NOTE:** There are two predefined user accounts in NetVault:

- **admin:** The administrator account for NetVault.
- **default:** A standard user account that can be used to perform various operations in NetVault.

After installing NetVault, you can use either the **admin** or the **default** user account to log in to NetVault. By default, no password is assigned to these user accounts. For more information about NetVault user accounts, see [Managing user and group accounts](#).

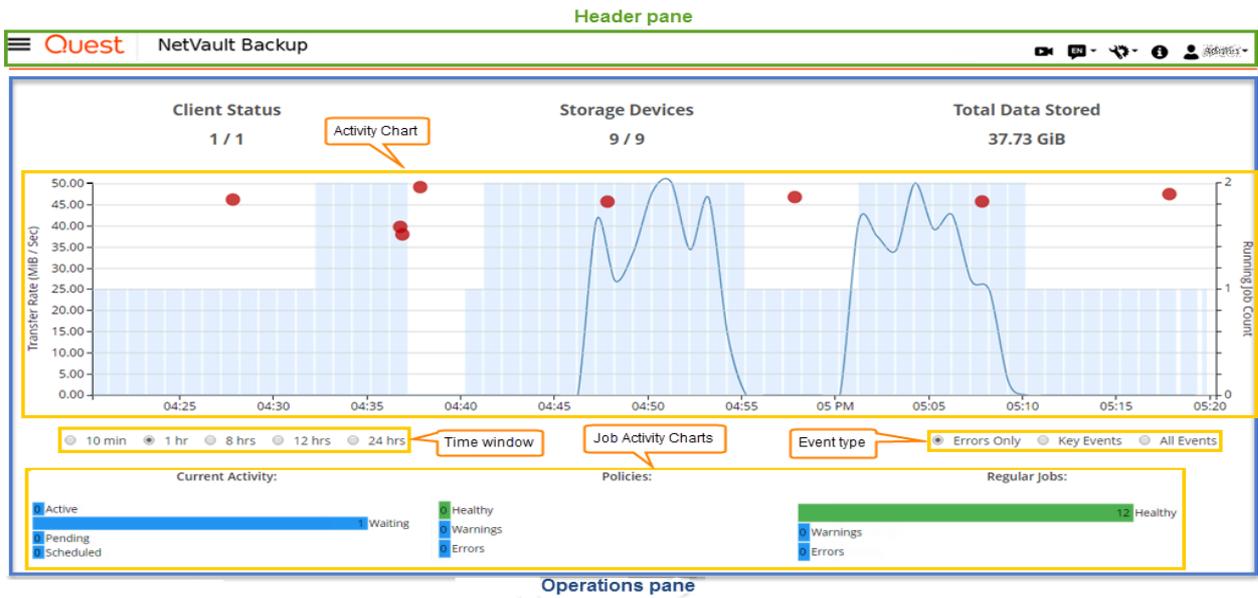
- To log in using Azure AD credentials, click **Sign in with Azure AD**, enter the appropriate credentials, and then click **Enter**.

After you log in, the WebUI opens the **Server Monitor** page in your browser window.

Overview of NetVault WebUI

The NetVault WebUI consists of the Header pane, Navigation pane, and Operations pane.

Figure 2. NetVault WebUI home page



The following table provides a brief description of the WebUI panes:

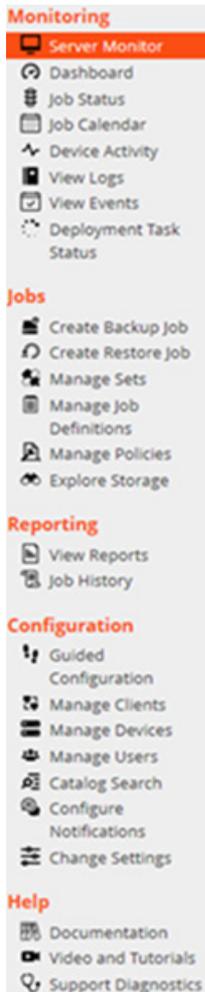
Table 2. NetVault WebUI panes

Pane	Description
Header pane	<p>This pane includes the following items:</p> <ul style="list-style-type: none"> • Video icon: Provides access to the video tutorial for the currently loaded page. The link opens in a new browser window or tab. • Language: Displays the list of available languages from which you can select the display and input language for NetVault WebUI. This option does not change the locale setting for NetVault. To change the locale for NetVault, see Changing language and locale settings. • Settings icon: Displays the list of options to change Application Settings, Clear NetVault cache, and Clear NetVault UI Preferences. <p>Application Settings</p> <p>To change the Application Settings (Navigation, Color Palette, Show Actions as, Time Formatting and Default Table Pagesize) for NetVault WebUI, follow these steps:</p> <ol style="list-style-type: none"> a In the header pane of NetVault WebUI, click the Settings icon and select Application Settings. Application Settings drawer is displayed on the right-side of the page. Configure the following options: <ul style="list-style-type: none"> ▫ Navigation: To remove the auto-hide setting of navigation bar in the NetVault WebUI, clear the Auto-hide Navigation Drawer check box. By default, the navigation bar is set to auto-hide. ▫ Color Palette: To change the color theme of NetVault WebUI to dark theme, select the Dark theme. By default, the NetVault WebUI is displayed in Light theme. ▫ Show Actions as: To view the Actions option as a page-level and table-level context menus, select Context menus. By default, the action options are displayed as buttons on bottom of the NetVault WebUI pages. When you select the Context Menus option, the page-based action buttons on the bottom of the page are replaced with a context menu shown as Horizontal Ellipsis icon at the top of the page (after the page title). Also, the table-based actions are shown in the 'Actions' column in the table. Click the Horizontal Ellipsis icon and select the required action. This option is applicable to some of the NetVault WebUI pages. ▫ Time Formatting: To change the time format, select 12-hour or 24-hour format from the application settings. By default, the application sets 24-hour time format from the server settings. Change the time format in following ways: <ul style="list-style-type: none"> - 12-hour (AM/PM): Sets 12-hour (AM/PM) format of the NetVault WebUI. - 24-hour: Sets 24-hour format for the NetVault WebUI. <p>NetVault applies time format configuration changes to all sections of NetVault WebUI except for report based features.</p>

To lock the Navigation pane:

- 1 In the header pane of NetVault WebUI, click the **Settings** icon and select **Application Settings**. **Application Settings** drawer is displayed on the right- side of the page.
- 2 In the Navigation section, clear the option '**Auto-hide Navigation Drawer**'. By default, the Navigation pane hides automatically when you move the pointer away from the icon for the navigation pane.
- 3 Click **OK**.

Figure 3. NetVault WebUI navigation pane



The following table provides a brief description of the links available in the Navigation pane.

Table 3. Navigation pane

Section	Item	Description
Monitoring	Server Monitor	Opens the Server Monitor page. Use this page to view the overall status of your NetVault Server. The Activity Chart shows the data transfer rate for jobs and the number of active jobs. You can also view the events that occurred during the selected time window. For more information, see Monitoring the NetVault Server .
	Dashboard	Opens the Dashboard page. Use this page to view the overall statistics of the NetVault. For more information, see Managing NetVault dashboard .

Table 3. Navigation pane

Section	Item	Description
	Job Calendar	<p>Opens the Job Calendar page.</p> <p>You can view your jobs for a month, a week, or a day on a calendar. It can be helpful to view your jobs in the calendar format to ensure that there are no scheduling conflicts. For more information, see Viewing job calendar.</p>
	Job Status	<p>Opens the Job Status page.</p> <p>Use this page to monitor the progress and status of your jobs. You can also use the page to manage your jobs and perform various job-related tasks, such as canceling a job, stopping or restarting a job, viewing job logs, or removing a job schedule. For more information, see Managing jobs.</p>
	Device Activity	<p>Opens the Device Activity page.</p> <p>Use this page to monitor data flows and data transfer rates for devices that are in use. For more information, see Monitoring device activity.</p>
	View Logs	<p>Opens the View Logs page.</p> <p>Use this page to view the current log messages. You can also use this page to perform various log-related tasks, such as downloading, exporting, or deleting log messages. For more information, see Monitoring logs.</p>
	View Events	<p>Opens the View Events page.</p> <p>Use this page to view events logs for NetVault. For more information, see Viewing event logs.</p>
	Deployment Task Status	<p>Opens the Deployment Task Status page.</p> <p>Use this page to monitor the progress and status of push installation tasks. The page shows the summary of current and completed tasks in the form of bar charts and provides the progress information for individual machines. For more information, see Monitoring deployment tasks.</p>

Table 3. Navigation pane

Section	Item	Description
Jobs	Create Backup Job	Starts the backup job wizard. For more information about creating and scheduling backup jobs, see Creating backup jobs .
	Create Restore Job	Starts the restore job wizard. For more information about creating and submitting restore jobs, see Creating restore jobs . You can also use this page to perform index management tasks. For more information, see Managing online backup indexes .
	Manage Sets	Opens the Set Management page. Use this page to view, modify, or delete existing sets. For more information, see Managing sets .
	Manage Job Definitions	Opens the Manage Job Definitions page. Use this page to view, modify, or remove job definitions. For more information, see Managing job definitions .
	Manage Policies	Opens the Policy Management page. Use this page to create and manage policy-based backups. For more information, see Managing policies .
	Explore Storage	Opens the Explore Storage page. Use this page to explore and manage disk- and tape-based storage media. You can also use this page to perform various media-related tasks, such as labeling media, scanning media, blanking media, or write-protecting tapes. For more information, see Managing storage media .
Reporting	View Reports	Opens the View Reports page. Use this page to access the predefined reports. For more information about these reports, see Available reports . You can do the following: view reports, customize report views, send reports by email, and export reports to PDF files. For more information, see Reporting in NetVault .
	Job History	Opens the Job History page. Use this page to view completed jobs. You can also use this page to view or modify job definitions. For more information, see Viewing job history .
Configuration	Guided Configuration	Starts the NetVault Configuration Wizard. The wizard guides you through the various aspects of setting up your backup system. You can use the wizard to install client and plug-in packages, add clients, configure devices, and create backup jobs. For more information, see Using the configuration wizard .
	Manage Clients	Opens the Manage Clients page. Use this page to add and manage clients, client groups, and virtual clients. For more information, see Configuring clients and Working with client clusters .
	Manage Devices	Opens the Manage Devices page. Use this page to add and manage disk- and tape-based storage devices. For more information, see Managing storage devices .

Table 3. Navigation pane

Section	Item	Description
	Users and Groups	Opens the Manage User and Groups Accounts page. Use this page to create and manage user and group accounts, create user notification profile, and set user password policy. If your environment uses Active Directory (AD), you can also use this page to integrate AD users with NetVault, as well as manage user groups. For more information, see Managing user and group accounts .
	Catalog Search	Opens the Catalog Search Configuration page. Use this page to configure and manage catalog search service on NetVault Server running on Windows and Linux operating systems. For more information, see Managing catalog search .
	Configure Notifications	Opens the Editing Global Notification Profile page. Use this page to configure global notification methods for one or more NetVault events. For more information, see Using global notification methods .
	Change Settings	Opens the NetVault Server Settings page. Use this page to customize the NetVault system and change default settings for the NetVault Server. Use the Choose Machine button at the bottom of the page to access settings for NetVault Client machines. For more information, see Configuring default settings for NetVault .
Help	Documentation	Provides access to the product documentation. The link opens in a new browser window or tab.
	Video and Tutorials	Provides access to the video tutorial page. The link opens in a new browser window or tab.
	Support Diagnostics	Opens the Support Diagnostics page. Use this page to download the detailed information of your NetVault environment to your local machine or upload the diagnostic data corresponding to a SR number and provide it directly to Quest Technical Support for further analysis. For more information, see Managing diagnostic data .

Quitting NetVault

To quit NetVault, move the pointer over the user area at upper-right corner of the NetVault WebUI, and click **Sign Out**.

Using the configuration wizard

The NetVault WebUI provides a configuration wizard that guides you through the various aspects of setting up your backup system. You can use the wizard to install client and plug-in packages, add clients, configure devices, and create backup jobs. The configuration wizard can be accessed from the **Guided Configuration** link in the Navigation pane.

To use the configuration wizard:

- 1 Start the NetVault WebUI, and log in to NetVault.
- 2 In the Navigation pane, click **Guided Configuration**.

- 3 Select the type of task that you want to perform.

Figure 4. NetVault Configuration Wizard

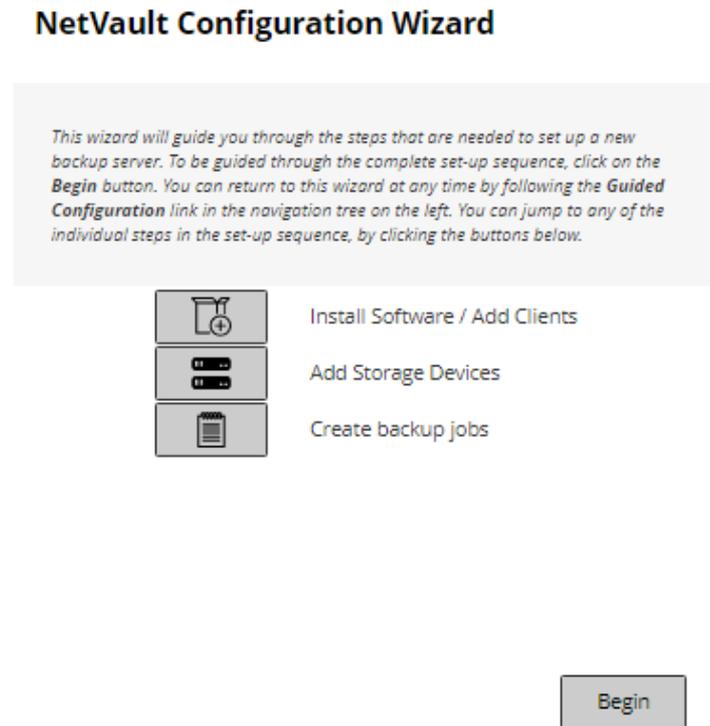


Table 4. Guided Configuration options

Option	Description
Install Software and Deploy Clients	Installs and upgrades client and plug-in packages on remote machines, and adds new machines as clients to the NetVault Server. For more information, see Performing push installations .
Add Clients	Adds a NetVault Client to the Server. For more information, see Adding a client to the NetVault Server .
Install Plugins	Installs a plug-in on one or more clients. For more information, see Installing plug-ins using the push installation method .

Table 4. Guided Configuration options

Option	Description
Add Storage Devices	<p>Adds a storage device to the NetVault Server. The available device types include the following:</p> <ul style="list-style-type: none"> • Single virtual disk device: Adds a virtual standalone drive. For more information, see Virtual standalone drives. • Virtual tape library/media changer: Adds a Virtual Tape Library (VTL). For more information, see Virtual Tape Libraries. • Shared virtual tape library: Adds a Shared Virtual Tape Library (SVTL). For more information, see Shared Virtual Tape Libraries. • Single physical tape device: Adds a standalone tape drive. For more information, see Physical tape devices. • Tape library/media changer: Adds a tape library. For more information, see Physical tape devices. • NetVault SmartDisk: Adds a Quest NetVault SmartDisk. For more information, see NetVault SmartDisk. • Quest RDA Device: Adds a Quest DR Series system or QoreStor. For more information, see Quest DR Series systems or Quest QoreStor. • Data Domain Boost Device: Adds an EMC Data Domain System. For more information, see EMC Data Domain Systems. • Snapshot Array Manager: Adds Snapshot Array Manager. For more information, see Snapshot Array Manager.
Create Backup Jobs	Select this option to create and schedule backup jobs. For more information, see Creating backup jobs .

- 4 Follow the instructions to complete the configuration steps.
- 5 After a task is completed successfully, a message is displayed.
- 6 To continue, click a button in the Operations pane. Alternatively, click a link in the Navigation pane to exit the configuration wizard and open a different page.

Monitoring the NetVault Server

You can monitor the overall status of your NetVault Server from the **Server Monitor** page. You can use this page to view the status of clients, devices, regular jobs, and policy jobs. The Activity Chart shows the data transfer rate for jobs and the number of active jobs. You can also view the events that occurred during the selected time window.

To monitor the NetVault Server:

- 1 In the Navigation pane, click **Server Monitor**.

i | **NOTE:** The **Server Monitor** page is automatically loaded when you log on to the WebUI.

- 2 On the **Server Monitor** page (see [Figure 2, NetVault WebUI home page](#)), you can view the following information.

Table 5. Server Monitor page

Item	Description
Client Status	<p>This area shows the number of online clients and total client count. Move the pointer over the count to view the total number of clients and number of offline clients.</p> <p>You can click this area to open the Manage Clients page.</p>
Storage Devices	<p>This area shows the number of online devices and total device count. Move the pointer over the count to view total count, the number of offline devices and number and type of online devices.</p> <p>You can click this area to open the Manage Devices page.</p>
Total Data Stored	<p>This area shows the total amount of data backed up from various clients.</p> <p>You can click this area to open the Explore Storage page.</p>
Activity Chart	<p>This chart displays the data transfer rate for active jobs and the number of active jobs. You can also view the events that occurred during the selected time window.</p> <p>You can use the following settings to modify the Activity Chart:</p> <ul style="list-style-type: none"> • Time window: This setting allows you to change the time window for the Activity Chart. The available options are 10 minutes, 1 hour, 8 hours, 12 hours, and 24 hours. By default, the time window is set to 1 hour. • Event type: This setting allows you to change the event type displayed on the page. The available options are Errors Only, Key Events, and All Events. By default, the event type is set to Errors Only.
Job Activity Charts	<p>This area shows the summary of current jobs, policy jobs, and regular jobs in the form of bar charts.</p> <ul style="list-style-type: none"> • Current Activity: The individual bars represent the number of jobs that are in active, waiting, pending, and scheduled states. • Policies: The individual bars represent the number of policy jobs that have completed successfully, completed with warnings, and failed. • Regular Jobs: The individual bars represent the number of regular jobs that have completed successfully, completed with warnings, and failed. <p>You can click a bar to open the Job Status page and view the job details for that category. For example, you can click the Active bar in the Current Activity area to view the jobs that are in progress. Similarly, you can click the Errors bar in the Regular Jobs area to view the regular jobs that have failed.</p>

3 To open a page, click the corresponding link in the Navigation pane.

Sorting records in NetVault WebUI

NetVault web application supports multiple-column sorting for the tables. However, the first sorted column takes the highest precedence over the sorting of other columns. For example, if the table is sorted by the 'Start Time' as the primary sort, selecting sort options on the 'ID/ Instance/ Phase' column does not affect the list order.

To sort by a column, click the column header; to change the sort direction, click the column header again. Repeat this step for other columns by which you want to sort the table.

The arrowhead next to the column name indicates the sorting order (up for ascending order and down for descending order). A blue arrowhead is used to indicate the sort order for the primary column.

To remove a sort, click the column header for the corresponding column until the arrowhead is no longer displayed.

To choose the following actions in a column, click the column menu:

Option	Description
Sort Ascending	Sorts the column in ascending order.
Sort Descending	Sorts the column in descending order.
Columns	Select/ clear to show/ hide the required column (s).

Customizing tables in NetVault WebUI

The following table provides a brief description of the icons at the lower- right corner of the table. You can use these icons to customize the records in the table.

Table 6. Customizing tables

Icon	Description
	<p>This icon is shown above the table as well as at the lower- right corner of the table. Use this icon to filter the records in the table. Click this icon to display the following two options:</p> <ul style="list-style-type: none"> Edit Filters: Click this option to set the filters for the records displayed in the table. <ul style="list-style-type: none"> 'Filters' drawer is displayed on the right- side of the page. Set the filter options and click Apply. For more information on filter options for the respective table, refer to Filter options. To clear the filter settings, click Clear. To close the 'Filters' drawer without editing or setting the filter options, click Cancel. Reset Filters: Click this option to remove the user-defined filter option settings for the records in the table. This option is not shown if filters are not applied in the table. <p>By default, the filter is not applied and all the records are displayed in the table. When you set the filters, the applied filter options are displayed above the table and the color of filter icon and table border is changed. For example, the filters are shown as View By: Current Activity, Run Status: Waiting, Scheduled, and so on.</p>
	<p>Move the pointer over this icon to view the page size setting, column sorting order, and filters applied in the table.</p>
	<p>To export the table data to a CSV format file, click this icon.</p>
	<p>Click this icon for Table Settings:</p> <ul style="list-style-type: none"> Set Pagesize: Click this to set the number of records per page to display in the table. Table Options dialog box is displayed. Specify the required number in the field "Specific number per page". Click OK. By default, table options are set to Automatically, based on what will fit. Restore Defaults: Click this to remove the user-defined settings for the table. A confirmation dialog appears. Click OK.

Filter options

The following table filter options are displayed in the 'Filters' drawer on the right side of the page. You can use one or more filters to display records that match the specified criteria. You can also include wildcards ("?" or "**") in the search filter string.

Table 7. Job status filter options

Option	Description
View By	<p>Use this option to filter the jobs based on their category. Select one of the following options:</p> <ul style="list-style-type: none"> • Current Activity • Regular Jobs • Policy Jobs • Policy Jobs by policy name <p>By default, Current Activity option is selected.</p>
Run Status	<p>Use this option to filter jobs by their Run Status. For example, for 'Current Activity' option in the View By field, shows the following run statuses:</p> <ul style="list-style-type: none"> • Active • Waiting • Pending • Scheduled <p>By default, all run statuses are selected. To clear a run status, click the corresponding button.</p>
Start Time	<p>From: To display the jobs from a specific date and time, do the following:</p> <ul style="list-style-type: none"> • Type the start date, or click the button next to the box, and select the start date. • Type the start time, or click the button next to the box, and select the start time. <p>To: To display the jobs up to a specific date and time, do the following:</p> <ul style="list-style-type: none"> • Type the end date, or click the button next to the box, and select the end date. • Type the end time, or click the button next to the box, and select the end time. <p>By default, the jobs are filtered by 'Start Time'. The default value for this setting is seven days.</p> <p>You can change the default number of days by modifying the configuration settings for the Schedule Manager. For more information, see Configuring default settings for Schedule Manager.</p> <p>NOTE: The Start Time filter only applies to the current session. This setting just hides the display of job status records on the Job Status page. It does not delete the records from the Scheduler Database.</p>
Job Title	Use this option to filter jobs where the job name matches the specified string.
Job Type	Use this option to filter jobs by job type. You can select backup, restore, or reporting job types.
Client	Use this option to filter jobs for a particular client.
Plugin	Use this option to filter jobs performed using a particular plug-in.
Job ID	Use this option to filter jobs by the job ID number.
Instance	<p>Use this option to filter jobs by the instance ID number.</p> <p>From: To filter jobs from a specific instance ID, type the instance ID number.</p> <p>To: To filter jobs up to a specific instance ID, type the instance ID number.</p>
Phase	<p>Use this option to filter jobs by the phase ID number (1 or 2).</p> <p>From: To filter jobs from a specific phase ID, type the phase ID number.</p> <p>To: To filter jobs up to a specific phase ID, type the phase ID number.</p>
Next Runtime	Use this option to filter jobs by the text in the 'Next runtime' column.
Current Status Text	Use this option to filter jobs by the text in the 'Current Status' column.

Table 7. Job status filter options

Option	Description
Last Exit Status Text	Use this option to filter jobs by the text in the 'Last Exit Status' column.
Exit Status	<p>This option is displayed when you select 'Regular Jobs', 'Policy Jobs', and 'Policy Jobs by policy name' in the View By option.</p> <p>Use this option to filter jobs by their exit status such as Healthy, Warnings, and Errors. By default, all statuses are selected. To clear a status, click the corresponding button.</p>
Select Policies	<p>This option is displayed when you select 'Policy Jobs by policy name' in the View By option.</p> <p>Use this option to filter the policy jobs by their policy name. Select the policy name to filter the jobs in the policy.</p>

Table 8. View logs filter options

Option	Description
Display Level	<p>Use this option to filter the messages based on their severity level. When you specify the severity level, the messages at that level and higher are displayed.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> • All • Background • Information • Job Messages • Warnings • Errors • Severe <p>By default, Display Level is set to 'Job Messages'.</p>
Date	<p>From: To filter logs from a specific date and time, do the following:</p> <ul style="list-style-type: none"> • Type the start date, or click the button next to the box, and select the start date. • Type the start time, or click the button next to the box, and select the start time. <p>To: To filter logs up to a specific date and time, do the following:</p> <ul style="list-style-type: none"> • Type the end date, or click the button next to the box, and select the end date. • Type the end time, or click the button next to the box, and select the end time. <p>NOTE: To display from the first log message, select only the To option. To display up to the last log message, select only the From option.</p>
Job ID	Use this option to filter logs by job ID number. Type the job ID number.
Instance	Use this option to filter logs by instance ID number. Type the instance ID number.

Table 8. View logs filter options

Option	Description
Classes	<p>Use this option to filter logs for a particular class. The following are the log classes:</p> <ul style="list-style-type: none"> • System • Schedule • Jobs • Media • Devices • Database • Plugins • UI <p>By default, all log classes are selected. To clear a class, click the corresponding button.</p>
Clients	Use this option to filter logs for a particular client. Select the client name.
Message text	Use this option to filter logs that contain a particular string. Type the filter string.

Table 9. View events filter options

Option	Description
Date	<p>From: To filter events from a specific date and time, do the following:</p> <ul style="list-style-type: none"> • Type the start date, or click the button next to the box, and select the start date. • Type the start time, or click the button next to the box, and select the start time. <p>To: To filter events up to a specific date and time, do the following:</p> <ul style="list-style-type: none"> • Type the end date, or click the button next to the box, and select the end date. • Type the end time, or click the button next to the box, and select the end time. <p>NOTE: To display from the first event, select only the To option. To display up to the last event, select only the From option.</p>
Classes	<p>Use this option to filter events for a particular class. Events are organized into the following categories or classes:</p> <ul style="list-style-type: none"> • Audit • NetVault Time • Device • Job • Licensing • Log Daemon • Machines • Media • Media Database • Policy • Scheduler Database • Stats Collection • Appliance <p>By default, all event classes are selected. To clear a class, click the corresponding button.</p>

Table 9. View events filter options

Option	Description
Event Name	Use this option to filter events by name.
Message Text	Use this option to filter event messages that contain a particular string. Type the filter string.

Table 10. Manage job definitions filter options

Option	Description
Job Title	Use this option to filter jobs where the job name matches the specified string.
Job ID	Use this option to filter jobs by the job ID number.
Policy	Use this option to filter jobs by policy name.
Job Type	Use this option to filter jobs by job type. You can select backup or restore job types.
Plugin	Use this option to filter jobs performed using a particular plug-in.
Client	Use this option to filter jobs for a particular client.
Selection Set	Use this option to filter jobs by selection set.
Plugin Options Set	Use this option to filter jobs by plug-in options set.
Schedule Set	Use this option to filter jobs by schedule set.
Source Set	Use this option to filter jobs by source set.
Target Set	Use this option to filter jobs by target set.
Advanced Options Set	Use this option to filter jobs by advanced options set.
Next Run Time Text	Use this option to filter jobs by the text in the column 'Next Run Time'.

Table 11. Manage policies filter options

Option	Description
Policy Name	Use this option to filter the policy jobs where the policy name matches the specified string.
Policy Status	<p>Use this option to filter policy jobs based on policy status. The following policy statuses are displayed:</p> <ul style="list-style-type: none"> • Completed • Warning • Failed <p>By default, all statuses are selected. To clear a policy status, click the corresponding button.</p>
Policy State	<p>Use this option to filter policy jobs based on policy state. The following policy states are displayed:</p> <ul style="list-style-type: none"> • Dormant • Active • Quiescing • Quiesced <p>By default, all states are selected. To clear a policy state, click the corresponding button.</p>

Table 12. Job history filter options

Option	Description
Job Title	Use this option to filter jobs where the job name matches the specified string.
Job ID	Use this option to filter jobs by the job ID number. From: To filter jobs from a specific job ID, type the job ID number. To: To filter jobs up to a specific job ID, type the job ID number.
Phase	Use this option to filter jobs by the phase ID number (1 or 2). From: To filter jobs from a specific phase ID, type the phase ID number. To: To filter jobs up to a specific phase ID, type the phase ID number.
Instance	Use this option to filter jobs by the instance ID number. From: To filter jobs from a specific instance ID, type the instance ID number. To: To filter jobs up to a specific instance ID, type the instance ID number.
Client	Use this option to filter jobs for a particular client.
Plugin	Use this option to filter jobs performed using a particular plug-in.
Policy	Use this option to filter jobs included in a policy.
Job Type	Use this option to filter jobs by job type. You can select backup, restore, or reporting job types.
Status	Use this option to filter jobs by their Run Status (Succeeded, Failed, Aborted, and others).
End Time	Use this option to filter jobs that completed at a particular time. From: To filter jobs from a specific date and time, do the following: <ul style="list-style-type: none"> • Type the start date, or click the button next to the box, and select the start date. • Type the start time, or click the button next to the box, and select the start time. To: To filter jobs up to a specific date and time, do the following: <ul style="list-style-type: none"> • Type the end date, or click the button next to the box, and select the end date. • Type the end time, or click the button next to the box, and select the end time.

Table 13. Create Restore Job - Choose Saveset filter options

Option	Description
Backup Time	Use this option to filter the savesets created during a specified period. From: To filter savesets from a specific date and time, do the following: <ul style="list-style-type: none"> • Type the start date, or click the button next to the box, and select the start date. • Type the start time, or click the button next to the box, and select the start time. To: To filter savesets up to a specific date and time, do the following: <ul style="list-style-type: none"> • Type the end date, or click the button next to the box, and select the end date. • Type the end time, or click the button next to the box, and select the end time.
Client	Use this option to filter the savesets created for particular clients. Select the client or clients to view the corresponding saveset or savesets.
Plugin	Use this option to filter the savesets created using a particular plug-in. Select the plug-in to view the corresponding savesets.

Table 13. Create Restore Job - Choose Saveset filter options

Option	Description
Job	Use this option to filter the savesets created for particular jobs. Select the job or jobs to view the corresponding saveset or savesets.
Selection Set	Use this option to filter the savesets that belong to particular selection sets. Select the select set or sets to view the corresponding saveset or savesets.

Table 14. Catalog Search Results filter options

Option	Description
Backed Up	Use this option to filter the catalog search records from a saveset that is backed up during a specified period. From: To display the records from backed up date and time, do the following: <ul style="list-style-type: none"> Type the start date, or click the button next to the box, and select the start date. Type the start time, or click the button next to the box, and select the start time. To: To display the records up to backed up date and time, do the following: <ul style="list-style-type: none"> Type the end date, or click the button next to the box, and select the end date. Type the end time, or click the button next to the box, and select the end time.
Clients	Use this option to filter records created for particular clients. To filter the records for a particular virtual machine, from the Plug-in for VMware, select the desired VM name (s). To hide the records of a client, clear the corresponding check box.
Plugins	Use this option to filter records for a particular plug-in. To hide the records of a plug-in, clear the corresponding check box.
Jobs	Use this option to filter records by the job ID number. To hide the records of a Job ID, clear the corresponding check box.

Deploying REST APIs

NetVault offers REST APIs managed Swagger, an open-source suite of API developer tools. To obtain the REST API files, contact Quest Support.

To deploy the REST APIs, on a Windows server or a Linux server, complete one of the following procedures.

i | **NOTE:** When deploying on Windows, run the command line as Administrator. When deploying on Linux, execute the commands as root.

To deploy REST APIs on Windows:

- 1 On the NetVault Server, download and unzip the REST API files.
- 2 On the NetVault Server, download and install Node.js from <https://nodejs.org/en/>.
- 3 If not already installed, download and install NPM from <https://npmjs.com/get-npm>.
- 4 Update config.js (for URL, DB_LOCATION, SERVER_PUBLIC_IP, CLI_PATH) present in the API folder.
- 5 In the root directory of the drive where you installed Node.js, create a folder named "db" (for example, C:\db).

- 6 On the command line as Administrator, change to the path that contains the unzipped API files and run the following command:

```
npm install
```

- 7 Start the node server and run the following command:

```
npm run dev
```

- 8 After the server starts, to start using the APIs and view API documentation, open the following Swagger link: <http://localhost:3000/v2/api-docs/index.html>.

To deploy REST APIs on Linux:

- 1 On the NetVault Server, download and unzip the REST API files.
- 2 On the NetVault Server, download and install Node.js from <https://nodejs.org/en/download/> or from the terminal.
- 3 If not already installed, download and install NPM according to your Linux distribution.
- 4 Install Nodemon by running the following command as a superuser on the terminal:

```
npm install -g nodemon
```

- 5 Update config.js (for URL, DB_LOCATION, SERVER_PUBLIC_IP, CLI_PATH) present in the API folder.
- 6 In the root directory of the drive where Node.js is installed, create a folder named “db” (for example, /usr/db).
- 7 On the terminal, change to the path containing the unzipped API files and run the following command:

```
npm install
```

- 8 Run the following command:

```
export PATH=/usr/netvault/util:$PATH
```

- 9 Start the node server and run the following command:

```
npm run dev
```

- 10 After the server starts, open one of the following Swagger links to start using the APIs and view API documentation:

- <http://localhost:3000/v2/api-docs/index.html>
- <http://yourmachineip:3000/v2/api-docs/index.html>

Configuring clients

- [About NetVault Clients](#)
- [Installing software from WebUI \(push installation\)](#)
- [Adding clients](#)
- [Managing clients](#)
- [Managing client groups](#)

About NetVault Clients

The NetVault Clients are machines that you want to protect using the NetVault solution.

These machines require at least the client version of NetVault and TCP/IP connectivity to the server. You can attach physical and virtual storage devices locally to a client after installing the NetVault SmartClient license on it.

To use a client in a backup or restore operation, you must first add the client to the NetVault Server. A single server and its associated clients form a NetVault Domain. A NetVault Server acts as a client to itself, and you can also add it as a client to other NetVault Servers.

Installing software from WebUI (push installation)

This section includes the following topics:

- [About push installation](#)
- [Prerequisites](#)
- [Managing package stores](#)
- [Performing push installations](#)
- [Monitoring deployment tasks](#)
- [Managing deployment tasks](#)

About push installation

The push installation method lets you install one or more software packages on remote machines. You can use this method to install or upgrade the NetVault Client software and NetVault plug-ins on multiple machines, including Windows-based and Linux-based machines. Push installation is available for Windows-to-Windows, Windows-to-Linux, Linux-to-Linux, and Linux-to-Windows scenarios.

You can run push installation tasks from the NetVault WebUI. When the task starts, a deployment agent is installed on the remote machines. The agent copies the installation packages from a shared location and performs silent

installations on the machines. The agent also sends status messages to the server. By default, the client logs are preserved on the target machine. For more information about how to modify this setting, see [Configuring default settings for Deployment Manager](#). You can view the task status from the WebUI. After the packages are successfully installed on the machines, the WebUI automatically adds the new clients to the NetVault Server.

By default, NetVault runs a maximum of 50 concurrent push-installation processes. You can modify the default settings to increase the number of concurrent processes. For more information, see [Configuring default settings for Deployment Manager](#).

Prerequisites

Before you start the push installation procedure, verify that the following requirements are met:

- **Copy the packages to a shared location:** Copy the client and plug-in packages to a shared location. Only CIFS shares, including Linux Samba Shares, are currently supported as package stores. The path must be accessible to the NetVault Server and all target machines where you want to install the packages.

Ensure that you use the original names of the installation packages. Renamed packages cannot be selected for push installations.

- **Configure a package store in NetVault:** After copying the installation packages, configure the shared location details in NetVault. For more information, see [Configuring a package store](#).
- **Configure the preferred network address setting on multihomed NetVault Servers:** On multihomed NetVault Servers, configure the **Preferred Network Address** setting to allow the remote clients to send status messages to the correct address. For more information about this setting, see [Configuring network addresses for multihomed machines](#).

If the remote clients fail to contact the server on the correct address, the task status is not updated on the server. In such cases, the clients are not added to the server and the task either remains in the Running state or completes with warnings.

- **Verify that the firewall is configured to allow traffic through the ports required for push installation:** To push the client and plug-in packages to a remote Windows machine, the NetVault Server establishes the initial WMI session with the remote machine using RPC over port 135. All further WMI traffic uses the dynamic port range of 49152 through 65535.

To push the client and plug-in packages to a remote Linux machine, the NetVault Server establishes an SSH connection with the remote machine over port 22.

If there is a firewall between the server and the remote client, ensure that these ports are opened to allow the WMI RPC traffic to pass through the firewall. Also ensure that an inbound rule is created for the HTTP or HTTPS listen port that is configured on the NetVault Server.

When pushing the NetVault client installation to a Linux machine, you need the following additional prerequisites:

- **Grant the proper access:** The user must have root-level access, with the NOPASSWD flag set in the sudoers file on the remote Linux machine. To complete this prerequisite, add one of the following entries to the sudoers file:

```
<username> ALL=NOPASSWD: ALL
```

or

```
<username> ALL=(ALL) NOPASSWD: <users home directory [/home/username or  
/home/domainname/username]>/nvpushinstall/nvclientinstaller
```

- **NOTE:** If user fails to get sudoers privilege from the sudoers file, the deployment task will always be in the running state with **Request sent to launch installation agent on the remote machine** message.

- **Affirm the password authentication:** In the `sshd_config` file, set the `PasswordAuthentication` entry to `yes`.

- If you are not using the default cipher setting, then support for aes128-ctr should be present. If it is not, then you must add `aes128-ctr` at the end of the list following a comma in the `etc/ssh/sshd_config` file.
- **Install libstdc++:** If you are pushing the hybrid installation package of the NetVault client on a 64-bit Linux machine, you must first install the 32-bit version of libstdc++.

When you push the NetVault client from a Linux machine to Windows machine, you need the following additional prerequisites:

- Ensure the **WinRM** version is greater than 2.0.
- Check WinRM on Windows client machine to verify Kerberos value:
 - 1 Run the following command as an administrator in command line interface to open the config file.


```
winrm get winrm/config
```
 - 2 Check the **Kerberos** value in the **Auth** section is set to **true**:


```
Kerberos = true
```
- When using push install as a local user, ensure that the following WinRM configurations exist on the target Windows machine:
 - Under Service:


```
AllowUnencrypted="true"
```
 - Under Service/Auth:


```
Basic="true"
```

For more information on Installation and configuration for Windows Remote Management, see,

<https://docs.microsoft.com/en-us/windows/desktop/winrm/installation-and-configuration-for-windows-remote-management>.

- Configure **Kerberos** on Linux machine:
 - Run the following command to open the **krb5 config** file:


```
vi /etc/krb5.conf
```

In the libdefaults section set the following value:

```
default_realm=<EXAMPLE.COM>
default_ccache_name = KEYRING:persistent:%{uid}
```
- When using domain user for push installation from a Linux machine to a Windows machine, mention the target machine in FQDN format. When using local user, mention the IP address of the target machine.
- During Push installation specify the domain suffix (for example: domain.com or domain.local) for the same account name or DL format.

Managing package stores

This section includes the following topics:

- [Configuring a package store](#)
- [Modifying a package store](#)
- [Removing a package store](#)

Configuring a package store

A package store is a shared location used for storing NetVault Client and plug-in binary files for push installations. Only CIFS shares, including Linux Samba shares, are currently supported as package stores. You can set up a package store in NetVault by providing the path and user credentials for the CIFS share.

i | **NOTE:** Quest recommends using shared folders instead of shared drives as package stores for NetVault.

To configure a package store:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Install Software/Add Clients**.
- 2 On the **Machines to Be Added as Clients** page, click **Install Software**.
- 3 Click **Manage Stores**, and provide the following details.

Option	Description
Store Name	Provide a display name for the package store. You cannot change the store name after it is configured.
Type	Select CIFS .
Location	Provide the Uniform Naming Convention (UNC) path of the CIFS (Common Internet File System) share. NOTE: Do not use hyphens in the folder names of the package store path. Underscores are acceptable. The format for specifying the path for a Windows-based server is: <code>\\<server>\<share name></code> The format for specifying the path for a Linux-based server is: <code>//<server>/<share name></code> NOTE: <code>//<server>/<share name></code> To reduce the time it takes to browse the package store for automatic upgrades, locate the shared folder in a minimal directory hierarchy. Verify that the path is accessible to the NetVault Server and all target machines where you want to install the packages. NOTE: On the Windows-based server, if the server IP address is used to configure a local package store (for example, <code>\\10.11.12.3\PkgStore</code>), the user credentials are not validated when the package store is added or updated. If you specify invalid credentials, no errors are reported. However, the deployment task fails if the authentication fails and the packages cannot be retrieved from the store. Therefore, Quest recommends that you use the server name when configuring the path to a local store (for example, <code>\\WinServer1\PkgStore</code>).
User Name	Specify a user account that can be used to access the CIFS share using one of the following formats: <ul style="list-style-type: none">• <code><AD domain>\<user name></code>• <code><NETBIOS name>\<user name></code>• <code><Local user name></code>
Password	Specify the password for the user account.

- 4 Click **Add** to add the store, and then click **OK**.

Modifying a package store

You can update the UNC path or user credentials for an existing package store.

To modify a package store:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Install Software/Add Clients**.
- 2 On the **Machines to Be Added as Clients** page, click **Install Software**.
- 3 Click **Manage Stores**, and in the package store list, select the applicable package store.
- 4 Under the package details section, type new values for the options that you want to change.
For more information about the options, see [Configuring a package store](#).
- 5 Click **Update**, and in the confirmation dialog box, click **OK**.

i | **NOTE:** When you edit a package store, it affects the existing deployment tasks that use the package store.

Removing a package store

If a package store is no longer required, you can delete it from the **Manage Stores** page.

i | **NOTE:** A package store can only be removed if it has not been used in a deployment task.

To remove a package store:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Install Software/Add Clients**.
- 2 On the **Machines to Be Added as Clients** page, click **Install Software**.
- 3 Click **Manage Stores**, and in the package store list, select the applicable package store.
- 4 Click **Remove**, and in the confirmation dialog box, click **OK**.

Performing push installations

You can deploy the NetVault Client and plug-ins to multiple machines by creating a deployment task from the WebUI to push the packages to the specified machines. You can use this method for both new and upgrade installations. After the packages are successfully installed on the machines, the new clients are automatically added to the NetVault Server.

To perform push installations:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Install Software/Add Clients**.
- 2 On the **Machines to Be Added as Clients** page, click **Install Software**.
- 3 To select a software package, next to **Software**, click the add icon (+).
- 4 In the **Select Packages for Deployment** window, select the installation packages that you want to use:
 - For NetVault Client packages, expand **Select core packages**, and then select the appropriate package.
i | **NOTE:** Be sure to select the package that corresponds to the operating system installed on the target machines. You can deploy to only one type of operating system at a time; for example, to all Windows target machines or to all Linux target machines.
 - For NetVault plug-in packages, expand **Select plug-in packages**, and then select the appropriate “.npk” binary file that you want to use.

Click **Apply**.

- 5 On the **Machines to Be Added as Clients** page, next to **Machines**, click the add (+) drop-down list, and select a method for adding the deployment targets.
- 6 In the **Add Machines** window, on the **Machine Details** tab, enter the following details based on your method selection.

Method	Description
By FQDN or IP	<p>To manually add the deployment targets, select this method.</p> <p>On the Machine Details tab, provide the following details:</p> <ul style="list-style-type: none"> • FQDN/IP Address: Specify the FQDN or IP address of the client. • User Name: Specify a user account that can be used to log in to the machine. The user name must be specified in the following format: <pre><AD domain>\<user name></pre> <p style="text-align: center;">— or —</p> <pre><NETBIOS name>\<user name></pre> <p style="text-align: center;">— or —</p> <pre><Local user name></pre> • Password: Specify the password for the user account. <p>Optionally, if you want NetVault to remember the password for this user name, select Save Credential.</p> <p>NOTE: Saved credentials are visible only to the user who saved them. They are not available to any other users.</p> <p>Click Next.</p> <p>NOTE: When configuring the User Name option, note the following:</p> <ul style="list-style-type: none"> • If the remote machine is a member of an Active Directory Domain, use a Domain Account that is in the local Administrators group of the machine. • If the remote Windows machine is in a workgroup, use the built-in Administrator account (<NETBIOS name>\<user name>). • If the remote machine is Linux-based, a Domain Account or local user with sudo privileges can perform a push installation. <p>If the built-in Administrator is disabled, you can use the following steps to enable the account and set a credentials for it:</p> <ol style="list-style-type: none"> 1 Start a command prompt with administrative privileges. 2 To enable the built-in Administrator account, type the following command: <pre>net user administrator /active:yes</pre> <p>Press Enter.</p> 3 To open the Local Security Policy editor, run secpol.msc. 4 Navigate to Security Settings > Local Policies > Security Options. 5 Ensure that the User Account Control: Use Admin Approval Mode setting is disabled for the built-in Administrator account. 6 Set a password for the built-in Administrator account, and restart the machine. <p>If the remote Windows machine is in a workgroup, you can also use a local account with administrative privileges. However, to use a local administrator account, disable the User Account Control: Run all administrators in Admin Approval Mode setting. For security reasons, disabling this setting is not recommended; using the built-in Administrator account is the recommended approach.</p>

Method	Description
From Active Directory	<p>To select machines from an Active Directory Domain, select this method.</p> <p>On the AD Details tab, provide the following details:</p> <ul style="list-style-type: none"> • AD Server Address: Type the host name or IP address of the Active Directory Server. • User Name: Specify a Domain Account that is in the local Administrators group of all target Windows machines. When pushing to Linux machines, the Domain account must have the sudo privilege in place on all Linux-based target machines. The user name must be specified in the following format: <pre><AD domain>\<user name></pre> <p>— or —</p> <pre><NETBIOS name>\<user name></pre> • Password: Specify the password for the user account. <p>Optionally, if you want NetVault to remember the password for this user name, select Save Credential.</p> <p>NOTE: Saved credentials are visible only to the user who saved them. They are not available to any other users.</p> <p>Click Connect.</p> <p>In the list of Active Directory Objects, select the target machines, and then click Next.</p>
From a file	<p>To import the target list from a file, select this method.</p> <p>In the Add machines from file dialog box, click Select a file. After selecting the file in the browse window, click OK to upload the file.</p> <p>For more information about the file format, see File format for specifying deployment targets.</p>

- 7 For each of the following tabs, to configure the applicable parameters, enter the details listed under Description.

i | **NOTE:** The **Installation Settings**, **Client Settings**, and **Firewall Settings** are only effective when a machine is being added as a client for the first time. If a machine is already added to the server, these settings are not used for that machine.

Tab	Description
Installation Settings	<p>On the Installation Settings tab, provide the following details:</p> <ul style="list-style-type: none"> Machine Name: Provide a NetVault name for the machine. The NetVault machine name can contain a maximum of 63 characters. If the machine name is longer than 63 characters, the NetVault Service may fail to start. The NetVault machine names can include uppercase letters (A–Z), lowercase letters (a–z), numbers (0–9), hyphens (“-”), underscores (“_”), and periods (“.”). For more information about NetVault names, see the <i>Quest NetVault Installation Guide</i>. If you leave the field blank, NetVault uses the host name. If multiple clients are selected from an Active Directory Domain, the label changes to Machine Name Prefix. The machine name that you provide is used as the base name, and NetVault adds “_n” to the base name while assigning client names. For example, if the Machine Name is “WinClient,” the clients are assigned the names WinClient_1, WinClient_2, ... WinClient_n. If you leave the field blank, NetVault uses the host names. Machine Password: Provide a password for the NetVault machine. The machine password can contain a maximum of 100 characters. It cannot contain the following characters: \ and spaces. Installation Folder: Specify the installation path. Installation Language: Select the language for your NetVault installation. <p>Click Next.</p>
Client Settings	<p>On the Client Settings tab, provide the following details:</p> <ul style="list-style-type: none"> Client Description: Type a descriptive text for the client. Client description can help you determine the location of the client or the role of the machine. Client Group: To add the client to one or more client groups, select the corresponding check boxes. The client groups that are configured to contain all clients are selected by default. You cannot change these selections. <p>Click Next.</p>
Firewall Settings	<p>If the client resides on a network that is outside the firewall, select the Client is Outside Firewall check box on the Firewall Settings tab, and provide the TCP/IP port specifications. For more information, see About firewall settings.</p>

8 To save the client details, click **OK**.

9 If you are performing an upgrade installation on any client, select the **Allow existing NetVault client installations to be upgraded** check box.

10 Before submitting the task, do the following:

- Click **Verify** to ensure connectivity to the clients. If any errors are reported, click **Edit**, and modify the applicable installation parameters.
- If you want to assign a user-defined task name, type the name in the **Task Name** box.

11 To submit the task, click **Install Software / Add Clients**.

You can monitor the progress and status of the task from the **Deployment Task Status** page. For more information, see [Monitoring deployment tasks](#).

i **NOTE:** The installer generates a log file in the system temporary directory. (The `TEMP` environment variable for system account determines the location of this directory; this path is typically set to `%windir%\Temp`.) The log file is named as follows: `netvault_{GUID}_install.log`, where `{GUID}` is unique for each installation occurrence (for example, `netvault_274bffb2-a3c1-457e-7f5e-221bf60a7689_install.log`).

File format for specifying deployment targets

i **NOTE:** When creating an input file, ensure the following:

- Use “.txt” or “.csv” filename extension for the input file.
- In the first row, specify the field names. Use a semicolon (“;”) to separate the fields.
- In the following rows, specify the values for the fields. Specify the field values for 1 machine per line. Use a semicolon (“;”) to separate the field values.

You can specify the following fields in the input file.

Table 15. Valid fields for input file

Field	Description
target	Specify the FQDN or IP address of the target machine. (The NetVault Server must be able to contact the machine on this address.)
targettype	Specify the machine type. Currently, only the fixed string “machine” is allowed.
targetusername	Specify a user account that can be used to log on to the machine. Use the following format to specify the user name: <ul style="list-style-type: none"> • Domain Account: <code><Domain>\<user_name></code> • User account that is local to the remote machine: <code><user_name></code>
targetuserpassword	Specify the password for the user account.
nvlanguage	Specify the installation language for NetVault. This setting is only used for new client installations.
nvmachname	Specify a NetVault name for the machine. The NetVault machine name can contain a maximum of 63 characters. If the machine name is longer than 63 characters, the NetVault Service may fail to start. The NetVault machine names can include uppercase letters (A–Z), lowercase letters (a–z), numbers (0–9), hyphens (“-”), underscores (“_”), and periods (“.”). For more information about NetVault names, see the <i>Quest NetVault Installation Guide</i> . NOTE: Make sure that you assign a unique name to each client. If you specify duplicate names, the existing client details are overwritten during client addition.
nvmachpassword	Specify a password for the NetVault machine. The password can contain a maximum of 100 characters. It cannot contain the following characters: <code>\</code> and spaces. The NetVault password is used to add and access the clients.
nvinstallfolder	Specify the installation folder for NetVault installation. This setting is only used for new client installations.
nvclientgroup	Specify the client groups to which you want to add the client. To specify multiple groups, you can use a comma (“,”) as a delimiter.
nvdesc	Type a descriptive text for the client.

Table 15. Valid fields for input file

Field	Description
outsidefirewall	Set this option to “true” if the client is outside the firewall, and specify the ports for communicating through the firewall. For more information about configuring firewall ports, see About firewall settings .
deviceslistenports	Ports to listen on for device requests.
devicesconnectports	Ports that plug-ins use to connect to remote storage devices.
msgchannellistenports	Ports for receiving messages during data transfers.
msgchannelconnectports	Ports for sending messages during data transfers.
ndmplistenports	Ports to listen on for NetVault devices operating as NDMP movers.
ndmpconnectports	Ports for sending NDMP messages (NDMP control channels).
intermachineconnectports	Ports for establishing initial contact (broadcast channels) while adding a NetVault Client, and later to ascertain its availability.

Sample input file

```
target;targettype;targetusername;targetuserpassword;nvlanguage;nvmachname;nvmachpassword;nvinstallfolder;nvclientgroup;nvdesc;outsidefirewall;deviceslistenports;devicesconnectports;msgchannellistenports;msgchannelconnectports;ndmplistenports;ndmpconnectports;intermachineconnectports

10.11.12.3;machine;testdomain\administrator;UserPwd;English;Client-A;nvpassword1;C:\Software; ClientGroup-2,default;NetVault Client-A>true;5000-5300;5000-5300;5000-5300;5000-5300;5000-5300;5000-5300;5000-5300;5000-5300

10.11.12.4;machine;testpc1\administrator;MyPassword;English;Client-B;nvpassword2;C:\Software;default;NetVault Client-B>true;5000-5300;5000-5300;5000-5300;5000-5300;5000-5300;5000-5300;5000-5300;5000-5300

10.10.25.225;machine;testpc2\administrator;MyPassword;English;Client-C;nvpassword3;C:\Software;default;NetVault Client-C>false;;;;;
```

Monitoring deployment tasks

You can monitor the status and progress of deployment tasks from the **Deployment Task Status** page. The page shows the summary of current and completed tasks in the form of bar charts and provides the progress information for individual machines.

To view the deployment task status:

- 1 In the Navigation pane, click **Deployment Task Status**.
- 2 On the **Deployment Task Status** page, you can view the following information:
 - **Task Name:** Name assigned to the task.
 - **Start Time:** Start time.
 - **Total:** Total number of clients.
 - **Success:** Number of clients on which package installation completed successfully.
 - **Warnings:** Number of clients on which package installation completed with warnings.
 - **Errors:** Number of clients on which package installation failed.
 - **Updated:** Status update date and time.
 - **Status:** Overall status of the task.

i **NOTE:** NetVault web application supports multiple-column sorting for the tables. For more information, see [Sorting records in NetVault WebUI](#).
To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower-right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

- 3 To view the status details for a particular task, select the task in the Task table, and click **View Task Target Status**.

On the Target Status page, you can view the following information:

- **Task status charts:** This area displays bar charts for the current and completed targets:
 - **Current Task Targets:** The individual bars represent the number of installation or upgrade tasks that running, deferred, or pending states.
 - **Completed Task Targets:** The individual bars represent the number of installation upgrade tasks that have completed successfully, completed with warnings, or failed.
 - **Task details table:** The table shows the following information:
 - **Task ID:** Task identification number
 - **Target:** Name of the target client
 - **Status:** Task status
- 4 To view the task logs, select the task in the Task table, and click **View Task Logs**.

(If you are on the **Deployment Task Status** page, select the task in the Task table, and click **View Task Logs**.)

On the **Deployment Task Logs** page, you can view the following information:

- **Severity:** Severity level
 - **Date:** Date and time when the log was generated
 - **Task ID:** Task identification number
 - **Target:** Name of the client for which the log was generated.
 - **Message:** Detailed log message or description.
- 5 To go back to the task status page, click **Back**.
 - 6 To cancel or retry a deployment task, select the task in the Task table, and click the corresponding button in the Operations pane.

Alternatively, click a link in the Navigation pane to open a different page.

Managing deployment tasks

This section includes the following topics:

- [Retrying failed deployment tasks](#)
- [Canceling a deployment task](#)
- [Cloning a deployment task](#)

Retrying failed deployment tasks

If a deployment task fails or completes with warnings, you can run the task again. The task runs only on those target machines where it failed previously.

To retry a deployment task:

- 1 In the Navigation pane, click **Deployment Task Status**.
- 2 In the task table, select the task that failed or completed with warnings, and click **Retry Task**.
- 3 In the confirmation dialog box, click **OK**.

Canceling a deployment task

You can cancel a deployment that is in progress from the **Deployment Task Status** page.

i **NOTE:** If a task is aborted after the deployment agent has started on the remote machine, the agent does not exit. It completes the software installation task. However, if the target is a new client, it is not added to the server.

To cancel a deployment task:

- 1 In the Navigation pane, click **Deployment Task Status**.
- 2 In the task table, select the task, and click **Abort**.
- 3 In the confirmation dialog box, click **OK**.

Cloning a deployment task

You can use the clone method to create a deployment task from an existing task. After the task is cloned, you can modify the installation packages, target clients, and other installation settings for the task.

To create a deployment task from an existing task:

- 1 In the Navigation pane, click **Deployment Task Status**.
- 2 In the task table, select the task, and click **Clone**.
- 3 Complete [Step 3](#) through [Step 11](#) in the section [Performing push installations](#).

Adding clients

This section includes the following topics:

- [Adding a client to the NetVault Server](#)
- [Adding multiple clients](#)
- [About firewall settings](#)
- [Locating a client](#)
- [Checking communication through a firewall](#)
- [Adding a server as a client](#)

Adding a client to the NetVault Server

To use a client in a backup or restore operation, you must first add the client to the NetVault Server. You can use the configuration wizard to add and configure a NetVault Client.

i | **NOTE:** A NetVault Server can only support clients of the same version as itself or older. The NetVault Client software version cannot be higher than the NetVault Server software version.

To add a client to the NetVault Server:

- 1 In the Navigation pane, click **Manage Clients**, and then click **Add Client**.

On Linux-based machines, you can also start the wizard from the **Guided Configuration** link. In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Clients**.

- 2 In the list of available clients, select the client that you want to add, and click **Next**.

The list includes machines that are automatically discovered on the network broadcast range. These are machines on which you have installed the NetVault Server or Client software.

You can view the following information about the machines:

- **Status:** Displays the status icons, which indicate whether the client is online or offline.
- **Client:** Displays the NetVault name assigned to the client.
- **Version:** Displays the NetVault version installed on the machine.
- **Description:** Displays the client description.

- 3 If the client is password-protected, type the NetVault password for the client. You can also provide the root or administrator password for the client.

If a NetVault password is not set for the client, provide the root or administrator password for the client.

Click **Next**.

This page is not displayed if security is disabled on the client. For more information, see [Configuring security settings](#).

- 4 Type a descriptive text for the client, and click **Next**.

- 5 If the client resides on a network that is outside the firewall, select the **Client is Outside Firewall** check box, and configure the firewall settings. For more information, see [About firewall settings](#).

Click **Next**.

- 6 After the client is successfully added to the server, a message is displayed.

Adding multiple clients

When adding clients to a Windows-based or Linux-based NetVault Server, can use the following procedure to add multiple clients at the same time.

i | **NOTE:** A NetVault Server can only support clients of the same version as itself or older. The NetVault Client software version cannot be higher than the NetVault Server software version.

To add multiple clients to a NetVault Server:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the NetVault Configuration Wizard page, select the **Install Software/Add Clients** option.
- 2 On the **Machines to Be Added as Clients** page, next to **Machines**, click the add (+) icon, and select a method for adding the deployment targets.
- 3 In the **Add Machines** window, on the **Machine Details** tab, enter the following details based on your method selection.

Method	Description
By FQDN or IP	To manually add the deployment targets, select this method. On the Machine Details tab, specify the FQDN or IP address of the client. Click Next .
From Active Directory	To select machines from an Active Directory Domain, select this method. On the AD Details tab, provide the following details: <ul style="list-style-type: none"> • AD Server Address: Type the host name or IP address of the Active Directory Server. • User Name: Specify a Domain Account that is in the local Administrators group of all target machines. The user name must be specified in the following format: <AD domain>\<user name> — or — <NETBIOS name>\<user name> — or — <user name>@<AD domain> • Password: Specify the password for the user account. <p>Optionally, if you want NetVault to remember the password for this user name, select Save Credential.</p> <p>NOTE: Saved credentials are visible only to the user who saved them. They are not available to any other users.</p> <p>Click Connect.</p> <p>In the list of Active Directory Objects, select the target machines, and click Next.</p>
From a file	To import the target list from a file, select this method. In the Add machines from file dialog box, click Select a file . After selecting the file in the browse window, click OK to upload the file. For more information about the file format, see File format for specifying deployment targets .
From available machines	To add a client from available machines, select this method. On the NetVault Machine Details tab, select the client that you want to add. Provide the NetVault password in the Machine Password box. Click Next .

4 Configure the applicable parameters.

Option	Description
Installation Settings	On the Installation Settings tab, provide the NetVault password. If no NetVault password is set for the client, specify the administrator password for the machine. Click Next .

Option	Description
Client Settings	<p>On the Client Settings tab, provide the following details:</p> <ul style="list-style-type: none"> • Client Description: Type a descriptive text for the client. Client description can help you determine the location of the client or the role of the machine. • Client Group: To add the client to one or more client groups, select the corresponding check boxes. The client groups that are configured to contain all clients are selected by default. You cannot change these selections. <p>Click Next.</p>
Firewall Settings	<p>If the client resides on a network that is outside the firewall, select the Client is Outside Firewall check box on the Firewall Settings tab, and provide the TCP/IP port specifications. For more information about these settings, see About firewall settings.</p>

5 To save the client details, click **OK**.

6 Before submitting the task, do the following:

- Click **Verify** to check connectivity to the clients. If any errors are reported, click **Edit**, and modify the applicable installation parameters.
- If you want to assign a user-defined task name, type the name in the **Task Name** box.

7 To submit the task, click **Add Clients**.

You can monitor the progress and status of the task from the **Deployment Task Status** page. For more information, see [Monitoring deployment tasks](#).

About firewall settings

Firewall settings are required to communicate with NetVault Clients that reside outside the firewall. You can use these settings to specify the TCP/IP ports for establishing data transfer channels, message channels, and broadcast channels through the firewall.

You can configure the firewall settings when deploying or adding a client, or update these settings from the **Change Settings** page. You can also use the NetVault WebUI or Txtconfig utilities to configure the firewall settings for new or existing clients.

i NetVault does not support firewalls using NAT (Network Address Translation)/IP Masquerading.

The following table provides a brief description of the firewall settings.

Table 16. Firewall settings

Option	Description
Listen ports for devices	Ports to listen on for device requests. Configure this option on the NetVault machines that have a locally attached device (for example, NetVault Server or NetVault Clients with SmartClient licenses). Requirement: Two ports per drive.
Connect ports for devices	Ports that plug-ins use to connect to remote storage devices. Configure this option on clients that connect to remote devices. Requirement: Two ports per drive.
Listen ports for NetVault message channels	Ports for receiving messages during data transfers. Configure this option on both the NetVault Server and the Client. NetVault requires a two-way connection between the Server and the Client for message channels. Requirement: Three ports per client. To run two or more plug-ins simultaneously on a client, configure two ports per plug-in and an extra port per client. For example, to run two plug-ins simultaneously, configure $(2 * 2) + 1 = 5$ ports for a client.
Connect ports for NetVault message channels	Ports for sending messages during data transfers. Configure this option on both the NetVault Server and the Client. NetVault requires a two-way connection between the Server and the Client for message channels. Requirement: Three ports per client. To run two or more plug-ins simultaneously on a client, configure two ports per plug-in and an extra port per client. For example, to run two plug-ins simultaneously, configure $(2 * 2) + 1 = 5$ ports for a client.
Connect ports for NDMP control channels	Ports for sending NDMP messages (NDMP control channels). Configure this option on the NetVault Server (on which the plug-in is installed) when a firewall separates an NDMP filer and the NetVault Server. By default, NetVault uses port number 10000. You can change it, if necessary.
Listen ports for NDMP data channels	Ports to listen on for NetVault devices operating as NDMP movers. Configure this option on the NetVault Server or Client to which the device is attached. These ports are used for data transfers between the NDMP filer and storage device when a firewall separates the two networks.
Connect ports for inter-machine setup	Ports for establishing initial contact (broadcast channels) while adding a NetVault Client, and later to ascertain its availability. Requirement: Two ports per client.

You can use the following formats to specify the ports or port ranges for data channels, message channels, and broadcast channels:

- A comma-separated list (example: 20000, 20050)
- A port-range separated by a dash (example: 20000-20100)
- A combination of comma-separated list and port-range (example: 20000-20100, 20200)

i | **NOTE:** You must configure the same port-range for data, message, and broadcast channels across all NetVault machines.

i | **IMPORTANT:**

- NetVault does not support firewalls using NAT (Network Address Translation)/IP Masquerading.
- NetVault uses port number 20031 for TCP and UDP messaging. Make sure that this port is open on the firewall.

For more information, see the following sections:

- [Firewall filtering rules](#)
- [Firewall configuration example](#)

Firewall filtering rules

When creating firewall rules on the server and client machines, make sure that you open the following ports to send and receive traffic from NetVault.

Table 17. Firewall filtering rules

From	To	TCP/UDP	Source port	Destination port
Server	Client	TCP	Connect ports for inter-machine connection setup specified on the NetVault Server	20031
Client	Server	TCP	Connect ports for inter-machine connection setup specified on the NetVault Clients	20031
Server	Client	TCP	Connect ports for message channels specified on NetVault Server	Listen Ports for Message Channels specified on the NetVault Clients
Client	Server	TCP	Listen ports for message channels specified on the NetVault Clients	Connect Ports for Message Channels specified on the NetVault Server
Server	Client	UDP	20031	20031
Client	Server	UDP	20031	20031
Server	Client	TCP	Listen ports for devices specified on the NetVault Server and Clients	Connect ports for devices specified on the NetVault Clients
Client	Server	TCP	Connect ports for devices specified on NetVault Clients	Listen ports for devices specified on NetVault Server and Clients

Firewall configuration example

This example illustrates the network port requirements for a NetVault system with the following configuration:

- Number of drives: 6
- Number of clients with one plug-in: 10
- Number of clients with two plug-ins: 2

Table 18. Example: Port requirement calculation

NetVault machine	Port type	Requirement	Total ports	Example port range
Server (with locally attached storage device)	Connect ports for inter-machine connection setup	Minimum two ports per client	24	50300-50323
	Connect ports for NetVault message channels	Minimum three ports per client	40	50200-50239
	Listen ports for devices	Minimum two ports per drive	12	50100-50111
Client	Connect ports for inter-machine connection setup	Minimum two ports per client (These ports can be the same as the ports specified on the server side.)	24	50300-50323
	Listen ports for NetVault message channels	Minimum three ports per client	40	50500-50539
	Connect ports for devices	Minimum two ports per drive	12	50400-50411

The following table illustrates the firewall filtering rules for this system.

Table 19. Example: Firewall filtering rules

From	To	TCP/UDP	Source port	Destination port
Server	Client	TCP	50300-50323	20031
Client	Server	TCP	50300-50323	20031
Server	Client	TCP	50200-50237	50500-50537
Client	Server	TCP	50500-50537	50200-50237
Server	Client	UDP	20031	20031
Client	Server	UDP	20031	20031
Server	Client	TCP	50100-50111	50400-50411
Client	Server	TCP	50400-50411	50100-50111

Locating a client

The clients that reside in a different subnet are not included in the list of available clients. You can use the **Find Machine** option to locate and add such clients.

To locate and add a client that is not discovered automatically:

- 1 Start the client addition wizard, and then click **Find Machine**.
- 2 On the Find Client page, type the Fully Qualified Domain Name (FQDN) or IP address of the client, and click **Find**.
- 3 After the machine is located, complete [Step 3](#) through [Step 6](#) in the section [Adding a client to the NetVault Server](#).

- i** **NOTE:** NetVault reports an error if it fails to locate the specified client on the network. An error can occur for reasons such as the following:
- The NetVault software is not installed on the machine.
 - The NetVault Service is not running on the machine.
 - The DNS lookup table or the machine's host table cannot be contacted.

Checking communication through a firewall

When adding a client that resides outside the firewall, you can perform a firewall test to check if the server and client can communicate through the firewall. You can also perform this test to check connectivity for existing clients.

To check if the server and client can communicate through a firewall:

- 1 To check connectivity for new clients, start the client addition wizard, and then click **Firewall Test**.

— or —

To check connectivity for existing clients, open the **Manage Clients** page. Select the client, and then click **Firewall Test**.

- 2 In the **Check Connection** dialog box, provide the following details.

Table 20. Check connection

Option	Description
NetVault Client Name	For new clients, type the NetVault name of the machine. For existing clients, the name is displayed as a read-only property.
NetVault Client Address	For new clients, type the IP address of the machine. For existing clients, the IP address is displayed as a read-only property.
UDP Port	The default UDP port for NetVault. It is set to 20031. If you have configured NetVault to use a different port, change this value.
TCP Port	The default UDP port for NetVault. It is set to 20031. If you have configured NetVault to use a different port, change this value.
Timeout	The timeout interval is set to 15 seconds by default.

- 3 Click **Test**. This command checks if TCP, UDP, and messaging connectivity is available between the server and client, and displays the result in the dialog box.

- i** **IMPORTANT:** The firewall test option checks the connectivity to the Stats Manager process on the specified server or client machine. If this process is not running on the client, the test fails. However, the client may still be accessible for backup. In any case, it is an abnormal situation and should be corrected.

Adding a server as a client

To allow a NetVault Server to function as a client to a different server, you must enable the setting **This machine may be added as a client to a server** on the server.

To allow a NetVault Server to function as a client to a different server:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **System and Security**, click **Security**. In the **Security** dialog box, configure the following setting:
 - **This machine may be added as a client to a server:** The security settings on a NetVault Server prevent the server from being added as a client to a different server. To allow a NetVault Server to function as a client to a different serve, select this check box.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Removing a client from the list of available clients

If the list of available clients includes a machine that is no longer in use, you can use the following procedure to remove it from the list.

To remove a defunct client from the list of available clients:

- 1 Start the client addition wizard.
- 2 In the **NetVault Clients** table, select the defunct client, and click **Remove**.
- 3 In the confirmation dialog box, click **Remove**.
- 4 If NetVault has been removed or stopped, the following message is displayed:
Client <client name> is not responding. Unable to remove this client.
If this error is displayed, click **Force Removal** in the **Error** dialog box to remove the client.

Managing clients

This section includes the following topics:

- [Viewing client details](#)
- [Setting client description](#)
- [Installing plug-ins](#)
- [Checking for upgrades](#)
- [Installing a product license file](#)
- [Checking client access](#)
- [Removing plug-ins](#)
- [Removing a client from the server](#)

Viewing client details

You can view all existing clients from the **Manage Clients** page. The page displays the client name, description, status, and product version. To view more information about a client, select the client, and click **Manage**.

To view client details:

- 1 In the Navigation pane, click **Manage Clients**.

On the **Manage Clients** page, you can view all NetVault Clients that are added to the server. The page also displays Virtual Clients. The page shows the following information about the machines:

- **Status:** Displays the status icons, which indicate the type of client and whether the client is online or offline.
- **Client:** Displays the NetVault name assigned to the client.
- **Version:** Displays the NetVault version installed on the machine.

- **Description:** Displays the client description.

The following table provides a brief description of the client icons.

Table 21. Client icons

Icon	Description
	Client is up and running.
	Client is online. It is in the process of being added, or the NetVault password for the client has changed since it was added.
	Client is unavailable. The system is offline or the NetVault Service is not running.
	Represents a Virtual Client that consists of a cluster of Clients. For more information about Virtual Clients, see Working with client clusters .

- 2 By default, the table is sorted by client name.

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

- 3 To quickly find entries containing specific field values or text, you can use the Search box at the upper-right corner of the table. The data in the table is filtered as you type the search text into the box. You can also include wildcards (“?” or “*”) in the search filter string.
- 4 In the **NetVault Clients** table, select the client that you want to view, and click **Manage**.
- 5 On the **View Client** page, you can view the following details:
 - **Client Summary:** The **Client Summary** area displays the following information: NetVault machine name, client description, NetVault version, machine ID, network name of the machine, IP address, release information, and OS.
 - **Server Capabilities:** To view the server license capabilities, click this link. You can view the license capabilities for clients, SmartClients, and various types of devices.
 - **License Key Details:** To view the license information for the server or client and any licensed plug-ins installed on the machine, click this link.
 - **Installed Plug-ins:** The Installed Plug-ins table displays the plug-ins installed on the selected client. The details include the plug-in name, version number, and installation date.
- 6 To perform a client-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

Setting client description

You can use the following procedure to set a description for a client. Client description can help you determine the location of the client or the role of the machine.

To set or change client description:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the **NetVault Clients** table, select the applicable client, and click **Manage**.
- 3 On the **View Client** page, under the **Client Summary** area, you can view the client description.

To set or change the client description, click the Edit icon. In the **Edit Description** dialog box, type a descriptive text for the client.

- 4 Click **Save** to save the details and close the dialog box.

Installing plug-ins

NetVault offers a selection of licensed plug-ins that you can install on the server and client machines to protect application-specific data. You can install the plug-ins simultaneously on multiple machines by using the configuration wizard. Alternatively, you can install a plug-ins on a single client from the **Manage Clients** page.

The following sections describe the different procedures that you can use to install the NetVault plug-ins:

- [Installing plug-ins using the push installation method](#)
- [Installing plug-ins from the Manage Clients page](#)

Installing plug-ins using the push installation method

On Windows-based and Linux-based machines, you can use the push installation method to install the plug-ins on several machines at the same time. You can perform push installations from the NetVault WebUI.

Before you start the push installation procedure, verify that the following requirements are met:

- **Copy the packages to a shared location:** Copy the client and plug-in packages to a shared location. Only CIFS shares (including Linux Samba shares) are currently supported as package stores. The path must be accessible to the NetVault Server and all target machines where you want to install the packages.

Make sure that you use the original names of the installation packages. Renamed packages cannot be selected for push installations.

- **Configure a package store in NetVault:** After copying the installation packages, configure the shared location details in NetVault. For more information, see [Configuring a package store](#).

To install one or more plug-ins:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Install Software/Add Clients**.
- 2 On the **Machines to Be Added as Clients** page, select **Install Software**.
- 3 In the **Select Packages for Deployment** window, select the installation packages that you want to use.
- 4 For NetVault plug-in packages, expand **Select plug-in packages**, and then select the appropriate “.npg” binary file that you want to use.
- 5 Click **Apply**.

i **NOTE:** Be sure to select the plug-in that corresponds to the operating system installed on the target machines. You can deploy to only one type of operating system at a time; for example, to all Windows target machines or to all Linux target machines.

- 6 On the **Machines to Be Added as Clients** page, next to **Machines**, click the add (+) drop-down list, and select a method for adding the deployment targets.
- 7 On the **NetVault Machine Details** tab, select the client that you want to add, and click **OK**.
- 8 To add more machines, repeat [Step 6](#) and [Step 7](#).
- 9 Before submitting the task, do the following:
 - Click **Verify** to ensure connectivity to the clients. If any errors are reported, click **Edit**, and modify the applicable installation parameters.
 - If you want to assign a user-defined task name, type the name in the **Task Name** box.
- 10 To submit the task, click **Install Software/Add Clients**.

You can monitor the progress and status of the task from the **Deployment Task Status** page. For more information, see [Monitoring deployment tasks](#).

Installing plug-ins from the Manage Clients page

From the **Manage Clients** page, you can install a plug-in on a single client. This method is available to both Linux and Windows systems.

To install a plug-in from the Manage Clients page:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the **NetVault Clients** table, select the client on which you want to install the plug-in, and click **Manage**.
- 3 At the lower-right corner of the Installed Plug-ins table, click the Install Plugin button ().
- 4 Click **Choose Plug-in File**, and in the browse window, navigate to the location of the “.npk” installation file for the plug-in (on the installation CD or the directory to which the file was downloaded from the website).
- 5 Select the platform-specific binary file for the plug-in, and click **Install Plugin** to begin installation.

After the plug-in is installed successfully, a message is displayed.

Checking for upgrades

Use this procedure to check whether there are upgrades available for NetVault or for plug-in packages.

To check for upgrades for the Core and plug-ins:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 Below the **NetVault Clients** table, click **Check Upgrades**.

The upgrades for the client software and plug-ins display on the right under **Available Software Upgrades**.

- 3 Enter the following details:

Option	Description
User Name	Enter the user name for the NetVault installation. NOTE: When upgrading Windows clients from a Linux NetVault Server, local administrator account credentials are required.
Password	Enter the password for the NetVault installation. NOTE: When upgrading Windows clients from a Linux NetVault Server, local administrator account credentials are required.
Save Credential	Optional. Select this option to save the credentials for the installation.
Package Store	Select the package store where the packages that can be upgraded are available. NOTE: To reduce the time it takes to browse the package store for automatic upgrades, the shared folder should exist at the highest possible level in the directory.
Task Name	Optional. Enter a name for the upgrade task. The default name uses the date and time of the task.
Search	Use this option to find a specific client that you want to upgrade.

- 4 Expand each client that you want to upgrade and select the upgrades that you want to include.

To include all available upgrades, select the client.

i | **NOTE:** All NetVault and plug-in packages that are available to upgrade are selected by default.

- 5 Optionally, to change the credentials for a client, complete the following steps:
 - a Click the icon for the client in the **Edit Credentials** column.
 - b In the **Enter credentials** dialog box, enter the user name and password for the client, optionally select **Save Credentials**, and then click **OK**.
- 6 Click **Verify** to check connectivity to the clients.
If any errors are reported, click **Edit**, and modify the applicable installation parameters.
- 7 To submit the task, click **Upgrade**.
You can monitor the progress and status of the task from the **Deployment Task Status** page. For more information, see [Monitoring deployment tasks](#).

Installing a product license file

The trial license for NetVault is valid for 30 days. A band across the top of the WebUI displays a message about when the license is scheduled to expire. The color of the band, as described in the following list, indicates the number of days remaining in the trial period:

- **Yellow:** The license has less than 21 days before it expires.
- **Light orange:** The license has less than 14 days before it expires.
- **Dark orange:** The license has less than 7 days before it expires.
- **Red:** The license expired.

i | **NOTE:** You can close the yellow and orange bands, but the red band remains until you install the permanent license.

To continue using the software after the expiry of the evaluation period, you must go to the **License Management** page and install the permanent license file for the server and installed plug-ins.

The server is licensed based on the NetVault Server Edition and the additional options that you have purchased. The clients do not require a permanent license file, unless there is a plug-in older than version 13.0 installed on the machine. Beginning with release 13.0, Clients and plug-ins do not require licenses installed on Client machines.

The permanent license file for NetVault is tied to the particular server where it is installed using the Machine ID. To obtain the permanent license file, locate the relevant Machine ID, and submit the request using the online licensing form.

fileInstalling the license file

To install the license file:

- 1 In the Navigation pane, click **License Management**.
- 2 On the **License Management** page, click **Install License**.
Because the NetVault WebUI is tied to a specific NetVault Server and all licenses are installed on the NetVault Server, the program automatically installs the license on the correct machine.
- 3 In the **Choose License File/s** dialog box, click **Select a file**, point to the “.dlv” file that contains the license key, and click **Open**.
- 4 In the dialog box, click **OK** to apply the selected license file.
After the file is installed successfully, a message is displayed.
- 5 To dismiss the message, click **X**.

Checking client access

You can use the **Check Access** option to determine the accessibility status of a client.

To check access to a client:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the **NetVault Clients** table, select the client, and click **Manage**.
- 3 On the **View Client** page, click **Check Access**.

The NetVault Server tries to connect to the client, and returns a message indicating the current accessibility status of the client. Click the Close button to close the dialog box.

Removing plug-ins

If a licensed plug-in is no longer required, you can remove it from the server or client machine.

To remove a plug-in from the NetVault Server or Client:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the **NetVault Clients** table, select the client, and click **Manage**.
- 3 In the Installed Plug-ins table, select the plug-in, and click the Remove Plugin button ().
- 4 In the confirmation dialog box, click **Remove**.

Removing a client from the server

If a client is no longer used, you can remove it from the NetVault Server.

To remove a client from the NetVault Server:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the **NetVault Clients** table, select the client, and click **Manage**.
- 3 On the **View Client** page, click **Remove Client**.
- 4 In the confirmation dialog box, click **Remove**.

Managing client groups

This section includes the following topics:

- [About client groups](#)
- [Creating a client group](#)
- [Viewing existing client groups](#)
- [Modifying a client group](#)
- [Removing a client group](#)

About client groups

NetVault lets you group the clients into one or more logical entities.

You can configure a client group to contain all clients or specific clients. The new clients automatically become members of the client groups that are configured to include all clients. You can create any number of client groups on the server and add a client to multiple client groups.

The NetVault Administrator can use client groups to control user access to clients. The administrator can grant access to specific clients by using a client group. Client groups are also useful when creating policy-based backups. Instead of individually selecting each client, you can select the client group to add multiple clients.

NetVault includes a pre-configured client group named **default**, which is originally configured to include **All Clients**. Therefore, when you add a client, it is automatically added to the default group, unless you have reconfigured this group to not include **All Clients**.

Creating a client group

You can group the NetVault Clients into one or more logical entities. A client group can contain all or specific clients. The new clients automatically become members for client groups that are configured to include all clients.

To create a client group:

- 1 In the Navigation pane, click **Manage Clients**, and then on the **Manage Clients** page, click **Manage Client Groups**.
- 2 On the **Client Group Management** page, click **New Group**. On the **Client Group** page, provide the following information
 - In **Group Name**, type a name for the client group.
 - In **Group Description**, provide a detailed description for the client group.
 - To add all clients to the group, select the **All Clients** check box. When you select this check box, the new clients are automatically added to the group.
To add a specific client, select the target client in the **Available Clients** table, and click the Add button (+) to the left of the item. When you click this button, the selected client is moved to the **Chosen Clients** table.
 - To remove a client from the group, select the target client in the **Chosen Clients** table, and click the Remove button (-) to the left of the item. When you click this button, the selected client is moved to the **Available Clients** table.
- 3 To add the group, click **Create Group**.

Viewing existing client groups

You can view information about the existing client groups from the **Client Group Management** page. The page displays the group name, number of clients, and the list of member clients.

To view the existing client groups:

- 1 In the Navigation pane, click **Manage Clients**, and then on the **Manage Clients** page, click **Manage Client Groups**.
- 2 On the **Client Group Management** page, you can view the existing client groups.

The table shows the following information about the client groups:

- **Group name:** Displays the name of the client group.
- **All Clients:** The check box is selected if the group is configured to include all clients.

- **Number of Clients:** Displays the number of clients in the group.
 - **Description:** Displays description for the client group.
- 3 You can click the client group name to display the list of clients included in the group. (The list of clients is only displayed for groups that do not have the **All Clients** check box selected.)
 - 4 To perform a client group-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

Modifying a client group

You can modify the group name, group description, or group members for an existing client group.

To modify a client group:

- 1 In the Navigation pane, click **Manage Clients**, and then on the **Manage Clients** page, click **Manage Client Groups**.
- 2 On the **Client Group Management** page, select the applicable client group, and click **Edit**.
- 3 Modify the required settings. For more information, see [Creating a client group](#).
- 4 To save the settings, click **Create Group**.

Removing a client group

If a client group is no longer required, you can remove it from the server. When you delete a group, the machines that are in the group are not removed.

To remove a client group:

- 1 In the Navigation pane, click **Manage Clients**, and then on the **Manage Clients** page, click **Manage Client Groups**.
- 2 On the **Client Group Management** page, select the applicable client group, and click **Remove**.
- 3 In the confirmation dialog box, click **Remove**.

Managing catalog search

- [About catalog search](#)
- [Configuring catalog search](#)
- [Modifying catalog search](#)
- [Disabling catalog search](#)
- [Removing catalog search](#)

About catalog search

The NetVault Server version 12.0 and later running on Windows and Linux operating system supports catalog searching on NetVault Server, and the NetVault Server version 12.2 and later running on Windows supports catalog searching on NetVault Server and NetVault Windows client of the same version. The NetVault catalog search service provides fast-search capabilities for the metadata (name, mod date, type etc) of items backed up by NetVault. Catalog search provides search capabilities for the tags (Filename) based on string data-type.

Catalog search results provides the following information:

- List of records matching the search query.
- Result paging: Retrieve results in pages for large result sets.
- Result collapsing: If a file is backed up multiple times in different backups, then only one match for that file is displayed with the option to drill- down to find the specific backups holding the file.
- Faceting and search narrowing: Displays the number of search results matching a particular facet. This information is available by applying filters to the search results. For example, the number of results from a particular plug-in or machine.

The following plug-ins support catalog searching:

- NetVault Plug-in *for FileSystem*
- NetVault Plug-in *for VMware*
- NetVault Plug-in *for Hyper-V*
- NetVault Plug-in *for SnapMirror to Tape*
- NetVault Plug-in *for NDMP* (plug-in version 12.2 or later)
- NetVault Plug-in *for SQL Server*
- NetVault Plug-in *for O365* (plug-in version 12.4 or later)

Prerequisites

Before you start configuring the NetVault, verify that the following requirements are met:

- Catalog search reserved RAM size (For Windows and Linux) must not be more than 50% of the physical RAM. For low memory systems (i.e. 8 GB or less), Quest recommends to set the reserved RAM size to no more than 2 GB.

- Virtual memory (only for Linux) mmap counts on Linux is 65530, which is a default operating system limit on mmap counts. Change this value to 262144 using the following command:

```
sysctl -w vm.max_map_count=262144
```

To permanently set this value to persist after the reboot, in `/etc/sysctl.conf`, update the parameter `vm.max_map_count`.

To verify the mmap count after rebooting, run the command `sysctl vm.max_map_count`

- Catalog search port requirement: 9200 - 9300.
- Elasticsearch version 6.2 is compatible with many NetVault Server supported platforms. For more information, see <https://www.elastic.co/support/matrix>.

i | **NOTE:** Catalog search installation is not supported on Debian 7 or lesser version.

- Determine NetVault catalog search Linux user account.

i | **NOTE:** Sometimes, due to policy restrictions on your system, the NetVault on Linux operating system is unable to create new catalog search Linux user account automatically.

While installing NetVault catalog search, you can select any of the following options for catalog search Linux user account:

- **Allow NetVault to create catalog search Linux user:** The default option to allow NetVault to create Catalog search Linux user account automatically (default value: netvault-catalog on Linux).
- **Linux user name:** Create a new Linux user account by referring to the corresponding operating system guide. Ensure that the specified catalog search Linux user belongs to the group <user name>. Create the group with the same name as of the Linux user name and add the user into the same group.

i | **NOTE:** While creating a catalog search Linux user account, do not use 'root' user in the field 'Linux use name'. In , catalog search cannot run by a 'root' user due to restrictions from Elasticsearch component.

i | **NOTE:** Catalog search sets the system level JAVA_HOME environment variable that is required for Elasticsearch. Hence, the other java applications may not work. It is assumed that the NetVault Server is a dedicated server.

Configuring catalog search

To enhance searching for files in savesets, you must install and configure the catalog search service on NetVault Server running on Windows or Linux operating systems.

To install and configure the catalog search:

- 1 In the Navigation pane, click **Catalog Search**.
- 2 On the **Catalog Search Configuration** page, provide the following details.

Table 22. Catalog search configuration options

Option	Description	
Host Information	Catalog Host	For Linux: Displays the name of the NetVault Server to install the catalog search server. For Windows: Displays the name of NetVault Server and Client to install catalog search. However, NetVault Server and Client must be of the same version.
	Catalog Search Password	Type the password for the catalog search database. It cannot contain the following characters: <code>& < > ^ spaces</code> By default, 'admin' is used as the password if this field is kept empty during catalog search installation.
	Confirm Catalog Search Password	Type the password again for confirmation.
	Existing Linux User (for Linux only)	For Linux operating system, type the existing Linux user name. Do not use the 'root' as a user name. For more information, see Prerequisites .
	Database Directory	Click Browse to select the database directory for catalog search: <ul style="list-style-type: none"> • 'Database Directory' drawer is displayed on the right-side of the page. • Open the parent node, browse through the directory tree, and locate the directory to store the catalog search database. • The selected directory path is displayed in the text box. • Click OK. <p>NOTE: To enter the Database Directory path manually, ensure that all the directory(s) and sub-directories must be already created and available.</p>
	Use same path for Snapshot and Log Directory	By default, the check box is selected and the Snapshot Directory and Log Directory fields are disabled, hence the Database Directory path is copied and same path is set for Snapshot Directory and Log Directory . Clear the check box to enable and edit the Snapshot Directory and Log Directory field.
	Snapshot Directory	Click Browse to select the location to store the catalog search database backup. <ul style="list-style-type: none"> • 'Snapshot Directory' drawer is displayed on the right-side of the page. • Open the parent node, browse through the directory tree, and locate the directory to store the catalog search database backup. • The selected directory path is displayed in the text box. • Click OK. <p>NOTE: To enter the Snapshot Directory path manually, ensure that all the directory(s) and sub-directories must be already created and available.</p>

Table 22. Catalog search configuration options

Option	Description
Log Directory	<p>Click Browse to select the location to store the catalog search logs.</p> <ul style="list-style-type: none"> • ‘Log Directory’ drawer is displayed on the right- side of the page. • Open the parent node, browse through the directory tree, and locate the directory to store the catalog search logs. • The selected directory path is displayed in the text box. • Click OK.
Server Snapshot Path	<p>NOTE: To enter the Log Directory path manually, ensure that all the directory(s) and sub-directories must be already created and available.</p> <p>Specify the shared directory path of Snapshot Directory. The format of Server Snapshot Path is <\\IP address of client>\shared snapshot directory>. Ensure the snapshot directory is shared to make it available for NetVault database catalog backup and restore operation on Windows Client.</p> <p>This field is available and visible only if the Catalog Host is selected as remote windows client.</p> <p>Once catalog search is installed on Windows Client, configure client host credentials:</p> <p>To configure client host credentials:</p> <ol style="list-style-type: none"> 1 Click Configure, to configure the connection details for accessing shared server snapshot path of client host. 2 Enter the following detail, and click Ok. <ul style="list-style-type: none"> • Domain: Type the Windows Domain name for the system that contains the shared folders. • Username: Type a Domain Administrator user name. The network share backups must be performed using a Domain Administrator account to fully retrieve the file and directory permissions during a restore. A user belonging to the Administrators Group does not have the Domain Administrator privileges.
HTTP Port	<p>NOTE: If a Non-Domain Administrator account is configured for backups, the file and directory permissions must be set manually after a restore.</p> <ul style="list-style-type: none"> • Password: Type the password for the user account. <p>Specifies the HTTP port for Catalog Search.</p> <p>You can specify the port number range from 9200-9299 for catalog search on which the NetVault Server will communicate. You can also configure the port after NetVault Server is installed, and then restart catalog search service.</p>
Configuration Reserved RAM Size (GB)	<p>This is the amount of RAM that the catalog search service attempts to pre-allocate before it starts. If the RAM isn’t available the service won’t start.</p> <p>By default, the value is set to 1GB. Type or select the required RAM size. For systems with Large RAM size, a higher value is recommended (2 – 16 GB, based on the item backup rate and availability of RAM, limited by the restrictions mentioned in the Prerequisites section).</p>

Table 22. Catalog search configuration options

Option	Description
Maximum simultaneously active catalog workers	<p>By default, 5 concurrent catalog workers can run on the server.</p> <p>To change the number of catalog workers that can be run in parallel at any time, type or select the new value.</p> <p>Minimum value: 1</p> <p>Maximum value: 20</p>
Add pre-existing and scanned backups to catalog	<p>Select this check box to include the pre-existing and scanned backup jobs for catalog searching. By default, the check box is clear.</p>
Include backups with offline indexes	<p>Select this check box to include the backup jobs with offline indexes for catalog searching. By default, the check box is clear.</p> <p>You can select this option only if the option Add pre-existing and scanned backups to catalog is selected.</p>
Offline index load wait (Minutes)	<p>This is the amount of time the catalog search service will wait for an offline index for a pre-existing backup to be loaded. If the media manager doesn't load the offline index within this wait time, cataloging of this pre-existing backup will be temporarily skipped. Eventually, the skipped pre-existing backup will be cataloged, with the timing depending on how many other pre-existing backups are available for cataloging.</p> <p>Type or select the duration (in minutes) to wait for loading the offline index during catalog indexing. By default, the minimum value is set to 10 minutes.</p>
Startup offline index load wait (Minutes)	<p>This is the minimum time from the last NetVault start time before the catalog search service requests the media manager to load an offline index. This delay from the last NetVault start is required to allow the media manager to complete its device scan to determine which devices are available. If a device is unavailable or hasn't been scanned to determine its availability at the time of receipt of an offline index request, the media manager fails the request without trying to load the offline index.</p> <p>Type or select the duration (in minutes). By default, the minimum value is set to 10 minutes.</p>
Restrict the age of backups added to the catalog	<p>Select this check box to restrict the pre-existing backup jobs for catalog searching.</p> <p>By default, the check box is clear.</p> <p>You can select this option only if the option Add pre-existing and scanned backups to catalog is selected.</p>
Add backups completed on and after	<p>To include the backup jobs completed on and after a specific date in the catalog search, type the date, or click the button next to the box, and select the date.</p>

i **NOTE:** You cannot change the directories (Database Directory, Snapshot Directory, and Log Directory) after catalog search configuration. If you must change the directories, you can uninstall the catalog search and then reinstall using different directories. However, you have to re-catalog your existing backups, which can take a long time.

3 Click **Install And Configure**.

Relocating the catalog search database directory

To relocate the catalog search database directory, use the procedures described in the following sections:

- [Changing the NetVault Catalog Database directory on a Windows-based machine](#)
- [Migrating NetVault Catalog Server to different Windows-based client machine](#)
- [Changing the NetVault Catalog Database directory on a Linux-based machine](#)

Changing the NetVault Catalog Database directory on a Windows-based machine

To change the Catalog Database Directory on a Windows-based NetVault Server:

- 1 Disable the catalog search, see, [Disabling catalog search](#).
- 2 Stop NetVault Catalog Search service from the Windows Services Management Console.
- 3 Manually move the **db** directory to the new location. Make sure that the user under which the **netvault-catalog** service runs has the ownership of the new Database Directory.

i | **NOTE:** Database is supported only on local drives of the system.

- 4 In the NetVault installation directory, open the **catalog.cfg** file to edit, available in **config** directory.
In the [Node1] section, change the following to specify the new location:

```
path.repo, path.repo.new, path.data, path.data.new, path.logs, path.logs.new
```

- 5 In the NetVault installation directory, open the **elasticsearch.yml** file to edit available in the following directory: **catalog\elasticsearch\config**.

Change the following to specify the new location:

```
path.repo, path.data, path.logs
```

- 6 Start NetVault Catalog Search service from the Windows Services Management Console.
- 7 To enable Catalog Search:

- a In the Navigation pane, click **Catalog Search**.
- b On the **Catalog Search Configuration** page, click **Enable**.

When catalog search indexing is enabled successfully, a message is displayed at the upper-right corner of the NetVault WebUI.

Migrating NetVault Catalog Server to different Windows-based client machine

To migrate the NetVault Catalog Search to Windows-based NetVault client:

- 1 Disable catalog search on NetVault Server, before performing catalog search uninstallation. For more information, see [Disabling catalog search](#).
- 2 Uninstall the existing catalog search from the NetVault Server. For more information, see [Removing catalog search](#).

- 3 Install and configure the catalog search on the Windows Client machine. For more information, see [Configuring catalog search](#).

i | **NOTE:** When installing catalog search on the Client machine, be sure to select **Add pre-existing and scanned backups to catalog**. Doing so includes pre-existing and scanned backup jobs for catalog searching.

When catalog search indexing is enabled successfully, a message is displayed at the upper-right corner of the NetVault WebUI.

Changing the NetVault Catalog Database directory on a Linux-based machine

To change the Catalog Database Directory on a Linux-based NetVault Server:

- 1 Disable the catalog search, see, [Disabling catalog search](#).

- 2 Run the following command to stop the **netvault-catalog** service

```
systemctl stop netvault-catalog
```

Or

Run the following command to stop the NetVault Catalog Search service, in case `systemd` is not supported on the system:

```
kill -9 <PID>
```

Read the `<PID>` value in the *catalog* section from the **catalog.cfg** file available in installation directory.

- 3 Manually move the **db** directory to the new location.

i | **NOTE:** Database is supported only on local drives of the system.

- 4 Change the file permissions and the owner and group for the new path:

```
chmod -R 750 <new DB path>/*
```

```
chown -R netvault-catalog:netvault-catalog <new DB path>/*
```

- 5 In the NetVault installation directory, open the **catalog.cfg** file to edit, available in **config** directory.

In the `[Node1]` section, change the following to specify the new location.

```
path.repo, path.repo.new, path.data, path.data.new, path.logs, path.logs.new
```

- 6 In the NetVault installation directory, open the **elasticsearch.yml** file to edit, available in the following directory **catalog/elasticsearch/config**.

Change the following to specify the new location:

```
path.repo, path.data, path.logs
```

- 7 Update the service script file in the following location: **/usr/lib/systemd/system/netvault-catalog.service** and change the database path **Environment=PID_DIR=<new location>**

Run the following command:

```
systemctl daemon-reload
```

- 8 Run the following command to start NetVault Catalog Search service:

```
systemctl start netvault-catalog
```

Or

Run the following command to start NetVault Catalog Search service, if `systemd` is not supported on the system:

```
su -s /bin/sh <CATALOG_USER_ACCOUNT> -c <netvault_installdir>/catalog/  
elasticsearch/bin/elasticsearch
```

The default value of `CATALOG_USER_ACCOUNT` is `netvault-catalog`, however if you have provided customized Linux account during installation then provide the same.

Once service is started, update the new PID in `catalog.cfg` file.

- 9 To enable Catalog Search:
 - a In the Navigation pane, click **Catalog Search**.
 - b On the **Catalog Search Configuration** page, click **Enable**.

When catalog search indexing is enabled successfully, a message is displayed at the upper-right corner of the NetVault WebUI

Modifying catalog search

To modify the catalog search configuration:

- 1 In the Navigation pane, click **Catalog Search**.
- 2 On the **Catalog Search Configuration** page, you can modify the following fields:
 - Catalog Search Password
 - Confirm Catalog Search Password
 - Reserved RAM Size
 - Maximum simultaneously active catalog workers
 - Add pre-existing and scanned backups to catalog
 - Include backups with offline indexes
 - Offline index load wait
 - Startup offline index load wait
 - Restrict the age of backups added to the catalog
 - Add backups completed on and after
- 3 Click **Configure**.

When catalog search is configured successfully, a message is displayed at the upper-right corner of the NetVault WebUI.

Disabling catalog search

After catalog search is installed and configured, you can disable the indexing.

To disable the catalog search indexing:

- 1 In the Navigation pane, click **Catalog Search**.
- 2 On the **Catalog Search Configuration** page, click **Disable**.

When catalog search indexing is disabled successfully, a message is displayed at the upper-right corner of the NetVault WebUI.

Catalog search is enabled automatically when you install and configure it.

Upgrading catalog search

To upgrade the catalog search service:

- 1 In the Navigation pane, click **Catalog Search**.
- 2 On the **Catalog Search Configuration** page, click **Upgrade**.
- 3 In the confirmation dialog, click **OK** to upgrade the catalog search package.

Removing catalog search

To remove the catalog search service:

- 1 In the Navigation pane, click **Catalog Search**.
- 2 On the **Catalog Search Configuration** page, click **Uninstall**.
- 3 In the confirmation dialog, click **OK** to remove the catalog search configuration.

Displaying Status of Catalog Search

The status gives you a quick look of catalog search server for cataloged data stored and system usage during the process. You can view the storage used, CPU usage, number of files, and savesets cataloged, and the last cataloged date and time. The status is displayed only if catalog search is installed and enabled; otherwise, no information displayed.

Table 23. Catalog Search Status

Status	Description
Storage Usage	<p>Displays the storage used for catalog search and the free disk space available, where catalog search is installed. The storage usage is displayed in percentage, and it also displays the current catalog search database store size in KB, MB, GB, or TB.</p> <p>To view the storage used by catalog search for NetVault, Click View Details.</p> <p>The graph displays monthly storage usage (KB, MB, GB, or TB). Month wise data points are plotted on this graph, these data points displays the value of used storage in KB, MB, GB, or TB once you hover the cursor over these data points.</p>
CPU Usage	<p>The total CPU usage during the catalog search process that is consumed by all the tasks. The CPU usage is calculated and displayed in percentage.</p>
Document Count	<p>Displays the total document/files count stored in the catalog search database.</p> <p>To view the document count of catalog search for NetVault, Click View Details.</p> <p>The graph displays monthly document count. Month wise data points are plotted on this graph, these data points displays the document count once you hover the cursor over these data points.</p>
Savesets Cataloged	<p>Displays the number of savesets that are cataloged against the total number of savesets available in the NetVault Database.</p> <p>For example: If there are 1000 Savesets available, but only 900 Savesets are cataloged, this field displays 900/1000</p>
Last Backup Cataloged	<p>Displays the recent catalog backup date and time. If there is no data that is been cataloged, in that case this information is not displayed.</p>

Configuring storage devices

- [About storage devices](#)
- [SAN considerations](#)
- [Quest DR Series systems](#)
- [Quest QoreStor](#)
- [NetVault SmartDisk](#)
- [EMC Data Domain Systems](#)
- [Snapshot Array Manager](#)
- [Virtual Tape Libraries](#)
- [Virtual standalone drives](#)
- [Shared Virtual Tape Libraries](#)
- [Physical tape devices](#)
- [Storage tiers](#)

About storage devices

NetVault supports a wide range of devices for storing backups. The supported devices include disk-based storage devices, deduplication appliances, physical tape libraries, autoloaders, tape drives, Virtual Tape Libraries (VTLs), and Shared Virtual Tape Libraries (SVTLs).

You can attach the storage devices to the NetVault Server, Clients, or NAS filers in a NetVault Domain. The physical storage devices can be configured for single or shared use, and connected through SCSI, iSCSI, IP, SAS, or Fiber Channel SAN interfaces. A NetVault SmartClient license is required to attach physical or virtual storage devices locally to a NetVault Client.

To use a storage device in a backup or restore operation, you must first add the device to the NetVault Domain. A device attached to a NetVault Client is only recognized after you add the client to the NetVault Server. Similarly, a device attached to a filer is only recognized after you add the filer to the server using NetVault Plug-in *for NDMP* (Plug-in *for NDMP*).

SAN considerations

The following requirements apply to a SAN environment:

- In a SAN environment, you must use persistent binding (also known as SCSI mapping, persistent reservation, or persistent naming).

NetVault cannot communicate with a library if its logical address changes as a result of changes in the SAN. Persistent binding assigns a fixed logical address to the device. This address does not change as devices are added or removed in the SAN.

For Fiber Channel Host Bus Adapters (HBAs), you can map the Fiber Channel device address (World Wide Name (WWN) or World Wide Identifier (WWID)) or Loop ID to the logical SCSI address. This configuration ensures that the changes in the SAN have no impact on the NetVault operations.

- You must also use persistent binding when the server and fiber devices are attached to separate switches or when zoning is implemented.

NetVault does not support multipathing to a tape library or device. To ensure consistent communication path, you must configure only one logical or physical channel for use.

You should not use tape libraries or drives on the same switch or in the same zone that has disk devices attached. Problems might be encountered if packets from both device types co-exist in a SAN environment. Therefore, you should use separate HBAs for these devices.

- Apple supports multipathing in FC Host Adapter and XserveRAID. Multipathing is often the default setting after installation. However, multipathing is not supported in NetVault. Therefore, the connections should be logically or physically separated to ensure a consistent communication path.

Quest DR Series systems

This section includes the following topics:

- [About Quest DR Series systems](#)
- [Quest DR Series system prerequisites](#)
- [Adding a Quest DR Series system](#)

About Quest DR Series systems

The Quest DR Series disk-based data protection appliances optimize utilization with in-line deduplication and compression, and reduce network bandwidth requirements with client-side deduplication processing and deduplicated replication.

These appliances incorporate innovative deduplication and compression technology to help you achieve a data-reduction level up to 15:1. As a result, you can retain more backup data for a longer period in the same footprint.

The Quest DR Series systems use the Plug-in *for Rapid Data Access (RDA)* to integrate its data storage operations with NetVault. This plug-in is automatically installed on the NetVault machines when you install the NetVault Server or Client software.

For more information about Quest DR Series systems, see the *Quest DR Series system Administration Guide*.

i | **NOTE:** The Quest DR Series systems do not support backups over a wide-area network (WAN).

Quest DR Series system prerequisites

Before adding a Quest DR Series system to a NetVault Server, verify that the following requirements are met:

- DR appliance must be running on DR OS version 4.0 or higher.
- **Configure the required storage options:** On the Quest DR Series system, configure the storage options for the container. For more information about the storage options, see the *Quest DR Series system Administration Guide*.
- **Configure ports for optimized replication:** To perform optimized replication operations across a firewall, configure the following fixed TCP ports on the Quest DR Series system to support RDA replication operations:
 - Port 9920

- Port 11000
- Port 10011

Adding a Quest DR Series system

To use a Quest DR Series system for backups and restores, you must first add the device to the NetVault Server. You can use the configuration wizard to add and configure this device.

- [Adding a Quest DR Series appliance](#)
- [Creating a container](#)
- [Adding a container as a media](#)

Adding a Quest DR Series appliance

To add a Quest DR Series system to the NetVault Server:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Storage Devices**.

— or —

In the Navigation pane, click **Manage Devices**, and then click **Add Device**.

- 2 On the Add Storage Devices page, select the **Quest RDA Device** option, and click **Next**.

Figure 5. Add Storage Devices page

NetVault Storage Configuration Wizard - Add Storage Devices

Select the type of device that you wish to add from the set below. If you select one of the 'virtual' device types you must also specify whether you want to create a new virtual device or whether you want to re-add one that has been created but has been removed from NetVault Backup.

- Single virtual disk device
 - Virtual tape library / media changer
 - Shared virtual tape library
 - Single physical tape device
 - Tape library / media changer
 - NetVault SmartDisk
 - Quest RDA Device
 - Data Domain Boost Device
 - Snapshot Array Manager
- Re-add previously generated virtual device

- 3 On the **Add Quest RDA Device** page, provide the following details.

Figure 6. Add Quest RDA Device page

Hostname	<input type="text"/>
Username	<input type="text"/>
Password	<input type="password"/>

Table 24. Add Quest DR Series system

Option	Description
Hostname	Type the Fully Qualified Domain Name (FQDN) or IP address of the Quest DR Series system. If the server is unable to resolve the host name, it fails to add the device.
Username	Specify a user account that can be used to log on to the device. On the Quest DR Series system, only one user account exists, and the user ID for that account is backup_user .
Password	Type the password for the user account.

- 4 Click **Add RDA Device** to add the device.

Creating a container

- 1 In the Navigation pane, click **Manage Devices**, and then click the plus icon corresponding to the DR device.
- 2 On the **Quest DR Device** page, verify that the **Containers** option is selected.
- 3 Click **Create Container**.
- 4 In the **Create Container** dialog box, use the drop-down list to select a storage group.
- 5 In the **Container Name** text box, enter a name for the storage container.
- 6 Click **Save**.

Adding a container as a media

- [Adding a container from the Manage Devices page](#)
- [Adding a container from the Quest DR Device page](#)

Adding a container from the Manage Devices page

- 1 In the Navigation pane, click **Manage Devices**, and then click the plus icon corresponding to the DR device.
- 2 On the **Quest DR Device** page, verify that the **Containers** option is selected.
- 3 Select the applicable device, and then click **Add As A Media**.
- 4 On the **Add Media** page, provide the following details:

Table 25. Add a container as a media

Option	Description
Container	The name of the selected container appears here.
Block size	The default block size is 512 KiB. The block size cannot be changed for Quest DR Series systems.

Table 25. Add a container as a media

Option	Description
Stream Limit	<p>The default value for this setting is 64 streams. This setting applies to all NetVault Servers to which the container is added. If the number of data streams exceeds the defined limit for the container, the Media Manager reports an error (“Device has too many streams”).</p> <p>You can set the soft stream limit to any value between 1 and 512.</p> <p>If the container is added to more than one NetVault Server, set the same soft stream limit on all servers.</p>
Force Add	<p>If the device is already added to another NetVault Server with the same name, select the Force Add check box. This option can be useful if you have performed a disaster recovery to rebuild the NetVault Server.</p>

5 Click **Add As A Media**.

Adding a container from the Quest DR Device page

To add a container as a media, see [Adding a container as a media](#).

- NOTE:** The Quest DR Series systems support three deduplication modes:
- **Passthrough:** When this mode is selected, deduplication processing occurs on the Quest DR Series system. The passthrough mode requires at least 200 MB of free memory on the NetVault Client.
 - **Optimized:** When this mode is selected, deduplication processing occurs on the NetVault Client. The Optimized mode requires at least 4 GB of free memory on the NetVault Client.
 - **Auto:** When this mode is selected, NetVault decides the deduplication mode (Passthrough or Optimized) based on the criteria for **Passthrough** and **Optimized**.

Quest QoreStor

Quest QoreStor is software-defined storage that provides deduplication, compression, and replication support. As a result, you can retain more backup data for a longer period in the same footprint.

- IMPORTANT:** QoreStor is only supported on Linux-based platforms.

QoreStor uses the Plug-in *for RDA* to integrate its data-storage operations with NetVault. This plug-in is automatically installed on the NetVault machines when you install the NetVault Server or Client software.

While QoreStor supports the RDAv2 and RDAv3 protocols, NetVault uses the RDAv2 protocol to communicate with QoreStor. QoreStor also uses a secured connection to NetVault, which ensures protection of your data as it moves between NetVault to QoreStor.

Before you can add and use QoreStor, you must install the applicable software. For more information, see the *Quest QoreStor Installation Guide*.

Secure Connect

Secure Connect encompasses a set of client and server components that creates a secure channel for QoreStor communication with WAN-connected clients that is also resilient to WAN outages. For more information, see the *Quest QoreStor Installation Guide*.

Secure Connect between NetVault is enabled by default upon setup; however, you can disable Secure Connect using the `sc_manager` script, if that is your preference.

For a Linux client, the location of the script is `/usr/netvault/dynlib/sc_manager`. To disable the script, set the `manager` status to `false`.

For a Windows client, the location of the script is C:\Program Files\Quest\NetVault\dynlib\sc_manager.

Adding QoreStor

After installing the QoreStor software, you can create the software-defined storage target. You can use the configuration wizard to add and configure this device.

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Storage Devices**.

—or—

In the Navigation pane, click **Manage Devices**, and then click **Add Device**.

- 2 On the **Add Storage Devices** page, select the **Quest RDA Device** option, and click **Next**.
- 3 On the **Add Quest RDA Device** page, provide the following details.

Table 26. Add Quest RDA Device

Option	Description
Hostname	Type the FQDN or IP address of the QoreStor Server.
Username	Specify a user account that can be used to log on to the QoreStor Server.
Password	Type the password for the user account.

- 4 Click **Add RDA Device**.

Adding a container as a Media for QoreStor

- [Adding a container for QoreStor from the Manage Devices page](#)
- [Adding a container from the Quest QoreStor page](#)

Adding a container for QoreStor from the Manage Devices page

- 1 In the Navigation pane, click **Manage Devices**, and then click the plus icon corresponding to the QoreStor device.
- 2 On the **Quest QoreStor** page, verify that the **Containers** option is selected.
- 3 For the applicable device, click the corresponding  icon in the **Actions** column, and select **Add As A Media**.
- 4 On the **Add Media** page, provide the following details:

Table 27. Add a container as a media for QoreStor

Option	Description
Storage Group Name	Select the name of the storage group.
Container	Select the name of the storage container.
Block size	The default block size is 512 KiB. The block size cannot be changed for Quest QoreStor systems.

Table 27. Add a container as a media for QoreStor

Option	Description
Stream Limit	<p>The default value for this setting is 64streams. This setting applies to all NetVault Servers to which the container is added. If the number of data streams exceeds the defined limit for the container, the Media Manager reports an error (“Device has too many streams”).</p> <p>You can set the soft stream limit to any value between 1 and 512.</p> <p>NOTE: You must consider target device capability before changing stream limit value for QoreStor.</p> <p>If the container is added to more than one NetVault Server, set the same soft stream limit on all servers.</p>
Force Add	<p>If the device is already added to another NetVault Server with the same name, select the Force Add check box. This option can be useful if you have performed a disaster recovery to rebuild the NetVault Server.</p>

5 Click **Add As A Media**.

Adding a container from the Quest QoreStor page

To add a container as a media, see [Adding a container as a Media for QoreStor](#).

- NOTE:** The Quest QoreStor Series systems support three deduplication modes:
- **Passthrough:** When this mode is selected, deduplication processing occurs on the Quest QoreStor system. The passthrough mode requires at least 200 MB of free memory on the NetVault Client.
 - **Optimized:** When this mode is selected, deduplication processing occurs on the NetVault Client. The Optimized mode requires at least 4 GB of free memory on the NetVault Client.
 - **Auto:** When this mode is selected, NetVault decides the deduplication mode (Passthrough or Optimized) based on the criteria for **Passthrough** and **Optimized**.

NetVault SmartDisk

This section includes the following topics:

- [About NetVault SmartDisk](#)
- [Adding a NetVault SmartDisk](#)
- [Migrating savesets and backup jobs from NetVault SmartDisk to QoreStor](#)

About NetVault SmartDisk

NetVault SmartDisk provides disk-based storage with optional data deduplication capability.

NetVault SmartDisk uses byte-level, variable block-based software deduplication, which packs up to 12 times more data into the same storage area for a 92 percent reduction in storage footprint. NetVault SmartDisk is installed and licensed separately from NetVault.

A NetVault SmartDisk instance consists of one or more Storage Pools and a set of processes that perform backup and deduplication operations. A Storage Pool consists of one or more file system volumes, which can be easily extended by adding more file system paths. An instance can accept data streams from heterogeneous platforms.

You can install a NetVault SmartDisk instance that has deduplication enabled (the default mode) on a dedicated NetVault SmartDisk Server or a NetVault Client. (If deduplication is enabled, you cannot install NetVault SmartDisk on the NetVault Server.) If deduplication is disabled, you can also select the NetVault Server as the host machine.

For more information about installing and configuring NetVault SmartDisk, see the *Quest NetVault SmartDisk Installation Guide* and *Quest NetVault SmartDisk Administrator's Guide*.

Adding a NetVault SmartDisk

To use a NetVault SmartDisk for backups and restores, you must first add the device to the NetVault Server. You can use the configuration wizard to add and configure this device.

To add a NetVault SmartDisk to the NetVault Server:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Storage Devices**.

— or —

In the Navigation pane, click **Manage Devices**, and then click **Add Device**.

- 2 On the Add Storage Devices page, select the **NetVault SmartDisk** option, and click **Next**.
- 3 On the **Add NetVault SmartDisk Instance** page, provide the following details.

Figure 7. Add NetVault SmartDisk Instance page

The screenshot shows the 'Add NetVault SmartDisk Instance' configuration page. It includes the following elements:

- Network name / IP address:** A text input field.
- Network port:** A dropdown menu with '0' selected, and a note: '(leave at zero to use the default connection port)'. There are up and down arrow icons next to the dropdown.
- Force Add:** A checkbox.
- Configure WebDAV credentials:** A checkbox.
- Username:** A text input field with placeholder text: 'Enter the WebDAV username configured on the SmartDisk instance.'
- Password:** A text input field with placeholder text: 'Enter the WebDAV password configured on the SmartDisk instance.'
- Confirm Password:** A text input field with placeholder text: 'Re-enter password'.

Table 28. Add NetVault SmartDisk

Option	Description
Network name/IP address	Type the Fully Qualified Domain Name (FQDN) or IP address of the host on which NetVault SmartDisk is installed. You must provide this information even if the device is deployed on the NetVault Server. If the server is unable to resolve the host name, it fails to add the device.
Network port	The default value for this setting is zero (0). If the device is listening on the default port, do not change this value. If the device is listening on a non-default port, type the port number configured as the Network Settings:Remote Listen Port in the percolator.cfg file. For more information about configuring a non-default port for a NetVault SmartDisk, see the <i>Quest NetVault SmartDisk Administrator's Guide</i> .

Option	Description
Force Add	If the device is already added to another NetVault Server with the same name, select the Force Add check box. This option can be useful if you have performed a disaster recovery to rebuild the NetVault Server.
Configure WebDAV credentials	<p>To prevent unauthorized access to data, you can set up WebDAV authentication for server requests on NetVault SmartDisk. NetVault SmartDisk uses Digest Access Authentication with WebDAV. For more information about setting up authentication on the NetVault SmartDisk Server, see the <i>Quest NetVault SmartDisk Administrator's Guide</i>.</p> <p>If WebDAV authentication is enabled on the NetVault SmartDisk, select this check box, and provide the following information:</p> <ul style="list-style-type: none"> • Username: Specify the user account configured on the NetVault SmartDisk Server. • Password: Type the password for the user account. • Confirm Password: Type the password again for confirmation. <p>NOTE: If you enable WebDAV authentication on the NetVault SmartDisk Server, but do not configure the authentication details on the NetVault Server, the backups and restore operations using that device fail without reporting any proper error messages. The scan operation also fails for the device.</p> <p>NOTE: NetVault does not provide any method to disable WebDAV authentication once it has been enabled for a NetVault SmartDisk. The only way to disable WebDAV authentication is to remove the NetVault SmartDisk from the NetVault Server, and add the device again.</p>

4 Click **Next** to add the device.

5 After the device is successfully added and initialized, a message is displayed.

Migrating savesets and backup jobs from NetVault SmartDisk to QoreStor

Use the migration feature to move savesets and backup jobs from NetVault SmartDisk device to QoreStor. As part of migration, all saved and scheduled backup jobs are redirected from the targeted NetVault SmartDisk device to the targeted QoreStor device. This process is completed before the data from existing backups is migrated.

Important notes

- QoreStor is only supported on Linux-based platforms.
- All savesets associated with a specific NetVault Server and NetVault SmartDisk are migrated to the new target. You cannot select specific savesets.
- Although migration moves the backup and saveset data to a QoreStor device, NetVault does not delete the data from NetVault SmartDisk. Additionally, you cannot use NetVault to delete the data from NetVault SmartDisk after migration is finished; you must manually delete the data.
- The **Expire All** function is automatically disabled during the migration process. However, it is not disabled on the new target.
- If migration fails or you manually stop the process, savesets that were already migrated automatically point to the new target. If you restart migration, only savesets that have not been migrated yet are addressed. Additionally, corresponding backup jobs are also automatically updated to point to the new target; they do not revert to the original target even if migration of the applicable savesets is interrupted.

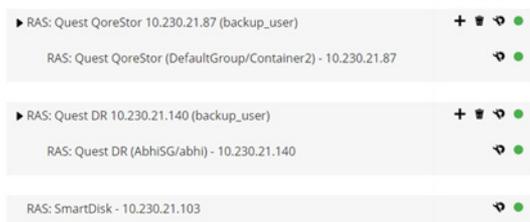
If migration of a specific saveset fails, you can run migration again after the first pass is finished. Depending on the reason for the failure, the remaining savesets might complete migration on the subsequent pass.

- If you interrupt the migration process to select a different target, you must manually update the backup jobs that have been redirected to point to the new target.
- Although NetVault does not limit the number of migrations that you can perform simultaneously, Quest recommends that you process each migration separately.
- You can migrate multiple savesets from different NetVault SmartDisk devices to the same QoreStor device if the target contains sufficient space. If space is not available, migration fails.
- You can migrate the same saveset to different QoreStor devices; however, the saveset points to the last QoreStor device that you selected.
- During the migration process, Quest strongly recommends that you avoid initiating a backup job that targets the NetVault SmartDisk device that you are migrating.
- You can create up to five storage groups and 16 containers across all storage groups.
QoreStor does not support multi-tenancy.
- NetVault does not support migration of a Secondary Copy from one type of RDA Device to another.
If you use a Duplicate method of a Phase 2 backup job to create a Secondary Copy that also targets a NetVault SmartDisk device, the job might fail after migration is finished. This occurs when both jobs targeted a NetVault SmartDisk device. During migration, the primary job is migrated to the new target, such as a QoreStor device, and the backup job completes successfully. The secondary job still identifies a NetVault SmartDisk device as the target, which causes it to fail.
- You can use the Secondary Copy feature with different RDA device types, but you cannot use it with an RDA device and a NetVault SmartDisk device. You must create a new Secondary Copy that uses an RDA device as the target.

Migrating savesets and backup jobs

- 1 Make sure that you have added the Linux-based QoreStor device to the NetVault Server.
For more information, see [Quest DR Series systems](#) or [Quest QoreStor](#).
- 2 Make sure that the NetVault SmartDisk device and the QoreStor device are online and assigned to the same NetVault Server.
 - a In the Navigation pane, click **Manage Devices**.
 - b On the **Manage Devices** page, note the status of the NetVault SmartDisk device and the QoreStor device.

Figure 8. Devices are online



- 3 Verify that the applicable backup savesets are pointing to the online NetVault SmartDisk device.
 - a In the Navigation pane, click **Explore Storage**.
 - b On the **Explore Storage** page, click **Explore Disk Storage**.
 - c On the **Explore Disk Storage** page, select the applicable device in the repository table, and click **Explore Repository**.
 - d On the **Explore Disk Storage Repository** page, review the information listed in the **Saveset table**.
 - e To close the dialog box, click **Close**.

- f Repeat these steps for each saveset that you want to migrate from NetVault SmartDisk to a QoreStor device.
- 4 When you are ready to start the migration process, click **Manage Devices** in the Navigation pane.
- 5 Click the **Manage Device** icon  for the NetVault SmartDisk device.
- 6 On the **RAS Device Management** page, click **Start Migration**.
- 7 In the **Select target device** dialog box, select the QoreStor device to which you want to migrate the NetVault SmartDisk saveset.
- 8 Select an existing Target Set or create a new Target Set that saved or scheduled backup jobs should target, and click **Confirm**.
- 9 To monitor the progress of the migration, click **Check logs**.

The **View logs** page displays all NetVault logs. To view only the NetVault SmartDisk, enter **NVSD** in the text box. The process is displayed as a percentage of the savesets that have been migrated.

EMC Data Domain Systems

This section includes the following topics:

- [About EMC Data Domain Systems](#)
- [Data Domain System prerequisites](#)
- [Adding a Data Domain System](#)
- [DD Boost commands](#)

About EMC Data Domain Systems

EMC Data Domain Systems provide disk-based storage with inline deduplication capabilities that reduce storage requirements by 10 to 30 times.

NetVault provides seamless integration with Data Domain systems through the EMC DD Boost software, allowing you to minimize your backup window and perform optimized disk-based backups while reducing your storage and network bandwidth requirements.

; | **NOTE:** The NetVault Starter Edition does not support DD Boost.

DD Boost components

The DD Boost software includes two components:

- **DD Boost Library:** This component runs on the NetVault Server, and provides the interface to communicate with the DD Boost Server running on the Data Domain system.
- **DD Boost Server:** This component runs on the Data Domain systems.

DD Boost features

DD Boost offers the following features:

- **Distributed segment processing:** DD Boost offloads parts of the deduplication process to the backup client or server, enabling the client or server to send only unique data segments to the Data Domain system.

Distributed segment processing offers the following advantages:

- It increases the aggregate backup throughput of the storage system and reduces the amount of data transferred over the network.
- It decreases processor utilization on the backup server because sending data over the network is more CPU-intensive than the distributed deduplication process.

Without this feature, the DD Boost Library sends all data (unique or redundant) to a Data Domain system for deduplication processing.

- **Advanced load balancing and link failover:** This feature lets you combine multiple Ethernet links into a group, and register a single interface with the backup application. The Data Domain system automatically balances the load for backup and restore jobs on multiple interfaces, and routes the jobs to the available interfaces if one of the interfaces in the group goes down.
- **File replication:** File-level replication enables transfer of deduplicated data directly between two or more DD Boost-enabled Data Domain systems, and thus reduces WAN bandwidth requirement by up to 99 percent. The Data Domain systems create and transfer the duplicate copies without using any resources on the backup server.

Replication requires optional DD Boost Replicator license. The license must be installed on all participating Data Domain systems.

If the source and target Data Domain systems are running different versions of the Data Domain OS, then for replication to be successful, the target system must be running the higher version of the OS.

Data Domain System prerequisites

Before adding a Data Domain system to a NetVault Server, verify that the following requirements are met:

- **Install DD Boost license on the Data Domain System:** To use a Data Domain system for backups and restores, install the required DD Boost license and enable DD Boost on the Data Domain systems.
- **Create a DD Boost user account:** On the Data Domain system, create a DD Boost user account that can be used to log on to the device for backups and restores.
- **Configure the required DD Boost features:** On the Data Domain system, configure the features that you want to use. For more information about enabling and configuring DD Boost features, see the **DD Boost** section in the *DD OS Administration Guide*.
- **Open the required firewall ports:** To perform DD Boost backups and replication across a firewall, open the following ports on the Data Domain system:
 - TCP 2049 (NFS)
 - TCP 2051 (Replication)
 - TCP 111 (NFS portmapper)
 - TCP xxx (select a random port for NFS mounted)
- **Install Microsoft Visual C++ 2005 SP1 Redistributable Package on Windows:** Install the Microsoft Visual C++ 2005 SP1 Redistributable Package on Windows-based NetVault Server. This requirement applies to all supported Windows platforms. The DD Boost library fails to load if you do not install this package on Windows.
- **Configure network time-outs:** Backup and restore jobs often take a long time to complete. Although the DD Boost Library can recover from temporary network interruptions, the operating system on the data protection application system might terminate a job prematurely if the data protection application time-outs are set too low. To avoid this, Data Domain recommends setting time-outs to at least 30 minutes (1800 seconds).

Adding a Data Domain System

To use a Data Domain system for backups and restores, you must first add the device to the NetVault Server. You can use the configuration wizard to add and configure this device.

To add a Data Domain system to the NetVault Server:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Storage Devices**.
— or —
In the Navigation pane, click **Manage Devices**, and then click **Add Device**.
- 2 On the Add Storage Devices page, select the **Data Domain Boost Device** option, and click **Next**.
- 3 On the **Add EMC DDBoost Storage** page, provide the following details.

Figure 9. Add EMC DDBoost Storage page

Network name / IP address:

Username:

Password:

LSU:

Block Size (in KiB):

Stream Limit:

Force Add

Table 29. Add Data Domain system

Option	Description
Network name/IP address	Type the Fully Qualified Domain Name (FQDN) or IP address of the Data Domain system. If the server is unable to resolve the host name, it fails to add the device.
Username	Specify a DD Boost user account that can be used to log on to the device for backups and restores. Verify that the user account is created on the Data Domain system before you add the device to the NetVault Server.
Password	Type the password for the user account.
Container	Type the name of the container. If the specified container does not exist on the Data Domain system, NetVault automatically creates it when you add the device to the server. You can configure multiple containers on a single Data Domain system. Each Data Domain system added to NetVault represents a container.
Block size	Type or select the block size for data transfers. The block size is specified in KiB. The default block size is 128 KiB.
Stream Limit	The default value for this setting is 32 streams. This setting applies to all NetVault Servers to which the container is added. If the number of data streams exceeds the defined limit for the container, the Media Manager reports an error (“Device has too many streams”). You can set the soft stream limit to any value between 1 and 256. If the container is added to more than one NetVault Server, set the same soft stream limit on all servers.
Force Add	If the device is already added to another NetVault Server with the same name, select the Force Add check box. This option can be useful if you have performed a disaster recovery to rebuild the NetVault Server.

- 4 Click **Next** to add the device.
- 5 After the device is successfully added and initialized, a message is displayed.

i **NOTE:** When you add a Data Domain system, NetVault creates several metadata files on the device. Each NetVault Server (to which you add the Data Domain system) creates its own set of metadata files. NetVault writes the data transfer statistics for Data Domain Systems to the **stats.stnz** file. The **nvstatsmng** process uses this file and requires that it is regularly updated. However, frequent updates can have a significant performance impact on the system. By default, NetVault updates the file after every 5 seconds or 10 blocks of data transfer. You can change this default setting from the **Change Settings** page. For more information, see [Configuring transfer update settings for foreign RAS devices](#).

DD Boost commands

This section provides a brief description of the DD Boost commands that you can use to manage the DD Boost features on a Data Domain system. For a detailed description of these commands, see the **DD Boost** section in the *DD OS Administration Guide*. For information about configuring DD Boost from the graphical-user interface-based Enterprise Manager, see the *DD OS Administration Guide*.

DD Boost Access

- Add clients to DD Boost access list:

```
ddboost access add clients <client-list>
```

- Delete clients from DD Boost access list:

```
ddboost access del clients <client-list>
```

- Reset DD Boost access list to factory default:

```
ddboost access reset
```

- Enable DD Boost:

```
ddboost enable
```

- Disable DD Boost:

```
ddboost disable
```

- Display DD Boost access list:

```
ddboost access show
```

- Display DD Boost status (whether enabled or disabled):

```
ddboost status
```

- Display number of active clients and connections:

```
ddboost show connections
```

This command displays the number of active clients, connections used for DD Boost, and connections used for a given group. It also provides an overview of the available interfaces.

- Delete all storage units and their contents from the Data Domain system:

```
ddboost destroy
```

This command removes all data from the storage units. The corresponding catalog entries must be removed manually.

DD Boost User

- Set DD Boost user:

```
ddboost set user-name <user-name>
```

- Display the current user:

```
ddboost show user-name
```

- Reset the DD Boost user:

```
ddboost reset user-name
```

Distributed Segment Processing

- Enable or disable Distributed Segment Processing:

```
ddboost option set distributed-segment-processing {enabled | disabled}
```
- Display status of the Distributed Segment Processing option (whether enabled or disabled):

```
ddboost option show distributed-segment-processing
```
- Reset Distributed Segment Processing to the default option (that is, enabled):

```
ddboost option reset distributed-segment processing
```

File Replication

- Enable file replication:

```
ddboost file-replication option set {encryption {enabled | disabled} | low-bw-optim {enabled | disabled}}
```

Note the following:

 - To enable file replication, this option should be set on both the source and destination Data Domain systems. Only an administrator can set this option.
 - To use encryption, the encryption option should be enabled on both the source and destination systems.
 - Low-bandwidth optimization option is only recommended for networks with less than 6 Mbps aggregate bandwidth. This option is disabled by default. For maximum filesystem write performance, leave this option disabled.
- Display status of the encryption or low-bandwidth optimization options (whether enabled or disabled):

```
ddboost file-replication option show [low-bw-optim] | [encryption]
```
- Reset the low-bandwidth optimization or encryption option for file replication:

```
ddboost file-replication option reset {low-bw-optim | encryption}
```
- Display file replication statistics:

```
ddboost file-replication show stats
```
- Reset file replication statistics:

```
ddboost file-replication reset stats
```
- Display the status of a DD Boost file replication transfer:

```
ddboost file-replication show active
```
- Display the data transfer history between the source and destination systems:

```
ddboost file-replication show history [duration duration{day | hr}] [interval hr]
```

This command displays the amount of pre- and post-compressed data, network transfer data, low-bandwidth optimization factor, and number of errors.

Interface Group (ifgroup)

- Add an interface:

```
ddboost ifgroup add interface <IP Address>
```
- Remove an interface from the group:

```
ddboost ifgroup del <IP Address>
```

Before you issue this command, verify that the interface that you want to remove is not in use by any backup or restore job.

- Enable Advanced Load Balancing and Link Failover:

```
ddboost ifgroup enable
```

- Disable Advanced Load Balancing and Link Failover:

```
ddboost ifgroup disable
```

- Remove the interfaces for Advanced Load Balancing and Link Failover and disable the ifgroup:

```
ddboost ifgroup reset
```

This command is equivalent to issuing the `ddboost ifgroup disable` command followed by multiple `ddboost ifgroup del interface ipaddr` commands.

- Display interfaces added to an ifgroup:

```
ifgroup show config
```

- Display Link Aggregation status:

```
ifgroup status
```

Storage Unit

- Create a storage unit:

```
ddboost storage-unit create <storage-unit-name>
```

- Delete a storage unit:

```
ddboost storage-unit delete <storage-unit-name>
```

The corresponding catalog entries should be removed manually.

- Display the names of all storage units or the names of all files in a specified storage unit:

```
ddboost storage-unit show [compression] [storage-unitname]
```

Use the compression option to display the original byte size, global compression, and local compression for all storage units.

Statistics

- Show the read-write statistics, including number of errors

```
ddboost show stats [interval seconds] [count count]
```

- Reset all statistics or clear all job connections when a network connection is lost

```
ddboost reset stats
```

Snapshot Array Manager

This section includes the following topic:

- [Supported Storage Array Manager](#)
- [Prerequisites](#)
- [Adding Snapshot Array Manager](#)

Supported Storage Array Manager

- Dell Enterprise Manager
- Huawei OceanStor Device Manager
- NetApp Device Manager

Prerequisites

Before you add Snapshot Array Manager, verify and consider the following requirements:

- NetVault supports snapshot-based backups using Plug-in *for FileSystem* on:
 - Dell Compellent storage center version 7.1.12.2 or later. (NetVault version 12.0 or later)
 - Huawei OceanStor Device Manager version V300R006C00 Patch Version SPC100 or later. (NetVault version 12.2 or later)
 - NetApp ONTAP Device Manager versions 8.x to 9.x
- Dell Compellent Storage arrays and Huawei OceanStor Device Manager supports the following Windows platforms: Windows Server 2008R2, 2012, 2012R2, and 2016.

Adding Snapshot Array Manager

The array manager provides a centralized management console to monitor and manage storage centers. To use the storage arrays for snapshot-based backups and restores, you must add the supported device manager to the NetVault Server running on Windows operating system.

You can use the configuration wizard to add and configure the Enterprise Manager.

To add snapshot array manager to the NetVault Server:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Storage Devices**.

— or —

In the Navigation pane, click **Manage Devices**, and then click **Add Device**.

- 2 On the **Add Storage Devices** page, select the **Snapshot Array Manager** option, and click **Next**.
- 3 On the **Add Snapshot Array Manager** page, provide the following details.

Table 30. Add Snapshot Array Manager

Option	Description
Device display name	Specify a display name for the Enterprise Manager.
Array Manager Type	Select the supported storage array manager from the drop down list.

Table 30. Add Snapshot Array Manager

Option	Description
Network name/IP address	Type the Fully Qualified Domain Name (FQDN) or IP address of the Enterprise Manager. If the server is unable to resolve the host name, it fails to add the Enterprise Manager.
Network port	Set the port number of supported devices to receive communication from all clients. Following are the default port numbers of supported device managers: <ul style="list-style-type: none">• Dell Enterprise Manager:3033• Huawei OceanStor Device Manager:8088• NetApp ONTAP Device Manager:443 If you have changed the Web Server port on the Enterprise Manager, specify the port number.
Username	Specify a user account that can be used to log on to the Enterprise Manager.
Password	Type the password for the user account.

4 Click **Next** to add the device.

5 After the Enterprise Manager is successfully added and initialized, a message is displayed.

Virtual Tape Libraries

This section includes the following topics:

- [About Virtual Tape Library](#)
- [Virtual Tape Library considerations](#)
- [Creating and adding a Virtual Tape Library](#)
- [Re-adding a previously created VTL](#)

About Virtual Tape Library

Virtual Tape Libraries (VTLs) emulate tape libraries on disk. VTLs are included in NetVault as a licensable option.

With VTLs, you have the flexibility to perform quick backups to disks, and during off-peak hours migrate or duplicate the backups to physical devices for off-site storage. The Media Manager does not distinguish between virtual and physical tapes. Therefore, you can use the same process to set up backup policies, including retirement period and rotation schemes.

VTLs are represented as directories on the disk. Each VTL contains three directories: **drives**, **slots**, and **media**. These directories contain numbered subdirectories. The virtual drives reside as files in the **drives** subdirectories. These files contain links to the media files. The virtual tapes reside as media files in the **media** directory. When a virtual tape is moved between slot and drive, the media file itself stays in the **media** directory, while the drives and slots files are modified to emulate the moving of the media.

A VTL can handle any number of concurrent NetVault Client backups. As with a physical library, the number of drives contained in the VTL dictate how many simultaneous operations can be performed. The number of slots should be the same as or more than the number of configured drives. Different operating systems may impose maximum file-size limits, which can affect the maximum VTL media size.

i | **NOTE:** VTLs are independent of file systems (for example, NTFS, UFS, ext3, and others) and disk systems (for example, IDE, SCSI, iSCSI, and others), but VTLs do not support file systems residing on removable drives.

Direct-to-cloud support

Beginning with release 12.3, NetVault supports StarWind Virtual Tape Library (StarWind VTL), which converts physical drives into virtual tapes to emulate tape hardware. It supports standard uplink protocols, including SMB3, NFSv4.1, iSCSI, and works with most virtual environments.

StarWind VTL is available on Microsoft Azure Marketplace, on AWS Marketplace, and as on-premises appliance

NetVault supports public and private cloud object storage replication with StarWind VTL through any of the following storage service providers:

- Amazon S3 and Amazon Glacier
- Amazon S3 Compatible
- BlackBlaze B2 Cloud Storage
- Microsoft Azure Cloud Storage
- Wasabi Cloud Storage

Virtual Tape Library considerations

Before creating and adding a VTL, review the following information:

- Before creating a VTL, NetVault performs a disk space check to ensure that the target disk has sufficient space to accommodate the new VTL.

On normal file systems you can use the disk space check feature to avoid errors during VTL creation. When creating a VTL on a third-party deduplication appliance or compressed file system, you should disable this feature. For more information, see [Configuring default settings for Disk Devices Plug-in](#).

- During disk space checks the free space required on the disk is calculated as follows:

`Number of Slots * Media Capacity + <x>`

Here <x> is the additional disk space considered for the following requirements:

- Disk space required to create the directory structure for VTL. The requirement varies for different file systems.
- Disk space required by other applications running on the system.

By default, the additional space requirement is set to 20MB. To change this setting, see [Configuring default settings for Disk Devices Plug-in](#).

- If the target disk does not have sufficient space to accommodate the specified VTL, the device emulation process is terminated and a message is displayed.

Creating and adding a Virtual Tape Library

You can use the configuration wizard to create a virtual tape library on the NetVault Server or a SmartClient. When creating a VTL, you can choose the number of virtual tape drives and slots, and specify the size of the virtual tapes. After the device is created, it is automatically added to the server.

To create and add a Virtual Tape Library:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Storage Devices**.

— or —

In the Navigation pane, click **Manage Devices**, and then click **Add Device**.

- 2 On the Add Storage Devices page, select the **Virtual tape library/media changer** option, and click **Next**.

- In the list of NetVault machines, select the machine on which you want to create the device, and click **Next**.
- On the **Add Virtual Tape Library** page, provide the following details.

Figure 10. Add Virtual Tape Library page

Choose the location on disk where you wish to create the new device:

Choose a name for the device:

Device Display Name:

Choose a 5-character media barcode prefix:

Specify the media capacity (in GiB):

Choose number of drives:

Choose number of slots:

Total disk space required: 512.00 GiB

Table 31. Add Virtual Tape Library

Option	Description
Choose the location on disk where you wish to create the new device	Enter the location where you want to create the VTL, or use the Browse button to select a location. When entering a location, make sure that the path is already created on the selected machine. NetVault does not automatically create any non-existing directories in the path. To create a path, click Browse , select a parent folder, and then click Add new folder . After you click OK , the new folder is added to the parent folder and can be used as the location for the new device.
Choose a name for the device	Specify a unique name for the device. NOTE: The VTL name does not support non-Latin characters.
Device display name	Specify a display name for the device.
Choose a 5 character media barcode prefix	The NetVault Server automatically generates a barcode prefix and assigns it to the VTL media. If you want to change it, type a unique code for the device.
Specify the media capacity	Type or select the size of the virtual tape. The media size is specified in GiB. Each slot contains a piece of media of the given size. The default value is 32GiB.
Choose number of drives	Type or select the number of drives for the VTL. You can create maximum of 100 drives. The default value is two drives.
Choose number of slots	Type or select the total number of slots that hold the media. You can specify a maximum of 999 slots. The default value is 16 slots.

- Click **Next** to create and add the device.

i | **NOTE:** The amount of time it takes to create a VTL depends on the media capacity and number of slots specified for the VTL. The time required is proportional to the VTL size.

- After the device is successfully added and initialized, a message is displayed.

Re-adding a previously created VTL

You can use the configuration wizard to find previously created virtual tape libraries and add the devices to the server. You can use this method to add a device that was removed from the server.

To re-add a previously created VTL:

- 1 Start the device configuration wizard.
- 2 On the Add Storage Devices page, select the **Virtual tape library/media changer** option and the **Re-add previously generated virtual device** check box.
- 3 In the list of NetVault machines, select the machine on which the device was created. Click **Next** to scan the selected client and display the discovered VTLs.
- 4 In the **Device** list, select the device that you want to add, and click **Next**.

After the device is successfully added and initialized, a message is displayed.

Virtual standalone drives

This section includes the following topics:

- [About virtual standalone drive](#)
- [Creating and adding a virtual standalone drive](#)

About virtual standalone drive

Virtual standalone drives emulate tape drives on disk. Virtual standalone drives are included in NetVault as a licensable option.

With these devices, you have the flexibility to perform quick backups to disks, and during off-peak hours migrate or duplicate the backups to physical devices for off-site storage. The Media Manager does not distinguish between virtual and physical tapes. Therefore, you can use the same process to set up backup policies, including retirement period and rotation schemes. Virtual standalone drives are represented as directories on the disk. The virtual tapes reside as media files in the directory.

Creating and adding a virtual standalone drive

You can use the configuration wizard to create a virtual standalone on the NetVault Server or a SmartClient. When creating the device, you can specify the size of the virtual tape. After the device is created, it is automatically added to the server.

To create and add a Virtual Standalone Drive:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Storage Devices**.

— or —

In the Navigation pane, click **Manage Devices**, and then click **Add Device**.

- 2 On the Add Storage Devices page, select the **Single virtual disk device** option, and click **Next**.
- 3 In the list of NetVault machines, select the machine on which you want to create the device, and click **Next**.
- 4 On the **Add Single Virtual Disk Device** page, provide the following details.

Figure 11. Add Single Virtual Disk Device page

Choose the location on disk where you wish to create the new device:

Choose a name for the device:

Device display name:

Choose a 5-character media barcode prefix:

Specify the media capacity (in GiB):

Table 32. Add virtual standalone drive

Option	Description
Choose the location on disk where you wish to create the new device	Specify the location where you want to create the device. Make sure that the path is already created on the selected machine. NetVault does not create any non-existing directories in the path.
Choose a name for the device	Specify a unique name for the device. NOTE: The single virtual disk device name does not support non-Latin characters.
Choose a 5 character media barcode prefix	The NetVault Server automatically generates a barcode prefix and assigns it to the virtual tape. If you want to change it, type a unique code for the device.
Specify the media capacity	Type or select the size of the virtual tape. The media size is specified in GiB. Make sure that sufficient space is available on the disk to create the virtual media. The default value is 32GiB.

- 5 Click **Next** to create and add the device.
- 6 After the device is successfully added and initialized, a message is displayed.

Shared Virtual Tape Libraries

This section includes the following topics:

- [About Shared Virtual Tape Libraries](#)
- [SVTL considerations](#)
- [SVTL prerequisites](#)
- [Creating and adding an SVTL](#)
- [Re-adding a previously created SVTL](#)

About Shared Virtual Tape Libraries

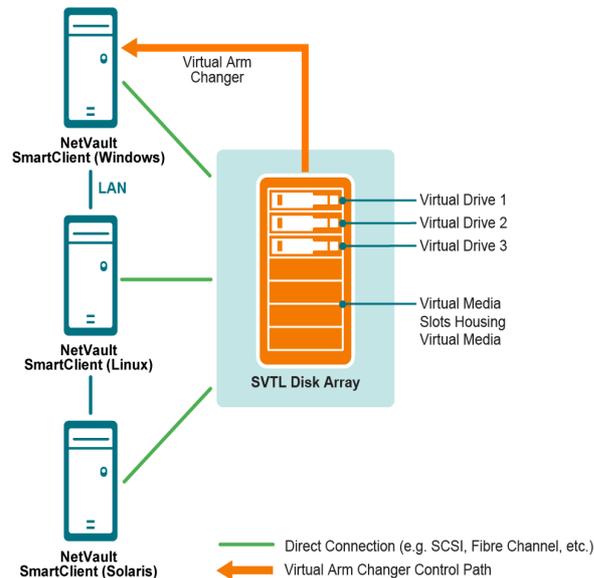
NetVault Shared Virtual Tape Libraries (SVTLs) extend the VTL implementation by allowing you to share a VTL with multiple NetVault machines for LAN-free backups.

SVTLs are supported on the following platforms:

- Windows
- Linux (x86 and x86-64)
- Solaris (SPARC and x86-64)

The interface can be Fibre Channel, iSCSI, or SCSI. On Linux and Solaris platforms, SCSI_FCP protocol is also supported. The disks can be of any size. However, operating system-imposed limitations do apply. The SVTL size can be changed during runtime using CLI utilities.

Figure 12. SVTL



SVTL considerations

Before setting up an SVTL, consider the following:

- The SVTL size depends on your disk size. Therefore, select a disk that meets your SVTL size requirements.
- Verify that the disk is physically connected to all clients that access the SVTL. The number of virtual drives for the SVTL depends on the number of machines that access the SVTL. However, it is not limited by the number of machines currently connected to the disk. You can configure additional drives for future use.
- Select the machine that controls the virtual arm changer. Although the SVTL drives can be shared or distributed among multiple clients, only one client controls the virtual arm changer.

SVTL prerequisites

Before creating an SVTL, verify that the following requirements are met:

- [General requirements](#)
- [Set up raw I/O on Linux-based systems](#)
- [Set up raw I/O on Red Hat Linux](#)
- [Set up raw I/O on Red Hat Enterprise Linux 5 or greater](#)
- [Set up raw I/O on SUSE Linux](#)

General requirements

- Connect the disk array to all NetVault Clients that will share the SVTL. The interface can be Fiber Channel, iSCSI, or SCSI. On Linux and Solaris platforms, SCSI_FCP protocol is also supported.

- Use an unformatted disk that contains no mounted partitions or volumes as an SVTL. A partition on a hard disk cannot serve as an SVTL. The additional requirements include the following:
 - The disk should allow multiple interfaces.
 - On Windows, any non-ejectable disk can serve as an SVTL.
 - NetVault does not support Multipath, Powerpath, or software RAID technologies.
- On Windows, delete the existing volumes on a disk or LUN before using it to create an SVTL.
- On Windows 2008, Windows 7, and Windows Vista, use the Disk Management administrative utility to bring the disk online. When prompted to initialize the disk, select “No.” If you do not perform these steps, NetVault cannot create an SVTL on the disk.
- On Linux, bind the pool of raw device nodes to a block device before performing any raw I/O on it. There is a raw device controller that acts as the central repository of raw to block device binding information. Binding is performed using a utility named `raw`, which is normally supplied by the Linux distributor.
- On Solaris systems hosting the SVTL or sharing the drives, specify the SCSI ID and LUN values for the applicable disks and volumes in the file `/kernel/drv/sd.conf`. This requirement is applicable only if you are using a disk or RAID volume on a SAN. Use the following format to specify the values:


```
name="sd" class="scsi" target=6 lun=5;
```
- On Solaris systems, create a single large partition named “Backup” on the hard disk. Use the applicable commands to set up the target hard disk so that it contains a single partition.
- Determine the client that controls the virtual arm changer.

Set up raw I/O on Linux-based systems

To set up raw I/O on Linux-based systems, you require the following:

- One or more free IDE or SCSI disk partitions.
- A raw device controller named `/dev/rawctl` or `/dev/raw`. If the controller is not present, type the following command to create a symbolic link:

```
ln -s /dev/your_raw_dev_ctrl /dev/rawctl
```

The following example shows how to set up raw I/O on Linux:

- 1 At the prompt, type the following command to display information from the file `devices.txt`. You can find this file in the `/usr/src/linux/Documentation` directory:

```
ls /dev/rawctl
```

— or —

```
ls /dev/raw/raw1
```

- 2 Logged in as root, type the following command to create the device:

```
mknod /dev/rawctl c 162 0
```

- 3 Set the following permissions:

```
crw-rw
```

If you require `/dev/raw/raw1` and `/dev/raw/raw2`, follow the same procedure using the proper numbers listed in the `devices.txt` file and set the same permissions.

Set up raw I/O on Red Hat Linux

The following example shows how to set up raw I/O on Red Hat Linux. The raw partition used is `/dev/sda`.

- 1 Calculate the number of 4096-byte pages in this partition, as shown in the following example:

```
fdisk /dev/sda
```

```
Disk /dev/sda: 255 heads, 63 sectors, 1106 cylinders
Units = cylinders of 16065 * 512 bytes
num_pages = floor( ((1106-524+1)*16065*512)/4096 )
num_pages = 11170736
```

- 2 Bind an unused raw device node to this partition. Binding is required each time the machine is rebooted. You must be logged in as root to run this command:

```
raw /dev/raw/raw1 /dev/sda
```

- 3 For persistent binding, open the `/etc/sysconfig/rawdevices` file and append the following line:

```
dev/raw/raw1 /dev/sda
```

Restart the system or type the following command:

```
/etc/rc.d/init.d/rawdevices start
```

- 4 Set appropriate read permissions on the raw device controller and the disk partition. Set appropriate read and write permissions on the raw device.

Set up raw I/O on Red Hat Enterprise Linux 5 or greater

The raw devices interface has been deprecated in Red Hat Enterprise Linux 5; the raw device mapping is now performed using **udev** rules. To correctly map the raw device, add the appropriate entries to the `/etc/udev/rules.d/60-raw.rules` file in the following format:

- For device names:

```
ACTION=="add", KERNEL=="<device name>", RUN+="/bin/raw /dev/raw/rawX %N"
```

- For major or minor numbers:

```
ACTION=="add", ENV{MAJOR}=="A", ENV{MINOR}=="B", RUN+="/bin/raw /dev/raw/rawX %M %m"
```

Here `<device name>` is the name of the device that you want to bind (for example, `/dev/sda1`), A and B are the major or minor numbers of the device you want to bind, and X is the raw device number that you want the system to use.

If you have a large pre-existing `/etc/sysconfig/rawdevices` file, convert it using the following script.

```
#!/bin/sh
grep -v "^ *#" /etc/sysconfig/rawdevices | grep -v "^$" |
while read dev major minor;
do
if [-z "$minor"]; then
echo "ACTION==\"add\", KERNEL==\"${major##/dev/}\"",
RUN+=\"/bin/raw $dev%N\""
else
echo "ACTION==\"add\", ENV{MAJOR}==\"$major\",
ENV{MINOR}==\"$minor\", RUN+=\"/
bin/raw $dev%M%m\"$dev%M%m\""
fi
done
```

Set up raw I/O on SUSE Linux

On SUSE Linux, administer the raw disk partitions in the `/etc/raw` file. This plain text file contains comments and examples for possible configurations. After creating the raw devices, bind the raw devices by starting them with the script `/etc/init.d/raw`. Use the `chkconfig(8)` utility to ensure that the raw device binding occurs after any restart.

Creating and adding an SVTL

You can use the configuration wizard to create and add an SVTL. When creating an SVTL, you can choose the number of virtual tape drives and slots, and specify the media capacity. After the device is created, you can assign the drives to different clients or share the drives with multiple clients.

To create an SVTL:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Storage Devices**.

— or —

In the Navigation pane, click **Manage Devices**, and then click **Add Device**.

- 2 On the Add Storage Devices page, select the **Shared virtual tape library** option, and click **Next**.
- 3 In the list of NetVault machines, select the machine on which you want to create the device, and click **Next**.
- 4 On the **Add Shared virtual tape library** page, provide the following details.

Option	Description
Library Device	Select the target drive. Depending on the operating system, the list includes the following items: <ul style="list-style-type: none">• Windows: PhysicalDrive1, PhysicalDrive2, and so on• Linux: /dev/raw/raw1, /dev/raw/raw2, and so on• Solaris: /dev/rdisk/c0t0d0s0, /dev/rdisk/c1t1d0s0, and so on
Device Description	Displays the disk type.
Device Size	Displays the disk size.
Device Block Size	Displays the block size.
Previously Formatted as SVTL?	Indicates whether the selected disk was previously formatted as an SVTL.
Barcode Prefix	The NetVault Server automatically generates a barcode prefix and assigns it to the media used by the device. If you want to change it, type a unique code for the device.
Number of Drives	Type or select the number of drives for the SVTL. The number of drives can be more than the number of NetVault Clients currently connected to the disk. The additional drives can be used in future to connect more clients.
Number of Media Items	Type or select the total number of slots that hold the media.
Media Capacity	Type or select the media size. The media size is specified in MiB. When creating an SVTL, NetVault uses a few megabytes of space to store some information about the SVTL on the disk. Take this requirement into consideration when you configure the media capacity.

- 5 To determine the disk space requirements for the SVTL, click **Calculate Size Required**. If the required disk size is larger than the actual disk size, reduce the **Media Items** and **Media Capacity**.
- 6 Click **OK**.

- 7 To create an SVTL, NetVault formats the specified disk. To complete this step, provide the following details:
 - **Password:** Type the NetVault password for the server.
 - **Confirmation Phrase:** Type the text **FORMAT SVTL**.

Click **Format** to format the disk and create the SVTL.

- 8 After the SVTL is created, and the tape drives are automatically discovered and assigned to the appropriate storage bays, a message is displayed.

You can use this default configuration if all drives are to be controlled by the client selected in [Step 3](#). In this case, click **Add Library**. Information message appears. Close the message to exit the configuration wizard.

- 9 To assign the drives to different clients or to share the drive with multiple clients, click **Add Drives Manually**, and complete the following steps:
 - a In the **Choose Machine** table, select the client to which the drive is attached. If the device is connected to multiple clients, select any one client. Click **Next** to scan the selected client and list the attached devices.

i | **NOTE:** To skip any bay and configure the next bay, click **Leave bay empty**.

- b In the **Choose drive for bay** table, select the device that you want to add, and click **Next**.
- c If the device is connected to multiple clients (for example, in a SAN setup), all host clients are listed in the **Choose Machines** table. To share the drive with multiple clients, select the additional clients in the **Choose Machines** table, and click **Next**.
- d After the drive is successfully assigned to the selected clients, a message is displayed.

To assign additional drives for the library, click **Add more devices**, and repeat [Step a](#) through [Step c](#).

If you do not want to add any other drives, exit the configuration wizard.

Re-adding a previously created SVTL

You can use the configuration wizard to find a previously created SVTL and add the device to the server. You can use this method to add a device that was removed from the server.

To re-add a previously created SVTL:

- 1 Start the device configuration wizard.
- 2 On the Add Storage Devices page, select the **Shared virtual tape library** option and the **Re-add previously generated virtual device** check box.
- 3 In the list of NetVault machines, select the machine on which the device was created. Click **Next** to scan the selected client and list the discovered SVTLs.
- 4 In the **Device** list, select the device that you want to add, and click **Next**.
- 5 After the SVTL is added, and the tape drives are automatically discovered and assigned to the appropriate storage bays, a message is displayed.

You can use this default configuration if all drives are to be controlled by the client selected in [Step 3](#). In this case, click **Add Library**. Information message appears. Close the message to exit the configuration wizard.

- 6 To assign the drives to different clients or to share the drive with multiple clients, click **Add Drives Manually**, and complete the following steps:
 - a In the **Choose Machine** table, select the client to which the drive is attached. If the device is connected to multiple clients, select any one client. Click **Next** to scan the selected client and list the attached devices.

i | **NOTE:** To skip any bay and configure the next bay, click **Leave bay empty**.

- b In the **Choose drive for bay** table, select the device that you want to add, and click **Next**.
- c If the device is connected to multiple clients (for example, in a SAN setup), all host clients are listed in the **Choose Machines** table. To share the drive with multiple clients, select the additional clients in the **Choose Machines** table, and click **Next**.
- d After the drive is successfully assigned to the selected clients, a message is displayed.
To assign additional drives for the library, click **Add more devices**, and repeat [Step a](#) through [Step c](#).
If you do not want to add any other drives, exit the configuration wizard.

Physical tape devices

This section includes the following topics:

- [Adding a standalone tape drive](#)
- [Adding a tape library](#)

Prerequisite

Before you add a tape device, confirm whether your operating system (OS) supports the NetVault default settings.

Changing the default Drive Transfer Buffer Size option

In NetVault, the Drive Transfer Buffer Size option is set to 8 MiB by default. To ensure that this configuration is supported on your system. If this value is too high, change it to 257 KiB in the template file by , completing the following steps.

To change the default setting for the Drive Transfer Buffer Size option:

- 1 Open the file “**drives.tpl**” in a text editor.
You can find this file in <NetVault home>\devices\drives on Windows and <NetVault home>/devices/drives on Linux and UNIX.
- 2 In the [**Device:Total Buffer Size**] section, set `Value` to 257.
For example:
`Value = 257`
- 3 Save the changes and close the file.

i | **NOTE:** If **Drive Transfer Buffer Size** is set too high and NetVault fails to allocate the transfer buffers, the backup jobs using tape devices show reduced performance.

Adding a standalone tape drive

You can use the configuration wizard to add and configure a standalone tape drive. The wizard lets you configure the device for shared access if it is connected to multiple hosts.

i | **NOTE:** On Windows-based machines, you must disable the Removable Storage Service before adding a device to the NetVault Server. For more information about disabling this service, see the *Quest NetVault Installation Guide*.

To add a standalone tape drive to the NetVault Server:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Storage Devices**.
— or —
In the Navigation pane, click **Manage Devices**, and then click **Add Device**.
- 2 On the Add Storage Devices page, select the **Single physical tape device** option, and click **Next**.
- 3 In the list of NetVault Clients, select the client that is physically connected to the target device. If the device is connected to multiple clients, select any one client.
Click **Next** to scan the selected client and list the attached devices.
- 4 In the **Choose drive** table, select the device that you want to add.
In the **Device display name** box, provide a display name for the tape library.
Click **Next**.
- 5 If the device is connected to multiple clients (for example, in a SAN setup), all host clients are listed in the **Choose Machines** table. To share the drive with multiple clients, select the clients in the **Choose Machines** table, and click **Next**.
This page is not displayed if the drive is connected to a single client.
- 6 After the device is successfully added and initialized, a message is displayed on the page.

Adding a tape library

You can use the configuration wizard to add and configure a physical tape library. The wizard lets you configure the device for shared access if it is connected to multiple hosts.

i | **NOTE:** On Windows-based machines, you must disable the Removable Storage Service before adding a device to the NetVault Server. For more information about disabling this service, see the *Quest NetVault Installation Guide*.

To add a tape library to the NetVault Server:

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Add Storage Devices**.
— or —
In the Navigation pane, click **Manage Devices**, and then click **Add Device**.
- 2 On the Add Storage Devices page, select the **Tape library/media changer** option, and click **Next**.
- 3 In the list of NetVault Clients, select the client that is physically connected to the target device. If the device is connected to multiple clients, select the client that you want to designate as the library controller.
Click **Next** to scan the selected client and list the attached devices.

i | **NOTE:** In NetVault, a library changer is always controlled by a single machine. The drives can be shared among multiple clients.
- 4 In the **Choose library** table, select the device that you want to add.
In the **Tape Library Display Name** box, specify a display name for the tape library.
Click **Next**.
- 5 After the tape drives are discovered and assigned to the appropriate storage bay, a message is displayed. To use this default configuration, no further action is required. You can exit the configuration wizard.
- 6 If the library is not configured automatically, click **Add drives manually**. You can also use this method to assign the drives to different clients or to share the drive with multiple clients.

To manually select the controlling client and drive for each drive bay in the library, complete the following steps:

- a In the **Choose Machine** table, select the client to which the drive is attached. If the device is connected to multiple clients, select any one client.

Select the applicable option:

- **Scan for devices:** To scan the client again for attached devices, select this option.
- **Do not scan for devices:** To use the previous scan results, select this option.
- **Scan only filer:** To scan for devices attached to a particular filer, select this option, and then select the applicable filer.

This option is only displayed if you have installed the Plug-in *for NDMP* on the server and added a filer to the plug-in.

Click **Next**.

i | **NOTE:** To skip any bay and configure the next bay, click **Leave bay empty**.
To leave the remaining bays empty and complete the library addition with the required drives, click **Leave remaining bays empty**.

- b In the **Choose drive for bay** table, select the device that you want to add, and click **Next**.

If a drive is available as both a SAN- and an NDMP-attached device, adding the drive as one type removes the other type from the list of available drives.

- c If the device is connected to multiple clients (for example, in a SAN setup), all host clients are listed in the **Choose Machines** table. To share the drive with multiple clients, select the clients in the **Choose Machines** table, and click **Next**.

This page is not displayed if the drive is connected to a single client.

- d After the drive is successfully assigned to the selected clients, a message is displayed.

To assign additional drives for the library, click **Add more devices**, and repeat [Step a](#) through [Step c](#) until all required drives are configured.

If you do not want to add any other drives, exit the configuration wizard.

i | **NOTE:** Inventory scanning during library initialization is faster for tapes with barcodes than tapes without barcodes. Therefore, using barcodes is recommended for all types of data and cleaning tapes.

Storage tiers

A storage tier is a logical grouping of storage media that can be used as a backup Target Set. You can base the grouping on similarities such as media type, performance, location, or capacity, or on other business requirements.

The storage tier feature supports the following storage formats:

- RAS (including Quest DR Series, EMC Data Domain Systems, NetVault SmartDisk, and Quest QoreStor)
- Virtual Tape Library (VTL)
- Physical Tape Library

For more information on creating and managing storage tiers, see the following topics:

- [Creating a storage tier](#)
- [Editing a storage tier](#)
- [Removing a storage tier](#)

After you create a storage tier, that tier becomes available to select as a Target Set when you create a backup job.

Creating a storage tier

To create a storage tier:

- 1 In the Navigation pane, click **Manage Devices**, and then click the **Storage Tiers** tab.
- 2 Click **Create Storage Tier**.
- 3 On the Create Storage Tier details drawer, enter the information described in the following table.

Detail	Description
Name	Required. Enter a name for the storage tier.
Description	Optional. Enter text that describes the storage tier, such as the types of servers that back up to it.
Selected Storage	Required. To select the storage that you want to include in the tier, do one of the following actions: <ul style="list-style-type: none">• Drag and drop one or more storage from the Available Storage list to the Selected Storage area.• Use the + button to locate a single storage device that you want to include in the tier.
Available Storage	Use the search bar to locate and add available storage to the list.

- 4 Click **Apply**.

The storage tier is created and appears on the Manage Devices and Storage Tiers page with the display name and storage count, which refers to the number of storage within the storage tier. Expanding the tier reveals the storage devices attached to it and a description. For more details, hover your cursor over the information icon.

The storage tier is available to be selected as a target storage set.

Editing a storage tier

To edit a storage tier:

- 1 In the Navigation pane, click **Manage Devices**, and then click the **Storage Tiers** tab.
- 2 Next to the relevant storage tier, click the **Edit** icon.
- 3 On the **Edit Storage Tier** details drawer, make changes to any of the details described in the following table.

Detail	Description
Name	Lets you change the name of the storage tier. NOTE: A name for the storage tier is mandatory.
Description	Lets you revise the description of the storage tier.
Selected Storage	To find a specific selected storage, use the Search bar. To remove storage, either use the - button or drag and drop the storage out of the Selected Storage area.
Available Storage	To add storage, either use the + button or drag and drop the storage into the Selected Storage area..

- 4 Click **Apply**.

The storage tier appears on the Manage Devices and Storage Tiers page with the display name and storage count, which refers to the number of storage within the storage tier. Expanding the tier reveals the

storage devices attached to it and a description. For more details, hover your cursor over the information icon.

Removing a storage tier

i | **NOTE:** You cannot delete a storage tier that is part of a target set.

To remove a storage tier:

- 1 In the Navigation pane, click **Manage Devices**, and then click the **Storage Tiers** tab.
- 2 Next to the relevant storage tier, click the **Delete** icon.
- 3 In the Remove Storage Tier dialog box, click **Yes** to confirm.

Backing up data

- [About backing up data](#)
- [Secondary Copy](#)
- [Snapshot-based backups](#)
- [Backup retirement](#)
- [About NetVault Sets](#)
- [Backup and recovery strategy](#)
- [Creating backup jobs](#)
- [Creating Schedule Sets](#)
- [Creating Target Sets](#)
- [Creating Source Sets for backup jobs](#)
- [Creating Backup Advanced Options Sets](#)
- [Managing sets](#)

About backing up data

A backup is a copy of data, which can be used to restore and recover the original data after a data loss event.

NetVault offers a selection of plug-ins, which integrate with the native application programming interfaces (APIs) to provide application-consistent backups and recovery of data. Depending on the application type, these plug-ins provide multiple methods and options to back up the selected data.

In general, NetVault supports the following features:

- Full and selective backups
- Primary and secondary backups
- Normal and deduplicated backups
- Encrypted and non-encrypted backups
- Repeating and non-repeating backups

Backup job definition

To back up your data, you must create and submit a backup job. The NetVault WebUI provides a configuration wizard that helps you to perform this task. You can run the wizard from the **Guided Configuration** or **Create Backup Job** link in the Navigation pane.

A backup job definition includes the following components:

- Selection list
- Plug-in options

- Scheduling options
- Source device options (available only to the Plug-in *for Consolidation*, Plug-in *for Data Copy*, and Secondary Copy jobs)
- Target device and media options
- Advanced backup options

These components are stored in NetVault Sets. For more information about NetVault Sets, see [About NetVault Sets](#).

Each backup job has a Job ID number and a Job Name. The Job ID number is an auto-generated number. The Job Name is a user-defined string, which allows you to easily identify the job when monitoring its progress, viewing the job logs, or selecting a backup to restore data. A backup is stored as a **Saveset** on the media.

Backup indexes

NetVault generates a backup index for each backup, and writes this index to the backup media and NetVault Database.

The backup index includes a header which contains information required for restoring data. There is no maximum limit on the index file size or the number of items that you can include in a backup job.

Backup indexes stored in the NetVault Database are called **Online Indexes**. Online indexes allow you to quickly scan through the contents of a saveset without loading the media.

Secondary Copy

With a backup job, you can choose to run a Phase 2 job to create a Secondary Copy, which can be used for off-site storage and disaster recovery purposes. NetVault provides two methods to create a Secondary Copy: Duplicate and Data Copy.

Duplicate

The Duplicate method creates an exact copy which is linked to the original backup. This method breaks down the backup into segments and copies the segments to the storage device. During restore, the segments from the primary backup and secondary copy are interchangeable. As it is not possible to mix unencrypted segments with encrypted segments during restore, you cannot enable or disable encryption for the Duplicate. If the original saveset is encrypted, the Duplicate method creates an encrypted copy. If the original saveset is not encrypted, this method creates an unencrypted copy.

Data Copy

The Data Copy method breaks down the backup into segments and copies the segments to the backup device. During restore, either the primary backup or the secondary copy is used to recover data; the segments from the primary backup and secondary copy are not interchangeable. Therefore, it is possible to enable encryption for the Data Copy when the primary copy is unencrypted. This option is useful when you want to use the deduplication option for primary backups.

The Data Copy job has two requests: one for read and another for write. When you set a priority to the Data Copy job, the first request (source request) receives the priority that the user assigns to it; the subsequent request (destination request) served as first priority in the background with priority 0 or 1. This implementation is to avoid waiting for the destination media request to be served if the source media is available.

; | **NOTE:** If the primary copy is encrypted, the Data Copy method automatically creates an encrypted saveset.

Snapshot-based backups

On supported Windows platforms, NetVault plug-ins can use a hardware or software VSS provider to create consistent point-in-time copies of volumes, and use these snapshots to perform backups. The plug-ins can also create persistent snapshots, which can be used for data recovery operations.

Currently, the NetVault Plug-in for *FileSystem* supports snapshot-based backups.

There are multiple ways in which you can configure the plug-in to perform VSS-based backups. You can configure the plug-in to:

- Create and use snapshots to perform backups to disk- or tape-based storage devices.
- Create and use snapshots to perform backups to disk- or tape-based storage devices, and retain snapshots as backups on supported disk arrays.

To create and use persistent snapshots, the data that you want to back up must reside on supported disk arrays. For more information about OS versions, plug-in versions, and disk arrays that support persistent snapshots, see the *Quest NetVaultCompatibility Guide*.

Backup retirement

A backup can be retained indefinitely or retired after a specified period. NetVault supports generation- and time-based retirement methods for backups. When a backup is retired, its index is deleted from the NetVault Database.

Backup retirement methods

You can specify generation- and time-based retirement methods for backups:

- **Generation-based backup retirement:** This method specifies the maximum generation count for a Full Backup (that is, the maximum number of Full Backups that are retained for the same data set). Generation-based retirement can only be set for Full Backups. When the number of Full Backups exceeds the specified generation count, the oldest Full Backup is retired.
- **Time-based backup retirement:** This method specifies the length of time a backup is retained. You can specify the retirement age for backups in number of days, weeks, or years. A backup is retired when the retirement age is reached. Time-based retirement can be set for all backup types (that is, Full, Incremental, and Differential).

i | **NOTE:** When a backup stored on a disk-based storage device (such as NetVault SmartDisk, Quest DR Series system, or Data Domain system) is retired, that backup is deleted from the device. You cannot import the deleted backup by scanning the device.

Retirement rules for dependent backups

If a backup has any dependent backups, you can use the following rules to delete the backup series:

- **Retire all backups when the last backup is retired:** By default, the retirement of a backup series is deferred until all backups from this series are ready for retirement. The backups from an Incremental series or a Differential series are retired when the last dependent backup from the series is retired.
- **Retire all backups when the first backup is retired:** NetVault lets you modify the retirement behavior to retire a backup series when the first backup is retired. When this rule is applied, the backups from an Incremental series or a Differential Series are retired when the first backup from the series is retired.

You can choose to apply this rule in the following ways:

- Use the **Retirement Timing Control** settings to change the default behavior for all backups
- Use the **saveset** retirement settings to override the default behavior for selected backups

The following examples illustrate the retirement behavior for different backup series:

- [Example 1: Incremental Backup series](#)
- [Example 2: Differential Backup series](#)
- [Example 3: Mixed backups \(Full, Incremental, and Differential Backups\)](#)

Example 1: Incremental Backup series

In this example, Full Backups are performed on Sunday and Incremental Backups are performed from Monday through Saturday. The retention period is set to seven days.

Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Full Backup	Incremental Backup 1	Incremental Backup 2	Incremental Backup 3	Incremental Backup 4	Incremental Backup 5	Incremental Backup 6

The Incremental Backups depend on the most recent backup of any type (Full, Incremental, or Differential). Therefore, these backups create a single backup series:

Full Backup <- Incremental Backup 1 <- Incremental Backup 2 <- Incremental Backup 3 <- Incremental Backup 4 <- Incremental Backup 5 <- Incremental Backup 6

Depending on the retirement rule that is applied, the backups are retired as follows:

- **Retire all backups when the last backup is retired:** When you apply this rule, all backups from this series are retired on Day 15.
- **Retire all backups when the first backup is retired:** When you apply this rule, all backups from this series are retired on Day 8.

Example 2: Differential Backup series

In this example, Full Backups are performed on Sunday and Differential Backups are performed from Monday through Saturday. The retention period is set to seven days.

Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Full Backup	Differential Backup					

Differential Backups depend on the recent Full Backup. Therefore, these backups create six separate backup series:

Full Backup <- Differential Backup 1

Full Backup <- Differential Backup 2

Full Backup <- Differential Backup 3

Full Backup <- Differential Backup 4

Full Backup <- Differential Backup 5

Full Backup <- Differential Backup 6

Depending on the retirement rule that is applied, the backups are retired as follows:

- **Retire all backups when the last backup is retired:** When you apply this rule, the backups from this example are retired as follows:

Day 8	Day 9	Day 10	Day 11	Day 12	Day 13	Day 14	Day 15
None	None	Differential Backup 1	Differential Backup 2	Differential Backup 3	Differential Backup 4	Differential Backup 5	Full Backup and Differential Backup 6

- **Retire all backups when the first backup is retired:** When you apply this rule, all backups from this example are retired on Day 8.

Example 3: Mixed backups (Full, Incremental, and Differential Backups)

In this example, Full Backups are performed on Sunday, Incremental Backups are performed on Monday, Tuesday, and Thursday, and Differential Backups are performed on Wednesday and Saturday. The retention period is set to seven days.

Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Full Backup	Incremental Backup 1	Incremental Backup 2	Differential Backup 1	Incremental Backup 3	Incremental Backup 4	Differential Backup 2

The Differential Backups depend on the recent Full Backup, while the Incremental Backups depend on the most recent backup of any type (Full, Differential, or Incremental). Therefore, these backups create three separate backup series:

Full Backup <- Incremental Backup 1 <- Incremental Backup 2

Full Backup <- Differential Backup 1 <- Incremental Backup 3 <- Incremental Backup 4

Full Backup <- Differential Backup 2

Depending on the retirement rule that is applied, these backups are retired as follows:

- **Retire all backups when the last backup is retired:** When you apply this rule, the backups from this example are retired as follows:

Day 8	Day 9	Day 10	Day 11	Day 12	Day 13	Day 14	Day 15
None	None	None	Incremental Backup 1 and Incremental Backup 2	None	None	Differential Backup 1, Incremental Backup 3, and Incremental Backup 4	Full Backup and Differential Backup 2

- **Retire all backups when the first backup is retired:** When you apply this rule, all backups from this example are retired on Day 8.

About NetVault Sets

NetVault Sets are used to create backup and restore jobs.

You can use sets to store data selections, backup and restore options, scheduling options, device and media options, and other advanced backup and restore options. Sets eliminate the need to manually select the data items or configure backup and restore options for each job, and allow you to quickly and easily apply the same data selections and options to multiple jobs.

For example, you can save the data selections in a Backup Selection Set, and use this set to create Full, Incremental, and Differential Backup jobs to ensure that the same data set is used for all jobs. Similarly, you can change the day, date, or time in a Schedule Set to automatically change the job schedule for multiple jobs, or specify a new device for backups by changing the Target Set.

Set types

The following table describes the set types that are available in NetVault.

Table 33. Set types

Set type	Descriptions
Backup Selection Set	<p>This set is used to specify data selections for backup jobs. The data items that are available for selection depend on the plug-in in use.</p> <p>For more information about these sets, see the relevant plug-in user's guide.</p>
Plugin Options Set	<p>This set is used to specify the backup method, backup type, and other backup options.</p> <p>The backup options that are available to a job depend on the plug-in in use. For more information about these sets, see the relevant plug-in user's guide.</p>
Schedule Set	<p>This set is used to specify scheduling options for backup and restore jobs. These options define when and at what intervals a job runs.</p> <p>NetVault includes the following predefined Schedule Sets:</p> <ul style="list-style-type: none"> • Immediate • Daily 10 PM • Friday 10 PM • Week Night 10 PM
Source Set	<p>This set is used to specify source device options.</p> <p>A Source Set is required for the following jobs:</p> <ul style="list-style-type: none"> • Plug-in <i>for Consolidation</i> jobs • Plug-in <i>for Data Copy</i> jobs • Secondary Copy jobs • Restore jobs <p>NetVault includes the following predefined Source Set:</p> <ul style="list-style-type: none"> • Any Device
Target Set	<p>This set is used to specify target device and media options for backup jobs.</p> <p>NetVault includes the following predefined Target Sets:</p> <ul style="list-style-type: none"> • Default Backup Target Options • Local Only • Reuse Old Media • Stand-alone
Backup Advanced Options Set	<p>This set is used to specify backup retirement settings, Secondary Copy job definitions, user-defined events, and other advanced options.</p> <p>NetVault includes the following predefined Backup Advanced Options Sets:</p> <ul style="list-style-type: none"> • Default Advanced Backup Options • Keep forever (Archive) • Encrypt
Restore Selection Set	<p>This set is used to specify data selections for restore jobs.</p> <p>The data items that are available for selection depend on the plug-in in use. For more information, see the relevant plug-in user's guide.</p>
Restore Advanced Options Set	<p>This set is used to specify restore type, user-defined events, and other advanced restore options.</p> <p>NetVault provides the following predefined Restore Advanced Options Set:</p> <ul style="list-style-type: none"> • Restore from selected backup

Backup and recovery strategy

The primary objective of backing up data is to recover from the damages caused by a data loss event and resume normal operations quickly. This objective requires a good backup strategy, which maximizes data availability and minimizes data loss and downtime, while balancing your business requirements with costs, resources, and other factors.

To create a good backup plan, consider the possible failure modes, like hardware failure, data corruption, human error, or loss of a data center, and select the suitable backup methods and features to recover from these scenarios.

Typically, your backup plan should define what backup methods are used, when and at what intervals the backups are performed, how the backups are stored, how long the backups are retained, and how the backup media are re-used.

Creating backup jobs

You can use the backup job wizard to create and submit backup jobs. The wizard can be accessed from the **Guided Configuration** or **Create Backup Job** link in the Navigation pane.

To create a backup job:

- 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**.

Figure 13. Create Backup Job page

The screenshot shows the 'Create Backup Job' page in the NetVault Configuration Wizard. The page has a header with the Quest and NetVault logos and a user profile 'admin'. The main content area is titled 'Create Backup Job' and contains a form with the following fields:

- Job Name:
- Selections: [+] [-] [trash]
- Plugin Options: [+] [-] [trash]
- Schedule: [+] [-] [trash]
- Source Storage: [+] [-] [trash]
- Target Storage: [+] [-] [trash]
- Advanced Options: [+] [-] [trash]

At the bottom right, there are three buttons: 'Back', 'Save', and 'Save & Submit'.

- 2 In **Job Name**, type a name for the job. Assign a descriptive name that allows you to easily identify the job for monitoring its progress or restoring data.

The job name can contain alphanumeric and non-alphanumeric 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 On the Create Backup Job page, the actions described in the following table are available for the details listed below Job Name.

Icon	Description
	Hover to view information about the corresponding detail.
	To create a new set for the corresponding detail, click this icon. The corresponding drawer opens. Enter the required information, and then click Save . Optionally, click Clone Existing Set , select a set and click Load , and then enter a name for the cloned set at the bottom of the page.
	To edit the information for the selected set, click this icon. The corresponding drawer opens. Update the required information, and then click Save . Optionally, click Clone Existing Set , select a set and click Load , and then enter a name for the cloned set at the bottom of the page.
	Click to delete the selected set. Click OK to confirm.

- 4 The details described in the following table are available for viewing, creating, editing, and deleting.

Detail	Description
Selections	Select an existing Backup Selection Set, or create a new set and select the items that you want to back up. NOTE: The data items that are available for selection depend on the plug-in in use. For more information about selecting data for backups, see the relevant plug-in user's guide.
Plugin Options	Select an existing Backup Options Set, or create a new set and configure the options that you want to use. NOTE: The backup options that are available to a job depend on the plug-in in use. For more information about these options, see the relevant plug-in user's guide.
Schedule	Select an existing Schedule Set, or create a new set and configure the schedule type and method. For more information, see Creating Schedule Sets . NOTE: The predefined set " Immediate " is selected by default. To run the job as soon as it is submitted, use this set.
Source Storage	This option appears only when creating a Plug-in <i>for Consolidation</i> and Plug-in <i>for Data Copy</i> job. Select an existing Source Set, or create a new set and configure the source device options for the job. For more information, see Creating Source Sets for backup jobs . NOTE: The predefined set " Any Device " is selected by default.
Target Storage	Select an existing target set, or create a new set and configure the target device and media options for the job. For more information, see Creating Target Sets . NOTE: The predefined set " Default Backup Target Options " is selected by default.
Advanced Options	Select an existing Backup Advanced Options set, or create a new set and configure the options that you want to use. For more information, see Creating Backup Advanced Options Sets . NOTE: The predefined set " Default Advanced Backup Options " is selected by default.

- 5 Select one of the following methods to save or schedule the job:
 - To save the definition without scheduling the job, click **Save**.
You can view, edit, or run this job from the **Manage Job Definitions** page. For more information, see [Managing job definitions](#). This job is not displayed on the **Job Status** page until you submit it.
 - 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. For more information, see [Viewing job activity](#) and [Viewing log messages](#).

i | NOTE: A job that uses the Schedule Type “Triggered” is only scheduled when you run the script.

Creating Schedule Sets

A Schedule Set is used to specify scheduling options for backup and restore jobs. These options define when and at what intervals a job runs. You can create Schedule Sets from the **Create Backup Job** and **Create Restore Job** pages.

To create a Schedule Set:

- 1 Start the job configuration wizard, and click **Create New** next to the **Schedule** list.
- 2 On the **Create Schedule Set** page, configure the following options.

Figure 14. Create Schedule Set page

The screenshot displays the 'Create Schedule Set' configuration page. At the top, there are four buttons for 'Schedule Type': 'Immediate', 'Once', 'Repeating' (which is highlighted in orange), and 'Triggered'. Below this is a section titled 'Schedule Options' with a grey header. It contains two input fields: 'Run at' with the value '12:19 PM' and a clock icon, and 'Starting from' with the value '5/11/2017' and a calendar icon. Underneath, the 'Schedule method' section has four buttons: 'Every day' (highlighted in orange), 'On days of week', 'On days of month', and 'Every'. The bottom section is titled 'Job Options' and includes a checkbox for 'Job Retries' (checked), a spinner set to '1', and a 'Retry After' field with '01:00' and a clock icon. Below that is a 'Job Priority' spinner set to '30' with the text '1 = Highest, 100 = Lowest Priority'.

Table 34. Schedule type

Schedule type	Description
Immediate	To run a job as soon as it is submitted, select this option.
Once	<p>To run a job once on the specified days, select this option, and configure the following options:</p> <ul style="list-style-type: none"> • Run at: Type or select the start time for the job. • Starting from: Type or select the date on which the schedule takes effect. • Schedule method: Select a scheduling method and configure the required options. The available methods are: Any day, On days of week, On days of month, and On specified date. For more information, see Scheduling methods and options for non-repeating jobs.
Repeating	<p>To create a recurring schedule for jobs that are performed regularly, select this option, and configure the following options:</p> <ul style="list-style-type: none"> • Run at: Type or select the start time for the job. • Starting from: Type or select the date on which the schedule takes effect. • Schedule method: Select a scheduling method and configure the required options. The available methods are: Every day, On days of week, On days of month, and Every. For more information, see Scheduling methods and options for repeating jobs. <p>NOTE: For repeating jobs, the first instance is scheduled when you submit the job. The next instance is scheduled when the current instance becomes active, and this procedure is repeated for each subsequent instance.</p> <p>You cannot use the Repeating schedule type to run Secondary Copy jobs.</p>
Triggered	<p>To schedule a job from an external script, select this option.</p> <p>The most common use of this option is to run a job independently of the NetVault Scheduler such as from a 3rd-party scheduler or an automation interface.</p> <p>To schedule a triggered job, do the following:</p> <ul style="list-style-type: none"> • Create an external script file, and include the following command in the script: <pre>nvtrigger <trigger_name></pre> <p>A trigger name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. On Linux OS, the names can have a maximum of 64 characters. On Windows OS, there is no length restriction, but a maximum of 20 characters is recommended. On Windows OS, the following characters are not supported:</p> <pre>" / \ : ; * ? < > ^</pre> <p>The nvtrigger utility is stored in the bin directory under the NetVault installation directory. If this path is not configured in the path variable, provide the complete file path. Alternatively, include commands to change to the appropriate directory in the script.</p> <p>You can run the script from the command line interface.</p> • When creating the Schedule Set, select the Triggered option. In the Trigger Name box, provide the same trigger name that was specified with the nvtrigger command in the external script file. <p>NOTE: You cannot use the Triggered schedule type to run Secondary Copy jobs.</p>

3 Under **Job Options**, configure the following settings.

Table 35. Job retry and priority settings for Schedule Set

Option	Description
Job Retries	<p>This setting is used to automatically reschedule a job after a failed attempt. To schedule retry attempts for a job, do the following:</p> <ul style="list-style-type: none"> Select the Job Retries check box, and in the value box type or select a value from 1 through 10. You can set a maximum of 10 retries for a job. In the Retry After box, type or select the interval between two attempts. By default, the job is scheduled to run immediately after a failed attempt. <p>NOTE: For each retry attempt, the same Job ID number is used, but the instance ID number is increased by 1.</p>
Job Priority	<p>The Schedule Manager assigns a default priority level to each job type (backup, restore, and report). These default settings are applied globally to all backup, restore, and report jobs. The priority settings are used to prioritize resource allocation when two or more jobs are scheduled to run at the same time.</p> <p>By default, the Schedule Manager assigns the following priority levels:</p> <ul style="list-style-type: none"> Backup job: 30 Restore job: 20 Report job: 50 <p>You can override the global priority setting for an individual job by configuring the Job Priority option in the Schedule Set.</p> <p>To override the default priority settings for an individual job:</p> <ul style="list-style-type: none"> In the Job Priority box, type or select a value from 1 through 100. 1 denotes highest priority, while 100 denotes lowest priority. A job with a priority level of zero (0) runs as a background task. <p>To change the priority level settings globally for all jobs, see Configuring default job priority settings.</p>

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

The set name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. On Linux OS, the name 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.

Click **Save** to save the Schedule Set.

Scheduling methods and options for non-repeating jobs

The Schedule Type “**Once**” offers the following methods and options.

Table 36. Scheduling methods for non-repeating jobs

Option	Description
Any day	Run a job on any day after the schedule takes effect.
On days of week	Run a job on specific days of the week.
	<p>Options:</p> <ul style="list-style-type: none"> Days: Select the days on which you want to run the job. Weeks: Select the weeks on which you want to run the job. To run the job on the last week of a month, select the Last option.

Table 36. Scheduling methods for non-repeating jobs

Option	Description
On days of month	Run a job on specific days of month. Options: <ul style="list-style-type: none"> Select the days on which you want to run the job. To run the job on the last day of a month, select the Last option.
On specified date	Run a job on a specific date.

Scheduling methods and options for repeating jobs

The Schedule Type “**Repeating**” offers the following methods and options.

Table 37. Scheduling methods and options for repeating jobs

Option	Description
Every day	Run a job daily at the specified time.
On days of week	Run a job on specific days of the week. Options: <ul style="list-style-type: none"> Days: Select the days on which you want to run the job. Weeks: Select the weeks on which you want to run the job. To run the job on the last week of a month, select the Last option.
On days of month	Run a job on specific days of month. Options: <ul style="list-style-type: none"> Select the days on which you want to run the job. To run the job on the last day of a month, select the Last option.
Every	Select this option to run a job at every <n> interval. The time interval can be specified in hours, days, weeks, or months. Options: <ul style="list-style-type: none"> Run every: Type or select the interval at which you want to run the job, and select Hours, Days, Weeks, or Months.

Creating Target Sets

A Target Set is used to specify target device and media options for backup jobs. You can create Target Sets from the **Create Backup Job** page.

To create a Target Set:

- 1 Start the backup job wizard, and next to the **Target Storage** list, click the **+** icon.

Figure 15. Create Backup Job page

- 2 On the **Backup Target** details drawer, click the corresponding button, and configure the options described in the following sections:
 - [Selecting the device](#)
 - [Specifying media options](#)
 - [Configuring media sharing options](#)
- 3 In the **Set Name** text box, enter a name for the Target Set.

The set name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. On Linux OS, the name 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.
- 4 Click **Save**.

Selecting the device

To select a device for a backup job:

- 1 On the **Backup Target** details drawer, click **Device Selection**, and configure the following settings.

Table 38. Device Selection options for Target Set

Option	Description
Use any device	This option is selected by default. If you do not specify a device, NetVault uses any suitable device for a job.
Specify devices	To use particular devices for a job, select this option. In the associated box, clear the check marks for the devices that you do not want to use. When you exclude a library, the associated drives are automatically excluded. You can also use the following buttons to select or clear devices: <ul style="list-style-type: none"> • Unselect all: Clears all devices. • Toggle selection: Clear the selected devices and select the cleared devices.
Use locally attached devices only	To use only devices that are locally attached to the target client, select this check box.

NOTE: Disk-based devices (RAS devices) are considered network-attached or non-local devices. A local device is given priority over a disk-based (RAS) device; a local VTL or SCSI device is given priority over a local NDMP filer-attached device.

If you select the **Any Device** option, NetVault tries to find a suitable local device (local to the NetVault client). To use any other device, you can set a higher media request weight for that device. For more information about media request weightings, see [Configuring media request weightings](#).

Specifying media options

To configure media options for a backup job:

- 1 On the **Backup Target** details drawer, click **Media Options**, and configure the following settings.

Figure 16. Media Options

Table 39. Media Options for Target Set

Option	Description
Storage Tier	Optionally, select a storage tier from the drop-down list. For information about how to create a storage tier, see Creating a storage tier . NOTE: If you select a Storage Tier and select Target Media By Media in group , then only media common to both selections will be used as a target.

Table 39. Media Options for Target Set

Option	Description
Target Media By	<p>Select one of the following options:</p> <ul style="list-style-type: none"> • Any media not in a group: This option is selected by default. To use media items that do not belong to any media group, leave this option selected. • Any media: To use any suitable media item regardless of its group association, select this option. • Specific media ID: To use particular media items, select this option. In the list of media items, select the label for the target media. • Media in group: To use media items that belong to a particular media group, select this option. In the list of media items, click the group label for the target media. Media group labels are case-insensitive. <p>NOTE: When selecting a target storage set, if you select Specific Media ID, then the storage tier will become disabled and set to the default empty storage tier.</p>
Performance	<p>To configure Media Manager to prioritize targets for optimal efficiency, select Optimize target selection for best deduplication.</p> <p>With this option, Media Manager examines the length of each segment in a backup job to determine the best media target to use.</p>
Label Media	<p>To automatically label blank media items during backup, select Automatically Label BLANK Media.</p> <p>By default, NetVault assigns a system-generated label to blank media. The label consists of the NetVault Server Name, the current date, and a seed number.</p> <p>You can configure NetVault to use media barcodes as the default labels. For more information, see Configuring general settings for Media Manager.</p> <p>NOTE: If a piece of media that previously appeared to contain data becomes unexpectedly blank, it is marked as “suspect” to ensure that auto-labeling does not occur for it. NetVault does not permit the use of same label, whether user-specified or system-generated, to ensure that only one label is associated with a piece of media. A notification event occurs when any piece of media becomes unexpectedly blank.</p>
Reuse media	<p>Select the appropriate option from the following:</p> <ul style="list-style-type: none"> • Never: This option is selected by default. To ensure that media marked for re-use are not used for backup, leave this option selected. • Any: To allow any suitable reusable backup media regardless of the group association, select this option. • With the same group label as target media: To reuse media that belong to same media group as the target media, select this option. This option can only be set if you have specified a group label for the Target media by option. <p>NOTE: A piece of media is automatically marked for reuse when the last saveset stored on it is retired. To manually mark a piece of media for re-use, see Marking a tape for reuse.</p>

Table 39. Media Options for Target Set

Option	Description
Media Request Timeout	<p>Type or select the amount of time NetVault waits for the backup media. The timeout value is specified in number of minutes. If the required media item is not available within the specified interval, the job is automatically aborted.</p> <p>The default value is zero (0). With the default setting, the job waits indefinitely until the media is provided or the job is manually aborted.</p> <p>The maximum timeout period that can be set is 1440 minutes (24 hours), which ensures that current instance is aborted before NetVault runs the next instance of a Daily job.</p> <p>NOTE: The timeout interval for media requests cannot be less than 5 minutes. The interval is automatically set to the nearest 5 minutes, if you specify a value that is not a multiple of 5. For example, if you specify 8 minutes, the timeout value is set to 10 minutes, and if you specify 23 minutes, the timeout value is set to 25 minutes.</p>

i | **NOTE:** Backups to disk-based devices are stream-oriented. The concept of media does not apply to disk-based devices. The media options **Target Media By**, **Label Media**, and **Reuse Media** do not apply to disk-based storage devices.

If a media group contains tape media and NetVault SmartDisk devices, NetVault always uses the physical or virtual tape media from that group. To write backups to a NetVault SmartDisk device, you must explicitly select the device.

Configuring media sharing options

To configure media sharing options:

- 1 On the **Backup Target** details drawer, click **Media Sharing**, and configure the following settings.

Table 40. Media Sharing options for Target Set

Option	Description
Ensure backup is the first on the target media	<p>To write a backup at the beginning of a tape, select this check box. The data can be restored quickly if the saveset is located at the beginning of a physical tape.</p> <p>Only new or blank media are eligible for backups that use this option.</p> <p>NOTE: When performing multi-stream backups, do not select the Ensure backup is the first on the target media check box. If you select this check box for multi-stream backups, each data stream targets a separate piece of media to exist as the first backup on the media item. Thus, if a backup generates five streams, the job tries to obtain five blank or new media items.</p>
Protect media from further writes after backup	<p>To write-protect a piece of media after the backup completes, select this option.</p> <p>NOTE: To mark any existing tapes as “read-only,” see Marking a tape as read-only.</p>
Only use media with a minimum of <x> gigabytes of free space	<p>To specify the minimum amount of space required on the target media, type or select the value. The value must be specified in GB.</p>

i | **NOTE:** The **Media Sharing** options do not apply to disk-based storage devices.

Creating Source Sets for backup jobs

Source Sets are used to specify source device options. Source Sets are required for Plug-in *for Consolidation* jobs, Plug-in *for Data Copy* jobs, and Secondary Copy jobs. You can create Source Sets from the **Create Backup Job** page.

To create a Source Set:

- 1 Start the backup job wizard, and click **Create New** next to the **Source Storage** list.
- 2 On the **Backup Job Wizard — Create Source Set** page, click **Device Selection**, and configure the following settings.

Table 41. Device Selection options for Source Set

Option	Description
Any Device	This option is selected by default. If you do not specify a device type, NetVault uses any suitable device for a job.
Specify Device	To use particular devices for a job, select this option. In the associated box, clear the check marks for the devices that you do not want to use. When you remove a library, the associated drives are automatically removed.
Local Drives Only	To use only devices that are locally attached to the target client, select this check box. NOTE: NetVault SmartDisk is considered a network-attached device or a non-local device.

- 3 Click **Save**, and in the **Create New Set** dialog box, type a name for the set.
The set name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. On Linux OS, the name 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.
Click **Save** to save the Restore Source Set.

Creating Backup Advanced Options Sets

An Advanced Options Set is used to specify backup retirement settings, Secondary Copy job definitions, and other advanced options. You can create a Backup Advanced Options Set from the **Create Backup Job** page.

To create a Backup Advanced Options Set:

- 1 Start the backup job wizard, and click **Create New** next to the **Advanced Options** list.
- 2 On the **Advanced Options** page, click the corresponding button, and configure the options described in the following sections:
 - [Setting backup retirement options](#)
 - [Specifying additional options](#)
 - [Creating a Secondary Copy](#)
 - [Configuring pre- and post script options for backup jobs](#)
 - [Configuring user-defined events for backup jobs](#)
- 3 Click **Save**, and in **Advanced Options Set Name**, type a name for the set.

The set name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. On Linux OS, the name 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.

Click **Save** to save the Advanced Options Set.

Setting backup retirement options

To set the backup retirement options:

- 1 On the **Advanced Options** page, click **Backup Life**, and configure the following settings.

Figure 17. Backup Life dialog box

Archive
 Backup

Backup Life Don't Discard based on Full Backup Count
 Discard After Full Backup Count

Backup Life Don't Discard based on Time
 Discard After

The backup will be discarded after both the Full Backup Count and Time specified above, if both criteria are active.

Force Expiry. If checked, the backup will be retired according to schedule and may cause early retirement of dependent incremental or differential backups. If unchecked, the backup will be marked for retirement on schedule, but removal will be deferred until all dependent backups are ready for retirement. See the NetVault Administrator's Guide "Backup retirement" section for information on retirement.

Offline Index After

Table 42. Backup Life options

Option	Description
Archive	To create an archive, select this option. An archive cannot be used as a base for Incremental or Differential backups. When archiving data, you must always select the Full Backup type. If you select Incremental or Differential Backup type, the restore job fails.
Backup	This option is selected by default. To create a backup, use this option.
Backup Life	This option specifies how long a backup is retained. A backup can be retained indefinitely or retired after a specified period. NetVault supports generation- and time-based retirement methods for backups. For more information about backup retirement methods, see Backup retirement . You can specify the retirement options in the Backup Advanced Options Set. Alternatively, you can use the Change Expiry method to set or change the expiry date or generation count for existing savesets. For more information about the Change Expiry method, see Configuring saveset expiry options . To set the Backup Life option, do the following: <ul style="list-style-type: none"> • Backup life — based on Full Backup count: To configure generation-based retirement, select the Discard After Full Backup Count option. In the associated list, type or select the maximum number Full Backups that you want to retain. • Backup life — based on time: To configure time-based retirement, select the Discard After option. Type or select the period, and in the associated list, select the Days, Weeks, Months, or Years option. By default, the Backup Life option is set to Discard After and the retirement period is set to three months.

Table 42. Backup Life options

Option	Description
	<ul style="list-style-type: none"> • Force Expiry: By default, if a backup has any dependent backups, its retirement is deferred until all dependent backups are ready for retirement. You can select this check box to retire a backup according to its retirement schedule. Forcing this behavior can cause early retirement of dependent Incremental and Differential Backups. To apply this rule globally to all backups, you can modify the Media Manager settings. For more information, see Configuring retirement rules for dependent backups. If the Retirement Timing Control option is set to Force Always, the Force Expiry is option is used, regardless of state of this check box. • Retain a backup indefinitely: To retain a backup indefinitely, select the Don't Discard based on Full Backup Count and Don't Discard based on Time options. <p>NOTE: If you set both Discard After Full Backup Count and Discard After options for a Full Backup, the backup is retired only when both conditions are met. For example, if you set the Discard After Full Backup Count option to four Full Backups and the Discard After option to 30 days, the backup is retired after five Full Backup counts and 30 days.</p> <p>NOTE: When you set time-based retirement, the time component (HH:MM) is automatically set to the job save or job submit time.</p> <p>NOTE: In time-based retirement, the time component (HH:MM) does not represent the actual retirement time. It only represents the time due for backup retirement. The actual time of retirement is determined by the interval at which Media Manager scans the Media Database to identify the backups that it needs to retire. The default interval between two scans is 60 minutes. Thus, if the retirement time is set to 10:20, the backup is actually retired at 11:00. You can change this default setting in the mediamgr.cfg file. For more information, see Configuring default interval for backup retirement scans.</p>
Offline Index After	<p>Online backup indexes allow you to quickly scan through the contents of a saveset without loading the media. However, these indexes can increase the overall size of the NetVault Database. To manage the database size, you can delete the online indexes.</p> <p>You can use different methods to delete online indexes:</p> <ul style="list-style-type: none"> • You can configure the Offline Index After option in the Backup Advanced Options Set to automatically delete the online index for a backup after the specified period. This method is described in this section. • You can use the Days of inactivity before an index is offlined setting to configure a global policy for deleting online indexes after a specified period of inactivity. For more information, see Configuring Media Manager settings for backup indexes. • You can manually delete the online indexes for one or more backups from the Manage Backup Indexes page. For more information, see Manually deleting online indexes. <p>To specify the index retention period in the Backup Advanced Options set, select Offline Index After option. Type or select the length of time that you want to retain the index, and in the associated list, select the Days, Weeks, Months, or Years option.</p>

2 Click **Set** to save the settings and close the dialog box.

- i** | **NOTE:** When a backup stored on a disk-based storage device (such as NetVault SmartDisk, Quest DR Series system, or Data Domain system) is retired, that backup is deleted from the device. You cannot import the deleted backup by scanning the device.

Specifying additional options

To specify additional options for a backup job:

- 1 On the **Advanced Options** page, click **Additional Options**, and configure the following settings.

Table 43. Additional backup job options

Option	Description
Enable Encryption	<p>NetVault provides two encryption plug-ins:</p> <ul style="list-style-type: none">• NetVault Plug-in for <i>Standard Encryption</i> (Plug-in for <i>Standard Encryption</i>)• NetVault Plug-in for <i>Advanced Encryption</i> (Plug-in for <i>Advanced Encryption</i>) <p>These plug-ins provide support for CAST-128, AES-256, and CAST-256 algorithms to meet regulatory backup security requirements. For more information about these plug-ins, see the <i>Quest NetVault Plug-in for Standard Encryption User's Guide</i> or <i>Quest NetVault Plug-in for Advanced Encryption User's Guide</i>.</p> <p>Once the Plug-in for <i>Standard Encryption</i> or Plug-in for <i>Advanced Encryption</i> is installed on a client, you can do either of the following:</p> <ul style="list-style-type: none">• Configure the plug-in to encrypt all backups originating for the client where the plug-in is installed. For more information about this setting, see relevant the plug-in user's guide.• Use the job-level encryption option to encrypt specific backups for the client. You can also select encryption only for the secondary copy. <p>The job-level encryption option is useful in the following situations:</p> <ul style="list-style-type: none">• Any plug-in installed on the server or client is incompatible with the encryption plug-ins.• Only specific backups on the server or client require encryption.• Primary backups do not require encryption, while secondary backups for offsite protection require encryption.• Primary backups are targeted to storage devices that support deduplication. <p>To perform job-level encryption for a primary backup, select the Enable Encryption check box. For more information about using job-level encryption for a secondary copy, see Encrypt Secondary Copy Only.</p> <p>NOTE: An encrypted backup can be restored to the original client or an alternate client. In both cases, the plug-in must be installed on the target client and it must be configured as it was when the backup was performed — using the same Encryption Key String and Encryption Algorithm.</p>

Table 43. Additional backup job options

Option	Description
Enable Deduplication	<p>Deduplication is enabled by default. Clear this check box if the target device does not support data deduplication.</p> <p>When performing backups to devices that support deduplication, we recommend that you clear this check box for the following jobs:</p> <ul style="list-style-type: none"> • Backups that use the encryption option. Encrypted backups do not deduplicate well and should not be deduplicated. • Incremental Backups that you want to consolidate using the <i>Plug-in for Consolidation</i>. By not selecting this option, you eliminate the unnecessary overhead of rehydrating the deduplicated Incremental Backups during the consolidation process. You can enable deduplication while backing up the Consolidated Full Backup. <p>NOTE: You cannot completely disable deduplication for a Quest DR Series system. The Quest DR Series systems provide a configuration mode for deduplication that controls whether deduplication is performed on the client or on the Quest DR Series system. You can turn off client-side deduplication by setting the Dedupe mode to Passthrough. For more information about this setting, see the <i>Quest DR Series system Administration Guide</i>.</p> <p>NOTE: When backups stored on the Quest DR Series systems are selected for backup consolidation jobs, the overhead of rehydrating the deduplicated data can have a negative impact on performance.</p>
Verify After Backup	<p>To verify the stream length written to the media and ensure that no blocks were dropped during backup, select this check box.</p> <p>Backup verification is performed as Phase 2 job after the actual backup is completed. If any dropped blocks are detected, the verification phase reports an error and fails. You must run the backup again if the verification phase fails.</p> <p>NOTE: The Phase 2 backup verification job does not verify the integrity of data. This phase only verifies that the backup was actually written to the media.</p> <p>By default, the verification job runs on the NetVault Server. To configure a different client to run the verification phase, see Configuring default settings for Verify Plug-in.</p>
Use Network Compression	<p>To use network compression while transferring data over the network, select this check box.</p> <p>The data is compressed on the backup client before being transferred over the network. On the machine to which the target device is attached, the data is decompressed before being written to the media.</p> <p>Network compression does not work for the following types of jobs:</p> <ul style="list-style-type: none"> • Backups to NetVault SmartDisk • Backups to devices attached to NDMP-based NAS filers • Backups performed using the NetVault Plug-in <i>for NDMP</i>, NetVault Plug-in <i>for NetWare</i>, and NetVault Bare Metal Recovery products.
Don't Add Saveset Items to the Restore Search Catalog	<p>To exclude a backup with a selected saveset from cataloging, select this check box.</p> <p>To check the cataloged status (Yes or No) of a saveset, see the 'Cataloged' column on the Create Restore Job - Choose Saveset page</p>

- 2 Click **Set** to save the settings and close the dialog box.

Creating a Secondary Copy

To create a secondary copy:

- 1 On the **Advanced Options** page, click **Secondary Copy**.
- 2 Click the **Create Secondary Copy** check box, and configure the following settings.

Figure 18. Secondary Copy dialog box

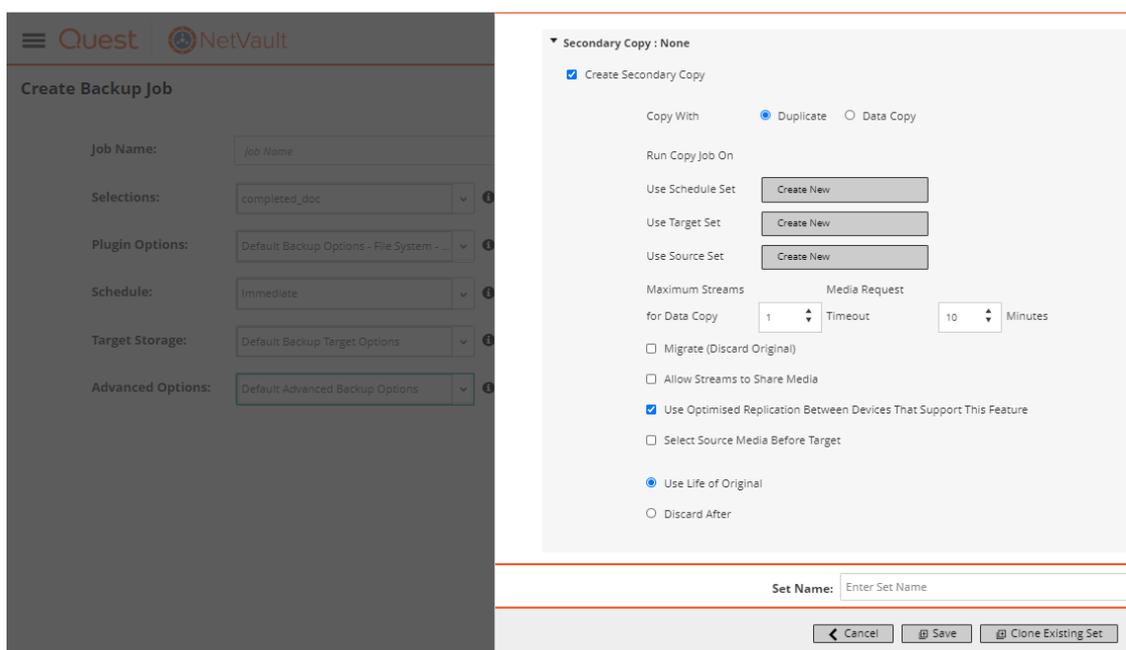


Table 44. Secondary Copy options

Option	Description
Copy with	<p>Select the method that you want to use to create the Secondary Copy. The available methods are:</p> <ul style="list-style-type: none"> • Duplicate • Data Copy <p>For more information about these methods, see Secondary Copy.</p> <p>NOTE: On NetVault Server 10.0.5 and later, the Plug-in <i>for Data Copy</i> creates index version 4, which is incompatible with earlier versions of NetVault. A client running NetVault 10.0.1 or an earlier version cannot read these index files. To restore Data Copy savesets generated from NetVault Server 10.0.5 or later, the client must also be running NetVault 10.0.5 or later.</p>
Run copy job on	<p>By default, the secondary copy job runs on the NetVault Server. If you want to run the job on a particular client, select the target client.</p> <p>You can use this option to perform backups on a client with a locally attached physical or virtual tape device.</p> <p>IMPORTANT:</p>

Table 44. Secondary Copy options

Option	Description
Use Schedule Set	<p>Select an existing Schedule Set, or click Create New, and configure the schedule type and schedule method. For more information, see Creating Schedule Sets.</p> <p>NOTE: You cannot use the Repeating and Triggered schedule types to run Secondary Copy jobs. A Secondary Copy job can run only after the Phase 1 job completes successfully. The Repeating and Triggered schedule types can cause a situation where a Secondary Copy job is configured to start before the source saveset is available.</p>
Use Target Set	<p>Select an existing Target Set, or click Create New, and configure the target device and media options for the job. For more information, see Creating Target Sets.</p> <p>NOTE: When performing Duplicate and Data Copy backups, you cannot use media items that contain the original saveset or a copy of the same backup. NetVault excludes these media items to ensure that copies and original backup do not exist on the same media item. The concept of media does not apply to disk-based devices. If the original saveset is stored on a disk-based device, NetVault does not exclude that device when you run a Duplicate or Data Copy job for that backup.</p> <p>NOTE: We recommend that you select the same drives for all secondary copy jobs. For example, in a library with four drives, select drives 1 and 2 for the primary backups and drives 3 and 4 for the secondary copies targeted to tape devices. This type of selection avoids deadlocks when running several duplication jobs at the same time.</p>
Use Source Set	<p>Select an existing Source Set, or click Create New, and configure the source device options for the job. For more information, see Creating Source Sets for backup jobs.</p>

3 Configure the required additional options.

Table 45. Additional options for Secondary Copy

Option	Description
Maximum Streams for Data Copy	Type or select the maximum number of parallel streams that can be generated for the Data Copy job. By default, a single data stream is generated to copy the data items in a sequential manner.
Media Request Timeout	<p>Type or select the amount of time NetVault waits for the backup media. The timeout value is specified in number of minutes. The default value is 10 minutes.</p> <p>If the required media item is not available within the specified interval, the job is automatically aborted.</p> <p>If you set the value to zero (0), the Secondary Copy job waits indefinitely until the media is provided or the job is manually aborted.</p> <p>The maximum timeout period that can be set is 1440 minutes (24 hours), which ensures that current instance is aborted before NetVault runs the next instance of a Daily job.</p> <p>NOTE: The media request timeout for Phase 1 Data Copy jobs (created using the Plug-in <i>for Data Copy</i>) is controlled by the Media Request Timeout setting in the Backup Options Set; it is not controlled by the Media Request Timeout setting in the Target Set. The media request timeout for Phase 2 (Secondary Copy) Data Copy and Duplicate jobs is controlled by the Media Request Timeout specified in the Backup Advanced Options Set.</p>
Encrypt Secondary Copy Only	<p>To perform job-level encryption for a secondary copy, select the Enable Encryption check box.</p> <p>This option can only be used with the Data Copy method. To use this option, the Plug-in <i>for Standard Encryption</i> or the Plug-in <i>for Advanced Encryption</i> must be installed on the client.</p> <p>Note the following:</p> <ul style="list-style-type: none"> • If the primary copy is encrypted, the Data Copy method automatically creates an encrypted saveset whether you select the Encrypt Secondary Copy Only check box or not. Therefore, this option is only useful when you want to create an encrypted secondary copy from an unencrypted primary copy. • Encrypted primary backups are not encrypted again if you select the Encrypt Secondary Copy Only check box for a Data Copy. • To restore data from an encrypted Data Copy, you must use the primary copy's Encryption Key. <p>For more information about using job-level encryption for a primary backup, see Enable Encryption.</p>
Migrate (Discard Original)	<p>To migrate the backup instead of creating a copy, select this check box. After copying the data, NetVault deletes the index for the original backup.</p> <p>NOTE: When creating copies of the Plug-in <i>for File System</i> backups, the Migrate option can only be selected for Full Backups that do not have any associated Incremental or Differential Backups. If you select this option for a Full Backup that has an associated Incremental or Differential Backup, NetVault creates the secondary copy successfully, but it does not delete the index for the primary or original backup. For such backups, after creating the copy, you must manually retire the primary or original backup.</p>

Table 45. Additional options for Secondary Copy

Option	Description
Allow Streams to Share Media	<p>This check box is not selected by default.</p> <p>You can use this option to convert multiple data streams into a sequential data stream. If you do not select this check box, each stream is written separately.</p> <p>If you are using a disk-based device to store the backup, multiple streams are recommended. For such jobs, there is no need to select this check box.</p> <p>If you do not select the Allow Streams to Share Media check box for tape devices, each data stream targets a separate media item. Any tape that is used to write a stream is excluded from the allowable media items. If enough tapes are not available for the job, the job fails. When using tape devices to store the backups, you can select this check box to use minimal media items.</p>

Table 45. Additional options for Secondary Copy

Option	Description
Use Optimised Replication Between Devices that Support this Feature	<p>Optimized replication enables transfer of deduplicated data directly from one device to another device of the same type during a Data Copy or Duplicate operation. It provides an efficient method to create secondary copies and offers the following advantages:</p> <ul style="list-style-type: none"> • Copies data in its deduplicated form, which greatly reduces the amount for data transferred over the network. • Copies data directly from the source to the destination without using any resources on the NetVault Server. <p>The following storage devices support optimized replication:</p> <ul style="list-style-type: none"> • Quest DR Series systems: To perform optimized replication, both the source and target Quest DR Series systems must be running the same release version of the DR OS. Replication is not supported between systems that run different releases of the OS. For example, to replicate data from a source system that is running DR OS 3.x, the target system must be running the same OS release version. Replication is unsuccessful if the target system is running DR OS release 2.0.x or 3.0.x. <p>NOTE: When optimized replication and backups are performed simultaneously on a Quest DR Series system, the backup throughput is affected.</p> <ul style="list-style-type: none"> • Quest QoreStor Servers • NetVault SmartDisk devices: To perform optimized replication, you require NetVault SmartDisk 2.0 or later. If the login credentials configured for the source and destination NetVault SmartDisk Servers do not match, replication fails. To ensure a successful replication, do one of the following: <ul style="list-style-type: none"> - Disable WebDAV authentication on both NetVault SmartDisk Servers. - Enable WebDAV authentication only on the source server. - Configure the same login credentials on both servers. <p>NOTE: When copying a backup from a NetVault SmartDisk to a different device type (for example, VTL, Quest DR Series system, or Data Domain system), you must clear this check box. If you do not clear this check box, the Data Copy or Duplicate job fails or stops responding.</p> <ul style="list-style-type: none"> • DD Boost-enabled Data Domain systems: The secondary copy backups between two DD Boost-Enabled Data Domain systems use the managed file-level replication feature provided by DD Boost. File-level replication requires the DD Boost Replicator license, which must be installed on both the source and destination Data Domain systems. <p>NOTE: If the source and target Data Domain systems are running different versions of the Data Domain OS, then for replication to be successful, the target system must be running the higher version of the OS.</p>
Select Source Media Before Target Media	<p>When you select this check box, NetVault tries to acquire the source media before it attempts to acquire the target media for the Data Copy and Duplicate backups.</p>

4 To set retirement period for the duplicate or data copy saveset, configure the following options.

Table 46. Retirement options for Secondary Copy

Option	Description
Use Life of Original	This option is selected by default. To use the original saveset's retirement period, leave this option selected.
Discard After	To configure a different retirement period for the Duplicate or Data Copy saveset, select this option. Type or select the length of time that you want to retain the backup, and in the associated list, select the Days , Weeks , Months , or Years option. You can set only time-based retirement period for the copy.
Force Expiry	This check box is displayed when you select the Discard After option. By default, if a Secondary Copy has any dependent backups, its retirement is deferred until all dependent backups are ready for retirement. You can select this check box to retire a Secondary Copy according to its retirement schedule. Forcing this behavior can cause early retirement of dependent Incremental and Differential Backups. To apply this rule globally to all backups, you can modify the Media Manager settings. For more information, see Configuring retirement rules for dependent backups . If the Retirement Timing Control option is set to Force Always , the Force Expiry option is used, regardless of state of this check box.

5 Click **Set** to save the settings and close the dialog box.

i | **IMPORTANT:** Data Copy or Duplicate of a persistent snapshot only backup, copies only the index and does not create the redundant copy of the data.

Configuring pre- and post script options for backup jobs

You can use the pre- and post-script options to run user-defined scripts before a job starts or after a job completes. You can use these scripts to perform tasks such as dismounting or shutting down a database before the job starts or mounting or starting the database after the job completes.

Before configuring pre- or post-scripts, review the following information:

- The script should be an executable file, for example, “.bat” files on Windows and “.sh” files on Linux.
- The script file should reside on the target client. It should be available in the **scripts** directory under the NetVault installation directory.
- The pre- and post-scripts can contain run-time parameters. These parameters are stored in the environment variable **NV_USER_ARG**. You can also use other NetVault environment variables in the scripts. For a list of available environment variables, see [Environment variables in NetVault](#).
- NetVault provides two predefined script files that can be used as post-scripts:
 - **psmail:** Use this script to send the job completion status to the specified email addresses.
 - **psmail_logs:** Use this script to send the job completion status and job logs to the specified email addresses.

On Linux and UNIX, the predefined scripts do not have any filename extension. On Windows, the scripts have the filename extension “.bat.”

To run these scripts, specify the following in the **Post Script** box:

- **Linux and UNIX:** psmail or psmail_logs
- **Windows:** psmail.bat or psmail_logs.bat

To specify pre- and post scripts:

- 1 On the **Advanced Options** page of the target client, click **Pre & Post Scripts**, and configure the following settings.

Table 47. Pre- and post-script options for backup jobs

Option	Description
Pre Script	<p>This option allows you to run a user-defined script before a job starts. You can use this script to perform any pre -backup preparation, such as dismounting or shutting down a database.</p> <p>To run a pre -script, do the following:</p> <ul style="list-style-type: none"> • In the Pre Script box, specify the script filename. • In the User Parameter box, provide the values for the run-time parameters. The value should be valid and conform to its usage in the script. NetVault does not perform any validity checks for the user parameters.
Post Script	<p>This option allows you to run a user-defined script after a job completes. You can use this script to perform any post -backup processing, like mounting or starting a database after a job completes.</p> <p>To run a post -script, do the following:</p> <ul style="list-style-type: none"> • In the Post Script box, specify the script filename. • In the User Parameter box, provide the values for the run-time parameters. The value should be valid and conform to its usage in the script. NetVault does not perform any validity checks for the user parameters.

- 2 Click **Set** to save the settings and close the dialog box.

The following table illustrates the effect of script exit status on the overall execution and job status.

Table 48. Script execution and backup job status

Process	Result				
Pre script	Success	Success	Success	Fail	Fail
Backup job	Success	Success	Fail	Does not run	Does not run
Post script	Success	Fail	Success	Success	Fail
Overall job status	Job completes successfully.	Job completes, but a post-script error is reported.	Job fails, but the post-script runs. An error is reported.	Job fails, and a pre-script error is reported.	Job fails; pre- and post-script errors are reported.

Configuring user-defined events for backup jobs

When creating a backup job, NetVault lets you configure user-defined events that can be raised when the job completes successfully, completes with warnings, or fails. You can use these options when you want to be notified about the completion status of individual jobs.

- i** | **NOTE:** NetVault includes several predefined events, which are raised for significant occurrences within the system. For more information, see [About NetVault events](#) and [Events types](#).
- You can also raise user-defined events for individual backup jobs, restore jobs, report jobs, policies, and log messages. You can use various notification methods to receive notifications when the predefined or user-defined events occur in NetVault. For more information, see [Using global notification methods](#).

You can configure the user-defined events in the Backup Advanced Options Set. Once you create a user-defined event, you can raise it for any job or backup policy. NetVault adds the user-defined events to the event class **Jobs — User defined**. You can then set up one or more notification methods for the event to receive notifications when the event is raised.

To raise user-defined events for a backup job:

- 1 On the **Advanced Options** page, click **Events**, and configure the following settings.

Table 49. User defined event types for backup jobs

Option	Description
Job is Successful	Type or select the event that you want to raise when the job completes successfully.
Job has Warnings	Type or select the event that you want to raise when the job completes with warnings.
Job Fails	Type or select the event that you want to raise when the job fails.

- 2 Click **Set** to save the settings and close the dialog box.
- 3 To receive a notification when an event occurs, set up a notification method for the event. For more information, see the [Setting up a global notification method](#).

Managing sets

This section includes the following topics:

- [Modifying a set](#)
- [Copying a set](#)
- [Deleting a set](#)

Modifying a set

You can modify the data selections or various job options stored in an existing set. When saving the changes, you can choose to overwrite the set or save the selections to a new set.

- i** | **NOTE:** When you edit a set, it affects the existing jobs that use the set.

To modify a set:

- 1 In the Navigation pane, click **Manage Sets**.
- 2 In the **Set Type** list, select the type of set that you want to modify.
- 3 You can use the **Search** option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string.
- 4 In the list of available sets, select the applicable set, and click **Edit**.
- 5 Modify the data selections or job options.
- 6 Click **Save**. In the **Edit Set** dialog box, click **Save** to overwrite the existing set. Alternatively, type a new name, and click **Save** to save the selections to a new set.

Copying a set

You can use the clone method to create a set from an existing set. After the set is cloned, you can change the data selections or various options stored in the set.

NetVault lets you clone a set when creating or modifying a job. You can also use the **Edit** option on the **Manage Sets** page to perform this task.

To copy a set:

- 1 Use the applicable method to access the set that you want to copy:
 - **To clone a set when creating a job:** Access the Create Backup Job or Create Restore Job page, and click the **Create New** button corresponding to the type of set that you want to create.
 - **To clone a set when editing a job:** Access the Edit Backup Job or Edit Restore Job page, and click the Create New button corresponding to the type of set that you want to create.
 - **To clone a set from the Manage Sets page:** In the Navigation pane, click **Manage Sets**. In the **Set Type** list, select the type of set that you want to modify. In the list of available sets, select a set, and click **Edit**.
- 2 On the page that is displayed, click **Clone Existing Set**.
- 3 In the **Choose a set to load** dialog box, select the set that you want to copy, and click **Load**.
- 4 Modify the data selections or job options, if required.
- 5 Click **Save**. In the **Edit Set** dialog box, type a name, and click **Save** to copy the selections to a new set.

Deleting a set

If a set is no longer required, you can remove it from the NetVault Database. You cannot delete a set if it is in use by an active job or a scheduled job.

; | **NOTE:** When you delete a set, it affects the existing jobs that use the set.

To delete a set:

- 1 In the Navigation pane, click **Manage Sets**.
- 2 In the **Set Type** list, select the type of set that you want to delete.
- 3 In the list of available sets, select the applicable set, and click **Delete**.

You can select multiple sets for deletion. To select consecutive items, hold down the Shift key while clicking with the mouse button; to select non-consecutive items, hold down the Ctrl key while clicking with the mouse button.
- 4 In the confirmation dialog box, click **OK**.

Managing policies

- [About policies](#)
- [Creating a policy](#)
- [Viewing existing backup policies](#)
- [Viewing job status of policy jobs](#)
- [Modifying a backup policy](#)
- [Quiescing a backup policy](#)
- [Deleting a backup policy](#)

About policies

A policy can be used to submit one or more jobs that target one or more similar clients.

You can use a policy to administer backup strategies such as the following:

- Daily Incremental and Weekly Full Backups of file servers
- Full backups of multiple Windows workstations
- Full and Incremental Backups of multiple databases.

The following plug-ins support policy-based backups:

- NetVault Plug-in *for FileSystem*
- NetVault Plug-in *for Consolidation*
- NetVault Plug-in *for Data Copy*
- NetVault Plug-in *for Databases* (plug-in for backing up the NetVault Database)
- NetVault Plug-in *for SQL Server*
- NetVault Plug-in *for Oracle*
- NetVault Plug-in *for MySQL*
- NetVault Plug-in *for PostgreSQL*
- NetVault Plug-in *for VMware*
- NetVault Plug-in *for Hyper-V*
- NetVault Plug-in *for Exchange*

You can create and submit backup policies from the **Manage Policies** page and monitor the policy jobs from the **Job Status** page.

Creating a policy

A policy can be used to submit one or more jobs that target one or more similar clients. You can create and submit policy-based backups from the **Manage Policies** page.

To create a policy:

- 1 In the Navigation pane, click **Manage Policies**, and then click **Add** to open the **Edit Policy** page.
- 2 In **Policy Name**, type a name for the policy.
- 3 To add a policy job, click **Add Job**.

On the **Create Policy Job** page, configure the following settings.

Table 50. Policy job definition

Option	Description
Job Name	Type a name for the job. Assign a descriptive name that allows you to easily identify the job for monitoring its progress or restoring data. The job name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms.
Selections	Select an existing Backup Selection Set, or click the + icon, and select the items that you want to back up. The selection tree is specific to a plug-in; it depends on the type of data that you are backing up. For more information about selecting data for backups, see the relevant plug-in user's guide.
Plugin Options	Select an existing Backup Options Set, or click the + icon, and configure the options that you want to use. The backup options are specific to a plug-in; the options depend on the type of data that you are backing up. For more information about these options, see the relevant plug-in user's guide.
Schedule	Select an existing Schedule Set, or click the + icon, and configure the schedule type and schedule method. For more information, see Creating Schedule Sets . The predefined set "Immediate" is selected by default. To run the job as soon as it is submitted, use this set.
Source Storage	Select an existing Source Set, or click the + icon, and configure the source device options for the job. For more information, see Creating Source Sets for backup jobs . This option is only available to Plug-in <i>for Consolidation</i> and Plug-in <i>for Data Copy</i> jobs.
Target Storage	Select an existing Target Set, or click the + icon, and configure the target device and media options for the job. For more information, see Creating Target Sets . The predefined set "Default Backup Target Options" is selected by default.
Advanced Options	Select an existing Backup Advanced Options Set, or click the + icon, and configure the options that you want to use. For more information, see Creating Backup Advanced Options Sets . The predefined set "Default Advance Backup Options" is selected by default.

Click **Save** to save the job definition.

- 4 The table on **Edit Policy** page shows the following information:
 - **Active:** By default, the check box is selected and the policy job is saved and submitted in the policy. To save the policy job without scheduling it, clear the check box.
 - **Job Name:** Name of the job.
 - **Selection Set:** Backup selection set.
 - **Plugin:** Name of the plug-in used to perform the job.
 - **Next Run Time:** Date and time when the next instance is scheduled to run OR 'Never' for non-scheduled policy jobs. This field shows the information only after the policy is saved.
- 5 To add more jobs, repeat [Step 3](#).
- 6 Click **Add Clients**.
- 7 To apply the policy to one or more clients or client groups, do the following:

Table 51. Client and client group selection for backup policy

Option	Description
Add clients or client groups	<p>In the Available table, select the clients and client groups that you want to add, and click Add.</p> <p>To select consecutive items, hold down the Shift key while clicking with the mouse button; to select non-consecutive items, hold down the Ctrl key while clicking with the mouse button.</p> <p>The selected items are moved to the Selected table.</p> <p>NOTE: The client status icons are displayed, which indicate whether the client is online or offline. The policy jobs are successful only if the selected client is online. For more information about client status icons, see Table 15.</p>
Remove clients or client groups	<p>In the Selected table, select the clients and client groups that you want to remove, and click Remove.</p> <p>The selected items are moved to the Selected table.</p>

- 8 To save the policy definition without configuring the events for Policy warnings and failures, click **Save Policy**, else Click **Add Events**.
- 9 When creating a backup policy, NetVault lets you configure user-defined events that can be raised when one or more policy jobs complete with warnings or fail.

You can use these options when you want to be notified about the completion status of a policy.

i NOTE: NetVault includes several predefined events, which are raised for significant occurrences within the system. For more information, see [About NetVault events](#) and [Events types](#).

You can also raise user-defined events for individual backup jobs, restore jobs, report jobs, policies, and log messages. You can use various notification methods to receive notifications when the predefined or user-defined events occur in NetVault. For more information, see [Using global notification methods](#).

To raise a user-defined event for the policy, configure the following settings. Once you create a user-defined event, you can raise it for any policy. NetVault adds these events to the event class **Policy — User defined**.

Table 52. User-defined events for backup policy

Option	Description
Raise event if policy has warnings	Type or select the event that you want to raise when one or more policy jobs complete with warnings.
Raise event if policy has errors	Type or select the event that you want to raise when one or more policy jobs fail.

To receive a notification when an event occurs, set up a notification method for the event. For more information, see the [Setting up a global notification method](#).

- 10 Click **Save Policy** to save the policy definition.

Viewing existing backup policies

You can view information about the existing policies from the **Manage Policies** page. The page displays the status, policy name, number of active jobs, state, total clients, scheduled jobs, job status, and total jobs.

To view the existing backup policies:

- 1 In the Navigation pane, click **Manage Policies**.
- 2 On the **Manage Policies** page, the **Available Policies** table is displayed.
 - **Status:** Displays the status of the policy using the following icons. Click the icon to display **Job Status** page showing the Job Activity table of all the jobs in the corresponding policy.

Table 53. Policy status icons

Icon	Description
	The corresponding policy has completed successfully. It indicates that all jobs have completed successfully.
	The corresponding policy has completed with warnings. It indicates that one or more jobs have completed with warnings.
	The corresponding policy has failed. It indicates that one or more jobs have failed.

- **Policy Name:** Displays the name of the backup policy.
- **Number of active jobs:** Displays the number of jobs that are currently active in the policy. To preview the job name and job ID of all the active jobs in the policy, move the pointer over the count.
- **State:** Displays the policy state. The policy state can be Active, Dormant, Quiescing, or Quiesced.
- **Total Clients:** Displays the total number of clients in the policy. To preview the client name, move the pointer over the count.
- **Scheduled Jobs:** Displays the total number of scheduled jobs in the policy. To preview the job name and job ID of all the scheduled jobs in the policy, move the pointer over the count.
- **Job Status:** Displays the last exit status and count of the policy jobs using the following icons. Click the corresponding icon to display **Job Status** page showing the Job Activity table of all the policy jobs that have 'completed', 'completed with warnings', or 'failed'.

Table 54. Policy jobs status icons and counts

Icon	Description
	The corresponding policy jobs have completed successfully. The count below the icon shows the number of healthy policy jobs. To preview the job name and job ID of all the 'completed' jobs in the policy, move the pointer over the count.
	The corresponding policy jobs have completed with warnings. The count below the icon shows the number of completed policy jobs with warning. To preview the job name and job ID of all the 'completed jobs with warnings' in the policy, move the pointer over the count.
	The corresponding policy jobs have failed. The count below the icon shows the number of failed policy jobs. To preview the job name and job ID of all the 'failed jobs' in the policy, move the pointer over the count.

- **Total Jobs:** Displays the total number of jobs in the policy. To preview the job name and job ID of all the jobs in the policy, move the pointer over the count.

- 3 By default, the table is sorted by Policy Name.

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To set the filter options (For more information on table filter options for **Manage Policies** page, see [Table 11](#)), view the page size setting, column sort order, applied filters in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table.

For more information, see [Customizing tables in NetVault WebUI](#).

- 4 You can use the **Search** option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string.
- 5 To perform a policy-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

Viewing job status of policy jobs

You can view the job status of all the policy jobs from the **Manage Policies** page.

To view the job activity of policy jobs in a policy:

- 1 In the Navigation pane, click **Manage Policies**.
- 2 In the **Available Policies** table, select the policy to view the status of policy jobs, and click **View Status**.
- 3 **Job Status** page is displayed showing the Job Activity table of **All Policy jobs**.

Modifying a backup policy

You can modify the job definitions, client or client groups selections, or event settings for an existing backup policy.

This section includes the following topics:

- [Modifying jobs in a backup policy](#)
- [Modifying clients in a backup policy](#)
- [Modifying events in a backup policy](#)

Modifying jobs in a backup policy

To modify the jobs in a policy:

- 1 In the Navigation pane, click **Manage Policies**.
- 2 In the **Available Policies** table, select the policy that you want to modify, and click **Manage Jobs**. The **Manage Job** page is displayed.

The table on **the** Mange Job page shows the following information: Active, Job Name, Selection Set, Plugin and Next Run Time.

- 3 To add or remove jobs or edit the job definitions, do the following:
 - **Add Job:** On the **Edit Policy** page, click **Add Job**, and create the job definitions. For more information about the job components, see [Policy job definition](#).
Click **Save** to the save the job definition.
 - **Edit Job:** In the Jobs table, select the job that you want to change, and click **Edit Job**.

On the **Manage Job** page, modify the required job components. For more information, see [Policy job definition](#).

Click **Save** to save the job definition.

- **Remove Jobs:** In the Jobs table, select the job that you want to delete, and click **Remove Jobs**. In the confirmation dialog box, click **OK**.
- 4 To save the policy definition without modifying client (s) and events, click **Save Policy**, else click **Edit Clients** or **Edit Events**.

Modifying clients in a backup policy

To modify the clients in a policy:

- 1 In the Navigation pane, click **Manage Policies**.
- 2 In the **Available Policies** table, select the policy that you want to modify, and click **Edit Clients**.
To change the client or client groups for the policy, see [Client and client group selection for backup policy](#). To save the policy definition without modifying the events and jobs, click **Save Policy**, else click **Edit Events** or **Manage Jobs**.
- 3 To change the user-defined events for the policy failures or policy warnings, see [User-defined events for backup policy](#).
- 4 Click **Save Policy** to save the policy definition.

Modifying events in a backup policy

To modify the events in a policy:

- 1 In the Navigation pane, click **Manage Policies**.
- 2 In the **Available Policies** table, select the policy that you want to modify, and click **Edit Events**.
To change the user-defined events for the policy failures or policy warnings, see [User-defined events for backup policy](#).
- 3 To save the policy definition without modifying the client (s) and jobs, click **Save Policy**, else click **Edit Clients** or **Manage Jobs**.
- 4 Click **Save Policy** to save the policy definition.

Quiescing a backup policy

An active backup policy is temporarily disabled when it is placed in a quiesced state.

To quiesce an active backup policy:

- 1 In the Navigation pane, click **Manage Policies**.
- 2 In the **Available Policies** table, select the policy, and click **Quiesce**.
In the policy table, the policy state is set to "**Quiescing**."
- 3 During this state, NetVault completes the following tasks:
 - Deletes all scheduled instances for the policy jobs.
 - Completes the jobs that are in progress.

- Complete phase 2 (for example, a Secondary Copy job) for the active jobs.
- 4 After these operations are completed, the policy state is set to **“Quiesced.”**
In this state, you can change the policy definition.
 - 5 A **“Quiesced”** policy remains in that state until you open and save the policy again. When you save the policy, all the policy jobs are scheduled again.

Deleting a backup policy

If a backup policy is no longer required, you can delete it from the database.

To delete a backup policy:

- 1 In the Navigation pane, click **Manage Policies**.
- 2 In the **Available Policies** table, select the policy that you want to delete, and click **Remove**.
- 3 In the confirmation dialog box, click **OK**.

Restoring data

- [About restoring data](#)
- [Creating restore jobs](#)
- [Restoring data using an existing Restore Selection Set](#)
- [Searching for files in savesets](#)
- [Viewing media list](#)
- [Creating Source Sets for restore jobs](#)
- [Creating Restore Advanced Options Sets](#)
- [Managing online backup indexes](#)
- [Managing Snapshots](#)

About restoring data

Restore refers to reconstructing all or part of a system from a backup.

You can perform a restore job for various reasons, such as the following:

- To recover lost or deleted data
- To recover database or files that have been corrupted
- To copy or move data to a different database or directory
- To recover to a previous point-in-time, if some operation goes wrong
- To migrate data when upgrading to a new system
- To copy or move data to a test or production server
- To recover from media failure, OS corruption, or loss of physical system

The NetVault plug-ins integrate with the native APIs to restore and recover application-specific data from backups. Depending on the application type, these plug-ins provide different methods and options to restore data.

In general, NetVault offers the following restore features:

- Full and selective restores
- Disaster recovery
- Restores to alternate location
- Restores to alternate server
- Explore Snapshot

Restore job definition

To restore data, you need to create and submit a restore job. You can create a restore job definition from the **Create Restore Job** link in the Navigation pane.

A restore job definition includes the following components:

- Selection list
- Plug-in options
- Target client name (when restoring to an alternate server)
- Scheduling options
- Source device options
- Advanced restore options

These components are stored in NetVault Sets. For more information about NetVault Sets, see [About NetVault Sets](#).

Each restore job has a Job ID number and a Job Name. The Job ID number is an auto-generated number. The Job Name is a user-defined string that allows you to easily identify the job when monitoring its progress or viewing the job logs.

Creating restore jobs

You can use the restore job wizard to create and submit restore jobs. The wizard can be accessed from the **Create Restore Job** link in the Navigation pane.

To create a restore job:

- 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.

Figure 19. Create Restore Job — Choose Saveset page

Status	Saveset Name	Selection Set	Plugin	Created	Size	Cataloged
●	fs back2 (Sav...	test fs set2	File System	03/19/2020 1...	48.50 MiB	✓
●	682_backup (L...	Full_DB_ORCL1	Plug-in for Or...	03/18/2020 1...	9.05 GiB	✗
●	RMAN contro...	NVBU Oracle...	Plug-in for Or...	03/18/2020 1...	11.50 MiB	✗
●	RMAN contro...	NVBU Oracle...	Plug-in for Or...	03/18/2020 1...	11.50 MiB	✗
●	RMAN contro...	NVBU Oracle...	Plug-in for Or...	03/18/2020 1...	11.50 MiB	✗
●	test 700 (Sav...	test tablespa...	Plug-in for Or...	03/18/2020 1...	2.25 MiB	✗
●	RMAN contro...	NVBU Oracle...	Plug-in for Or...	03/18/2020 1...	11.50 MiB	✗
●	test 642 (Sav...	test tablespa...	Plug-in for Or...	03/18/2020 1...	2.25 MiB	✗
●	RMAN contro...	NVBU Oracle...	Plug-in for Or...	03/18/2020 1...	11.50 MiB	✗

The table shows the saveset status, saveset name (Job Title and Saveset ID), selection set, plug-in, creation date and time, saveset size, and cataloged status.

The saveset status is indicated using the following icons.

Table 55. Saveset status 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. For more information, see [Sorting records in NetVault WebUI](#).

The **Create Restore Job — Choose Saveset** page loads a maximum of 5000 records initially. The total number of records fetched is displayed at the lower-right corner of the table.

You can click **Load more** to load the next set of records, if available. Each load action fetches a maximum of 5000 records. This button is disabled if there are no more records to load.

- To set the filter options, view the page size setting, change the sorting order in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#) and [Table 13](#)).

You can use the **Search** option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string. For more information, see [Searching for files in savesets](#).

- Select the saveset that you want to use.

i | **NOTE:** If the online backup index is unavailable for the selected saveset, the **Confirm load** dialog box is displayed. To load the index from the backup media, click **OK**, and then in the **Load Index** dialog box, type or select the number of days you want to store the index in the NetVault Database.

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, saveset size, and cataloged status.

- Click one of the following options:
 - Restore All Using Defaults.** This option selects the entire saveset and pre-populates the details on the Create Restore Job page. It then navigates you to the Create Restore Job page where you can optionally change the details by using the buttons beside each text box. The job name appears in the following format: Restore_Saveset savesetid_system date and time _jobid+1. Optionally, you can change the job name. For more information about changing the restore details, continue to the next step. To restore using the pre-populated defaults, skip to [Step 11](#).
 - Restore.** This option opens the **Create Selection Set** page. On the **Create Selection Set** page, select the items that you want to restore.

The data items that are available for selection depend on the plug-in in use. For more information about selecting data for restores, see the relevant plug-in user's guide.

Click **Edit Plugin Options**, and configure the options that you want to use, and then click **Next**.

i | **NOTE:** The restore options that are available to a job depend on the plug-in in use. For more information about these options, see the relevant plug-in user's guide.

- On the **Create Restore Job** page, specify a name for the job.

Assign a descriptive name that allows you to easily identify the job for monitoring its progress. The job name can contain alphanumeric and non-alphanumeric 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, select the restore target. To restore data to the same client (from which data was backed up), use the default setting.

To restore data to an alternate client, select the target client in the list. Alternatively, click **Choose**. In the **Choose the Target Client** dialog box, select the client, and click **OK**.

- 7 In the **Schedule** list, select an existing Schedule Set, or click **Create New**, and configure the schedule type and schedule method. For more information, see [Creating Schedule Sets](#).

The predefined set “**Immediate**” is selected by default. To run the job as soon as it is submitted, use this set.

- 8 In the **Source Options** list, select an existing Source Set, or click **Create New**, and configure the source device options. For more information, see [Creating Source Sets for restore jobs](#).

The predefined set “**Any Device**” is selected by default. To select any available device for the job, use this set.

- 9 In the **Advanced Options** list, select an existing Restore Advanced Options Set, or click **Create New**, and configure the options that you want to use. For more information, see [Creating Restore Advanced Options Sets](#).

The predefined set “**Restore from selected backup**” is selected by default.

i | **NOTE:** To create a restore job, you can understand more about the sets by previewing the set information. To preview the set summary, move the pointer over the fields on **Create Restore Job** page.

- 10 To save your selections but not submit the restore job for scheduling, click **Save**.
- 11 To save any changes you may have made and submit the restore 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. For more information, see [Viewing job activity](#) and [Viewing log messages](#).

Restoring data using an existing Restore Selection Set

You can use the following procedure to restore data using an existing Restore Selection Set. This procedure is useful if a user role is not granted the “Jobs — Administer backup/restore sets” privilege, which is required to create or edit a Restore Selection Set.

To restore data using an existing Restore Selection Set:

- 1 In the Navigation pane, click **Manage Sets**.
- 2 In the **Set Type** list, select **Restore Selection**.
- 3 In the list of available sets, select the set that you want to use, and click **Edit**.
- 4 On the **Edit Selection Set** page, click **Next**.

-OR-

On the **Edit Selection Set** page, modify the data selections, set name, or plug-in option.

Click **Next**. In the Confirm overwrite dialog box, click **OK** to overwrite the existing set.

i | **NOTE:** To preview the selection set summary and plug-in options set summary, move the pointer over the information icons on **Edit Selection Set** page.

- 5 Complete [Step 5](#) through [Step 11](#) in the section [Creating restore jobs](#).

Searching for files in savesets

The **Search All / Search Selected Saveset(s)** 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.

To conduct a catalog search within a saveset, click the bulb icon shown near **Search All or Search Selected Saveset(s)** button on the **Create Restore Job — Choose Saveset** page.

To search for data items in savesets:

- 1 On the **Create Restore Job — Choose Saveset** page, select a saveset and then click **Search Selected Saveset(s)**.

If you do not want to select a saveset, the default option to click is **Search All**.

- 2 In the **Search for files in savesets** details drawer, enter the following criteria.

i | **NOTE:** By default, the simple search uses the wildcards "*" meaning "match anything" and "?" meaning "match one character." To use a full-featured Portable Operating System Interface for Unix (POSIX) regular expression, select **Regular expression search** next to each applicable option.

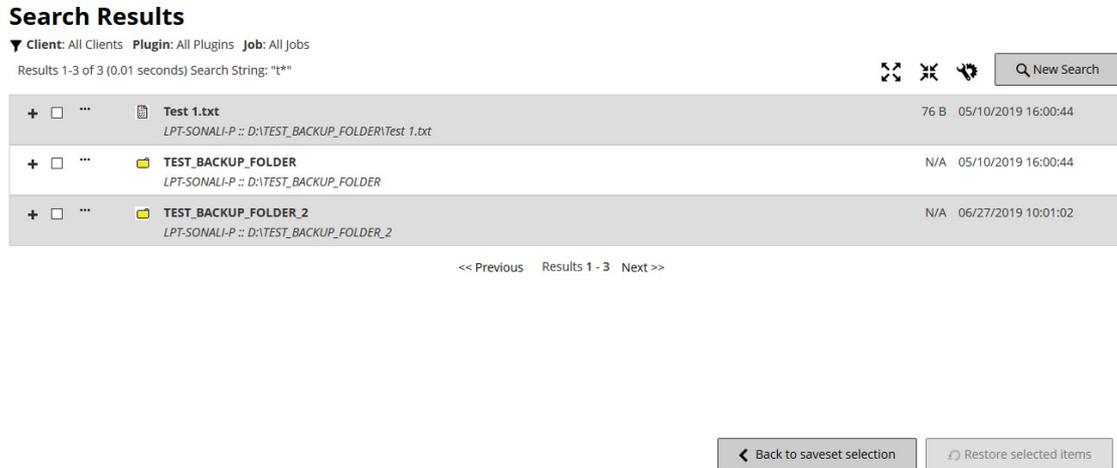
Table 56. Search criteria for files in savesets

Option	Description
Name	Enter the name of the file that you want to find.
Path	Enter the location of the file that you want to find.
Search	In the text box, enter a string of words that are associated with the selected criteria.
Job Name	Enter the job name that applies to the file that you want to find.
Backup Selection Set	Enter the name of the backup selection set that contains the file that you want to find.
Backup Date	Under From , enter the earliest time and date of your search range in the respective text boxes. Under To , enter the latest time and date of your search range in the respective text boxes.
Modified Time	Under From , enter the earliest time and date of your search range in the respective text boxes. Under To , enter the latest time and date of your search range in the respective text boxes.
Size in Bytes	Enter the Min (minimum) and Max (maximum) size, in bytes, for the search range.
Plugin fields	Expand to reveal the available options, and then select the plug-in or plug-ins that relate to your search.
Client fields	Expand to reveal the available options, and then select the client or clients that relate to your search.
Job fields	Expand to reveal the available options, and then select the job or jobs that relate to your search.
Use legacy search method	In cases when a saveset that is not cataloged is included in the search, the option to use the legacy search method appears. With this option, only a search bar is available to conduct the search. Catalog filters do not apply to savesets that are not cataloged.

3 Click **Search**.

The details drawer closes and redirects you to the Search Results page.

Figure 20. Search Results with catalog search enabled



The following information about the search results is displayed:

- **Filter** icon: To filter the records displayed in the search results table, click this icon. For more information on filter options, see [Table 14](#).
 - 'Search filters' drawer is displayed on the right- side of the page. Set the following filter options and click **Apply**.
 - To clear the filter settings, click **Clear**.
 - To edit the applied filters, clear the filter settings, and then apply new filter settings.
 - To close the 'Search filters' drawer without editing or setting the filter options, click **Cancel**.
- The total number of search results and number of records shown on the page may not be equal because, the search results shows the number of occurrences however, the rows shows the number of non-identical items found. If copies of an item are found, then the row displays the copy icon. To view the copies of the files and select for restore, click the copy icon. Copies are defined as items with identical machine, path, name, size, and modified time.
- **Expand** icon: To expand all result items, click the expand icon shown on top of the results table.
- **Collapse** icon: To collapse all result items, click the collapse icon shown on top of the results table.
- **Settings** icon: To set the number of search results per page, click the settings icon. 'Settings' drawer is displayed on the right- side of the page. Type or select the value and click **OK**. By default, the page displays 25 records. You can view a minimum 10 and maximum 100 records on the **Search Results** page.
- **New Search** button: To search the data items in the selected saveset, click **New Search**.
- The **Search for Files in Savesets** details drawer opens. Repeat steps 1-3.To expand an item and view the file metadata, such as plug-in, job name, job ID, saveset ID, and backup date and time, click the corresponding plus icon.
- To view the media information for a saveset, click the corresponding context- menu icon and then click 'show media list for the saveset'. Status (online or offline) of the storage media and other information is displayed in the dialog box. To close the dialog box, click **Close**.
- Copy icon is displayed if the corresponding result is present in multiple savesets.

- Filename
 - The **Filepath** link appears below the filename. To search within this location, click the link. The **Search for Files in Savesets** drawer opens with the **Path** automatically populated.
 - Filesize
 - Last modified date and time
- 4 Select the items you want to restore. You can only restore items from one saveset. Click **Restore selected items**.
 - 5 Complete [Step 5](#) through [Step 11](#) in the section [Creating restore jobs](#).

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.

To view the media details for a saveset:

- 1 On the **Create Restore Job — Choose Saveset** page, select the applicable saveset.
- 2 In the **Saveset Information** area, click **Media List**.
- 3 In the dialog box that appears, 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.
- 4 Click **Close** to close the dialog box.

Creating Source Sets for restore jobs

A Source Set is used to specify source device options. You can create Source Sets for restore jobs from the **Create Restore Job** page.

To create a Source Set:

- 1 Start the restore job wizard, and click **Create New** next to the **Source Storage** list.
- 2 On the **Create Restore Source Set** page, click **Device Selection**, and configure the following settings.

Table 57. Device Selection options for Restore Source Set

Option	Description
Any Device	This option is selected by default. If you do not specify a device type, NetVault uses any suitable device for a job.

Table 57. Device Selection options for Restore Source Set

Option	Description
Specify Device	To use particular devices for a job, select this option. In the associated box, clear the check marks for the devices that you do not want to use. When you remove a library, the associated drives are automatically removed.
Local Drives Only	To use only devices that are locally attached to the target client, select this check box. NOTE: NetVault SmartDisk is considered a network-attached device or a non-local device.

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

The set name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. On Linux OS, the name 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.

Click **Save** to save the Restore Source Set.

Creating Restore Advanced Options Sets

A Restore Advanced Options Set is used to specify restore type, pre- and post-script, and other advanced options. You can create a Restore Advance Options Set from the **Create Restore Job** page.

To create a Restore Advanced Options Set:

- 1 Start the restore job wizard, and click **Create New** next to the **Advanced Options** list.
- 2 Configure the options described in the following sections:
 - [Setting restore type](#)
 - [Specifying additional options](#)
 - [Configuring pre- and post -scripts for restore jobs](#)
 - [Configuring user-defined events for restore jobs](#)
- 3 Click **Save**, and in the **Create New Set** dialog box, type a name for the set.

The set name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. On Linux OS, the name 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.

Click **Save** to save the Restore Advanced Options Set.

Setting restore type

To specify the restore type:

- 1 On the **Advanced Options** page, click **Restore Type**, and select one of the following options.

Table 58. Restore type

Option	Description
Restore from selected backup	This option is selected by default. To restore data from the selected backup, leave this option selected.
Restore from latest backup	Select this option to restore data from most recent backup that was performed using a particular Backup Selection Set regardless of which saveset you use to create the restore job.

The following example illustrates the difference between these two options:

- a Create a test directory and create a Backup Selection Set “SelectionSet-1” to back up the test directory.
 - b Create a file named “first.txt” in the test directory.
 - c Perform a Full Backup (TestBackup1) using “SelectionSet-1.”
 - d Delete “first.txt” and create a file named “last.txt” in the test directory.
 - e Perform a second Full Backup (TestBackup2) using “SelectionSet-1.”
 - f Restore TestBackup1 using the **Restore from selected backup** option. This job restores the file “first.txt.”
 - g Restore TestBackup1 using the **Restore from latest backup** option. This job restores the file “last.txt.”
- 2 Click **Set** to save the settings and close the dialog box.

Specifying additional options

To specify additional options for a restore job:

- 1 On the Advanced Options page, click **Additional Options**, and configure the following option:
 - **Use Network Compression:** To use network compression while transferring data over the network, select this check box. The data is compressed on the server or client to which the source device is attached before being transferred over the network. On the target client, the data is decompressed before being restored to the original or alternate location.

Network compression does not work for the following types of jobs:

 - Restores from NetVault SmartDisk
 - Restores from devices attached to NDMP-based NAS filers
 - Restores performed using the Plug-in *for NDMP*, Plug-in *for NetWare*, and NetVault Bare Metal Recovery products
- 2 Click **Set** to save the settings and close the dialog box.

Configuring pre- and post -scripts for restore jobs

You can use the pre- and post-script options to run user-defined scripts before a job starts or after a job completes. You can use these scripts to perform tasks such as dismounting or shutting down a database before the job starts or mounting or starting the database after the job completes.

Before configuring pre- or post-scripts, review the following information:

- The script should be an executable file, for example, “.bat” files on Windows and “.sh” files on Linux.

- The script file should reside on the target client. It should be available in the **scripts** directory under the NetVault installation directory.
- The pre- and post-scripts can contain run-time parameters. These parameters are stored in the environment variable **NV_USER_ARG**. You can also use other NetVault environment variables in the scripts. For a list of available environment variables, see [Environment variables in NetVault](#).
- NetVault provides two predefined script files that can be used as post-scripts:
 - **psmail**: Use this script to send job completion status to the specified email addresses.
 - **psmail_logs**: Use this script to send job completion status and job logs to the specified email addresses.

On Linux and UNIX, the predefined scripts do not any filename extension. On Windows, the scripts have the filename extension “.bat.” To run these scripts, specify the following in the **Post Script** box:

- **Linux and UNIX:** **psmail** or **psmail_logs**
- **Windows:** **psmail.bat** or **psmail_logs.bat**

To specify pre- and post -scripts:

- 1 On the **Advanced Options** page, click **Pre & Post Scripts**, and configure the following settings.

Table 59. Pre and post script options for restore jobs

Option	Description
Pre Script	<p>This option allows you to run a user defined script before a job starts. You can use this script to perform any pre -restore preparation, like dismounting or shutting down a database.</p> <p>To run a pre -script, do the following:</p> <ul style="list-style-type: none"> • In the Pre Script box, specify the script filename. • In the User Parameter box, provide the values for the run -time parameters. The value should be valid and conform to its usage in the script. NetVault does not perform any validity checks for the user parameters.
Post Script	<p>This option allows you to run a user-defined script after a job completes. You can use this script to perform any post restore processing, such as mounting or starting a database after a job completes.</p> <p>To run a post -script, do the following:</p> <ul style="list-style-type: none"> • In the Post Script box, specify the script filename. • In the User Parameter box, provide the values for the run-time parameters. The value should be valid and conform to its usage in the script. NetVault does not perform any validity checks for the user parameters.

- 2 Click **Set** to save the settings and close the dialog box.

Table 60. Script execution and restore job status

The following table illustrates the effect of script exit status on the overall execution and job status.

Process	Result					
Pre script	Success	Success	Success	Fail	Fail	Fail
Restore job	Success	Success	Fail	Does not run	Does not run	Does not run

Process	Result					
Post script	Success	Fail	Success	Success	Not configured	Fail
Overall job status	Job completes successfully	Job completes, but a post-script error is reported	Job fails, but the pre-script and post-script runs. An error is reported	Job fails, and a pre-script error is reported, but post-script runs	Job fails, and a pre-script error is reported	Job fails. Pre-script and post-script errors are reported

Configuring user-defined events for restore jobs

When creating a restore job, NetVault lets you configure user-defined events that can be raised when the job completes successfully, completes with warnings, or fails. You can use these options when you want to be notified about the completion status of individual jobs.

- i** **NOTE:** NetVault includes several predefined events, which are raised for significant occurrences within the system. For more information, see [About NetVault events](#) and [Events types](#).
- You can also raise user-defined events for individual backup jobs, restore jobs, report jobs, policies, and log messages. You can use various notification methods to receive notifications when the predefined or user-defined events occur in NetVault. For more information, see [Using global notification methods](#).

You can configure the user-defined events in the Restore Advanced Options Set. Once you create a user-defined event, you can raise it for any job or backup policy. NetVault adds the user-defined events to the event class **Jobs — User defined**. You can then set up one or more notification methods for the event to receive notifications when the event is raised.

To raise user-defined events for a restore job:

- 1 On the Advanced Options page, click **Events**, and configure the following settings.

Table 61. User-defined event types for restore jobs

Option	Description
Job is Successful	Type or select the event that you want to raise when the job completes successfully.
Job has Warnings	Type or select the event that you want to raise when the job completes with warnings.
Job Fails	Type or select the event that you want to raise when the job fails.

- 2 Click **Set** to save the settings and close the dialog box.
- 3 To receive a notification when an event occurs, set up a notification method for the event. For more information, see the [Setting up a global notification method](#).

Managing online backup indexes

This section includes the following topics:

- [About managing online indexes](#)
- [Manually deleting online indexes](#)
- [Loading offline indexes](#)

- [Manually compressing online indexes](#)
- [Uncompressing online indexes](#)

About managing online indexes

Online backup indexes allow you to quickly scan through the contents of a saveset without loading the media. However, these indexes can increase the overall size of the NetVault Database. To manage the database size, you can delete or compress the online indexes.

The following sections provide a brief description of the methods that are available for deleting and compressing online indexes:

- [Deleting online indexes](#)
- [Compressing online indexes](#)

Deleting online indexes

You can use the following methods to delete the online indexes for backups:

- You can configure the **Offline Index After** option in the Backup Advanced Options Set to automatically delete the online index for a backup after the specified period. For more information, see [Setting backup retirement options](#).
- You can use the **Days of inactivity before an index is offlined** setting to configure a global policy for deleting online indexes after a specified period of inactivity. For more information, see [Configuring Media Manager settings for backup indexes](#).
- You manually delete the online indexes for one or more backups from the **Manage Backup Indexes** page. For more information, see [Manually deleting online indexes](#).

The indexes for the savesets can be loaded from the backup media if required.

i | **NOTE:** Deleting a backup index is not the same as retiring a backup. When a backup is retired, NetVault discards all information about it from the NetVault Database. When you scan the media to retrieve the index for a retired saveset, it is loaded as a new index in the NetVault Database. When the online indexes are deleted, NetVault still retains some information about the saveset. This information allows NetVault to quickly reload the saveset index from the backup media.

Compressing online indexes

Online indexes are automatically compressed after 30 days of inactivity. You can modify the **Days of inactivity before an index is compressed** setting to customize this policy. For more information, see [Configuring Media Manager settings for backup indexes](#).

You can also manually compress the online indexes for one or more backups from the **Manage Backup Indexes** page. For more information, see [Manually compressing online indexes](#).

Manually deleting online indexes

You manually delete the online indexes for one or more backups from the **Manage Backup Indexes** page.

To manually delete online indexes:

- 1 In the Navigation pane, click **Create Restore Job**, and then on the Choose Saveset page, click **Manage Indexes**.
 - 2 Click **Choose Client**, and select the client for which the backup was created.
- OR-

To quickly locate the client (s), you can use the Search box at the upper right corner of the table. The data in the table is filtered as you type the search text into the box.

- 3 Click **OK** to close the dialog box.
- 4 Optionally, click **Choose Plugin**, and select the plug-in that was used to create the backup. Click **OK** to close the dialog box.
- 5 Click the View icon to display the savesets for the selected client and plug-in.
- 6 In the savesets list, all items are selected by default.

To delete indexes for specific savesets, clear the check marks for the savesets that you want to exclude. Alternatively, click the check box in the header row to remove all check marks, and select the individual savesets.

- 7 Click **Offline**.

Loading offline indexes

When restoring data, you can load the indexes from the backup media if the online indexes are unavailable for the savesets. You can also use the **Manage Backup Indexes** page to load the backup indexes for one or more backups from the backup media. The indexes are stored in the database for a specified period.

To load offline indexes:

- 1 In the Navigation pane, click **Create Restore Job**, and then on the Choose Saveset page, click **Manage Indexes**.
- 2 Click **Choose Client**, and select the client for which the backup was created.

-OR-

To quickly locate the client (s), you can use the Search box at the upper right corner of the table. The data in the table is filtered as you type the search text into the box.

- 3 Click **OK** to close the dialog box.
- 4 Optionally, click **Choose Plugin**, and select the plug-in that was used to create the backup. Click **OK** to close the dialog box.
- 5 Click the View icon to list the savesets for the selected client and plug-in.
- 6 In the savesets list, all items are selected by default.

To load indexes for specific savesets, clear the check marks for the savesets that you want to exclude. Alternatively, click the check box in the header row to remove all check marks, and select the individual savesets.

- 7 Click **Load**, and in the **Load Index** dialog box, configure the following option:
 - **Days to Keep Index:** Type or select the number of days you want to store the index in the NetVault Database. By default, the index is retained for one day.
- 8 Click **OK** to close the dialog box.

Manually compressing online indexes

Online indexes are automatically compressed after 30 days of inactivity. You can also manually compress the online indexes for one or more backups from the **Manage Backup Indexes** page.

To manually compress online indexes:

- 1 In the Navigation pane, click **Create Restore Job**, and then on the Choose Saveset page, click **Manage Indexes**.

- 2 Click **Choose Client**, and select the client for which the backup was created.

-OR-

To quickly locate the client (s), you can use the Search box at the upper right corner of the table. The data in the table is filtered as you type the search text into the box.

- 3 Click **OK** to close the dialog box.
- 4 Optionally, click **Choose Plugin**, and select the plug-in that was used to create the backup. Click **OK** to close the dialog box.
- 5 Click the View icon to list the savesets for the selected client and plug-in.
- 6 In the savesets list, all items are selected by default.

To compress indexes for specific savesets, clear the check marks for the savesets that you want to exclude. Alternatively, click the check box in the header row to remove all check marks, and select the individual savesets.

- 7 Click **Compress**.

Uncompressing online indexes

When you try to browse or restore data from savesets with compressed indexes, the indexes are automatically de-compressed to a temporary directory. You can also use the **Manage Backup Indexes** page to manually de-compress the indexes for one or more backups.

The temporary directory is deleted after the operation completes.

To manually uncompress online indexes:

- 1 In the Navigation pane, click **Create Restore Job**, and then on the Choose Saveset page, click **Manage Indexes**.
- 2 Click **Choose Client**, and select the client for which the backup was created.

-OR-

To quickly locate the client (s), you can use the Search box at the upper right corner of the table. The data in the table is filtered as you type the search text into the box. Click **OK** to close the dialog box.

- 3 Optionally, click **Choose Plugin**, and select the plug-in that was used to create the backup. Click **OK** to close the dialog box.
- 4 Click the View icon to list the savesets for the selected client and plug-in.
- 5 In the savesets list, all items are selected by default.

To de-compress indexes for specific savesets, clear the check marks for the savesets that you want to exclude. Alternatively, click the check box in the header row to remove all check marks, and select the individual savesets.

- 6 Click **Uncompress**.

Managing Snapshots

Explore snapshots is a feature that allows you to view, mount, and unmount the existing snapshots.

Explore snapshot feature is currently supported for Dell Storage Manager, Huawei OceanStor Manager, and NetApp ONTAP Device Manager. To add these storage manager's refer [Managing the Snapshot Array Manager](#). Currently, NetVault Plug-in for *FileSystem* supports snapshot-based backups.

You can perform following operations:

- [Restoring snapshots](#)

- [Setting expiry for Snapshots](#)
- [Mounting Snapshots](#)
- [Unmounting Snapshots](#)

i | **NOTE:** If you create snapshots in version 12.0 and upgrade NetVault server to 12.1, the snapshot mount and unmount operation does not support. In this case you need to upgrade NetVault client to 12.1.

Restoring snapshots

Restoring data from a snapshot involves a simple copy operation from the snapshot directory to the target/secondary system. The restore operation overwrites the existing volume configuration. Any changes that are made to the data in the volume after the Snapshot copy is created, are lost.

To restore snapshots refer, [Creating restore jobs](#).

Setting expiry for Snapshots

Each snapshot has an expiration date and time, you can preserve a snapshot to prevent it from expiring. Snapshot expires when the backup Saveset expires, even if the snapshot expiry date is greater than Saveset Expiry date. However, if this snapshot is already mounted, the snapshot goes in the pending state if the saveset has expired.

To set an expiry date of snapshot

- 1 In the navigation pane, click **Create Restore Job**.
- 2 On the **Create Restore Job — Choose Saveset** page, select the SaveSet Name and then click Explore Snapshots.
- 3 Select the snapshot from the **Saveset Name** list and click **Set Expiry**.
- 4 Select **Change Expiry Date**, and choose the options to set the expiry date and click **Ok**.
 - **On:** Set the date from the calendar and set the time.
 - **Now:** Snapshot expires Immediately.
 - **Expire with Saveset:** Snapshot never expires.

i | **NOTE:** In case of Huawei OceanStor Device Manager and NetApp ONTAP Device Manager, you can only set an option to expire snapshot immediately.

In the **Explore Snapshot** page, the **Expires** column updates the date and time.

Mounting Snapshots

You can mount any available snapshot, to access the data available in the snapshot. Snapshots must be mounted on a MediaAgent with the same operating system as the snapshot volumes. Currently, snapshots are mounted only on Windows machine. You can mount these snapshots on a particular drive or a specified empty folder. Once the snapshot is mounted you can view the details of each snapshot, which helps during the restore operation.

i | **NOTE:** In the case of Huawei OceanStor Device Manager, to mount a snapshot, the host where the LUN is mounted and host where the snapshot will be mounted must be in same host group.

NOTE: In the case of NetApp ONTAP Device Manager, to mount a snapshot, the host on which the LUN is mounted and the host where the snapshot will be mounted must be in the same initiator group.

To mount a snapshot

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

- 2 On the **Create Restore Job — Choose Saveset** page, select the **SaveSet Name** and then click **Explore Snapshots**.
- 3 Select the snapshot from the **Saveset Name** list to mount and click **Mount**.
- 4 Select the machine to mount the selected snapshot, and click **Next**.

You can mount the snapshot on multiple client machines.

- 5 Select the drive or folder path to mount the snapshot.

To mount the snapshot on drive:

- Select **Mount using drive letter**, and select the drive from the list, and click **Mount**.

i | **NOTE:** The drive letter of the respective client machine is not available which is already used by the client machine.

To mount the snapshot in a folder:

- Select **Mount in folder**, and browse the folder for the mount point, and click **Mount**.
You can select empty folder from any of the drives available for the selected machine.
- In the Explore Snapshot page, the status in **Mounted** column changes to **Yes**.

Unmounting Snapshots

The unmount snapshot operation removes the snapshot that is mounted on the drive or folder.

To unmount a snapshot

- 1 In the navigation pane, click **Create Restore Job**.
- 2 On the **Create Restore Job — Choose Saveset** page, select the SaveSet Name and then click Explore Snapshots.
- 3 Select the snapshot from the **Saveset Name** list to unmount and click **Unmount**.

i | **NOTE:** Check status of the snapshot in **Mounted** column.

- 4 Select the client machine where the snapshot is mounted to unmount, and click **Ok**.

In the Explore Snapshot page, the status in **Mounted** column changes to **No**.

Managing NetVault dashboard

- [About NetVault dashboard](#)
- [Configuring Dashboard](#)
- [Viewing NetVault dashboard](#)
- [Managing a widget on NetVault dashboard](#)

About NetVault dashboard

You can view the overall statistics of the NetVault system from the **Dashboard** page.

The dashboard is a set of graphical widgets that displays the data relevant to your NetVault system and the machines you protect. By default, the data in the widgets is refreshed according to a pre-defined refresh interval time however, you can also manually enter the refresh interval.

The dashboards feature enables you to create multiple configurable graphical views of your environment. You can construct and configure a dashboard and load it to access data of the monitored widgets for NetVault WebUI. A default dashboard is created and loaded by NetVault. A dashboard can have maximum of eighteen widgets. You can perform various actions on the dashboard like share, clone, delete etc.

Viewing NetVault dashboard

To view the NetVault dashboard, in the Navigation pane, click **Dashboard**.

The dashboard includes the following widgets:

- **Job Status:** This widget shows the status of NetVault jobs, and is represented by  icon. The individual bars in this widget represent the total number of successful jobs (green), completed jobs with warnings (yellow), and failed jobs (red). To view the job summary, move the pointer over the bar. By default, the chart displays the data for the last 7 days. However, you can set the widget filter to change the period. To view the jobs on the **Job Status** page, click the corresponding bar. Default refresh interval time for this widget is 900 seconds.
- **Job Duration:** This widget shows the job duration (in seconds) of the top 10 (default) time consuming NetVault jobs and is represented by  icon. You can set the widget filter to change the number of jobs. To view the job summary, move the pointer over the bar. To view the job details of any job on the **Job History** page, click the corresponding bar. Default refresh interval time for this widget is 3600 seconds.
- **Media Space Usage:** This widget shows the amount of media space used and available in your NetVault system and is represented by  icon. By default, 40 media are displayed. You can set the widget filter to change the number of media for the widget. To view the media summary, move the pointer over the bar. To explore the media item, click the corresponding bar. Default refresh interval time for this widget is 1800 seconds.

- **Client Status:** This widget shows the online or offline status of existing NetVault Clients and is represented by  icon. This widget appears as a pie chart. To view the client details summary, move the pointer over the corresponding slice. To view the client details on **Manage Clients** page, click the corresponding slice on the pie chart. Default refresh interval time for this widget is 900 seconds.
- **Unusable Media:** This widget shows the details of unusable media in your NetVault system and is represented by  icon. Each unusable media is displayed as a bubble in the chart with different colors, where the color indicates information for unusable media, see [Table 62](#). To view the media summary, move the pointer over the bubble/circle. To explore the media item, click the corresponding bubble chart. Default refresh interval time for this widget is 1800 seconds.

Table 62. Color legends for unusable media

Unusable media	Color
ASF	Purple
Bad Media	Red
Foreign Media	Brown
Media is Full	Grey
Media is Offline	Orange
Other Media	Blue
Media is read Only	Green
Media is Unusable	Yellow

- **Device Status:** This widget shows the device status of RAS devices, tap libraries, and tape drives and is represented by  icon. Each device is represented as a separate slice in the pie chart. To view the device summary, move the pointer over the pie chart. To explore the device status, click the corresponding slice on the pie chart. Default refresh interval time for this widget is 1800 seconds.
- **Policy Status:** This widget shows the details of backup policies in your NetVault system and is represented by  icon. Each policy is represented as a separate slice in the pie chart. To view the policy summary, move the pointer over the pie chart. To explore the policy status on the **Manage Policies** page, click the corresponding slice on the pie chart. Default refresh interval time for this widget is 900 seconds.
- **Event Stream:** This widget shows the events that occurred in NetVault Server during a time interval and is represented by  icon. By default, the error events for last 10 minutes are displayed. You can set the widget filter to change the time interval or event type (Error Only, Key Events, and All Events). To view the event summary, move the pointer over the bubble. To explore an event, click the corresponding bubble on the chart. Default refresh interval time for this widget is 60 seconds.
- **Log Stream:** This widget shows the log stream for all NetVault processes and is represented by  icon. By default, the log stream for last 10 minutes is displayed. You can set the widget filter to change the time interval, level, or class of the logs. To view the log summary, move the pointer over the bubble. To explore a log on the **View Logs** page, click the corresponding bubble on the chart. Default refresh interval time for this widget is 60 seconds.
- **Running Jobs:** This widget shows the status and job history of NetVault running jobs and is represented by  icon. The individual bars in this widget represent the running job which displays the current job's last three instances information like, duration, bytes transferred, and exit status. This information is displayed by hovering the mouse on the pointer over bar. To view the job details of any job on the **Job Status** page, click the corresponding bar. Default refresh interval time for this widget is 900 seconds. You can also set different filter like Client, Plugin, Policy, Job Type, Record Limit and Duration.

Configuring Dashboard

Configure dashboard allows you to edit the loaded dashboard. You can load any dashboard from the manage dashboard panel.

To configure the loaded dashboard

- 1 In the Navigation pane, click **Dashboard**.
- 2 In the dashboard page, click the ellipsis-horizontal  icon and click **Configure Dashboard**.
- 3 Edit the required fields:
 - **Dashboard Title:** Name of the dashboard.
 - **Share with usergroups and users:** Specifies to select the list of available users and user groups to share the dashboard.
 - **NOTE:** Usergroups are displayed only if NetVault Server is connected to Active Directory. For more information see, [Managing AD groups](#).
 - **Widgets to be updated parallely:** Set the number of widgets to be updated at the same time.
 - **NOTE:** Widgets are updated in parallel when the **Refresh Interval** for any of widgets matches a specific time slot.

Managing Dashboard

You can manage the dashboard available in manage dashboard panel. The default dashboard is identified by the opaque star icon. You can set any dashboard as default, this dashboard is loaded when you select the dashboard from the navigation pane. You can perform various actions for the available dashboards:

- Add Dashboard
- Share Dashboard
- Clone Dashboard
- Delete Dashboard

Adding Dashboard

You can create a new dashboard, and select the widgets to monitor, these widgets are mapped to that dashboard. However, if you do not select any widget while creating a dashboard, a blank dashboard is created. You can add widgets to this blank dashboard later.

To add a new dashboard

- 1 In the Navigation pane, click **Dashboard**.
- 2 In the dashboard page, click the ellipsis-horizontal  icon and click **Manage Dashboard**.
- 3 In the **Manage Dashboard** dialog box, click **Add Dashboard**.
- 4 In the Add Dashboard dialog box, enter and select the following details, and click **OK**.
 - **Dashboard Title**
 - **Select the widgets for the dashboard**
 - Select the **Set as a default dashboard**, to set this as a default dashboard.

The newly created dashboard is listed in Manage Dashboard list. Once you create a new dashboard, you can perform following actions:

- Share dashboard with others
- Clone dashboard
- Delete Dashboard

Sharing dashboard

You can share your owned dashboard that is available in manage dashboard list with any users or usergroups. The user or user group can only view the shared dashboard, refresh and resize the widgets. However they are not permitted to re-share or modify the dashboard, The user can only refresh and resize the available widgets. The user is allowed to clone shared dashboard, and the user will be the owner of cloned dashboard.

To share a dashboard

- 1 In the Navigation pane, click **Dashboard**.
- 2 In the dashboard page, click the ellipsis-horizontal  icon and click **Manage Dashboard**.
- 3 Click the  icon, and select **Share Status** for **Share with usergroups and users**, to share the dashboard, and click OK.

Once the dashboard is shared, the icon is changed to , which indicates that the dashboard is shared with some user or usergroup. The  icon is displayed to the user or user group with whom the dashboard is shared.

Cloning dashboard

Clone dashboard allows you to duplicate the dashboard where all the widgets mapped with the dashboard are also duplicated/cloned.

To clone a dashboard

- 1 In the Navigation pane, click **Dashboard**.
 - 2 In the dashboard page, click the ellipsis-horizontal  icon and click **Clone Dashboard**.
- or--

Click **Manage Dashboard**, and click **Clone**  icon, against the required dashboard.

Deleting dashboard

Delete dashboard removes or deletes the available dashboard from the manage dashboard list. Deleting the dashboard also removes all the customizations made in the dashboard and in the widgets which are associated with the dashboard. In case, the owner of the dashboard deletes the dashboard which is shared with many users or usergroups, the dashboard is removed or deleted for all the users and usergroups. However, if the user is not the owner of dashboard and deletes the dashboard, it would be removed for the particular user. You can delete all the dashboards, however a default dashboard is created and loaded by NetVault, as there is no dashboard available.

To delete a dashboard

- 1 In the Navigation pane, click **Dashboard**.

- 2 In the default dashboard page, click the **ellipsis-horizontal**  icon and click **Manage Dashboard**.
- 3 Click **Delete**  icon against the required dashboard to delete.

Managing a widget on NetVault dashboard

On the NetVault dashboard page, you can perform the following operations to customize a widget:

- Refresh a widget
- Set and clear the filters for a widget
- Clone a widget
- Configure a widget
- Delete a widget
- Resize a widget

To manage a widget, move the pointer over the widget title area. The last updated time and icons to manage a widget are displayed on top of a widget. To manage a widget, click the corresponding icon and follow the steps as described in the following table:

Table 63. Manage widget icons and descriptions

Icon	Description
 Refresh	To forcefully update the data in a widget in between of a pre-defined refresh interval time, click refresh icon. Refresh interval can be in range of 30 - 86400 seconds for all widgets.
 Filter	To filter the data in a widget, click filter icon. Set the filter options and click Apply . To set the number of records to display on a widget, type the number in the 'Record limit'. To clear the filters on a widget, click Clear . To close the filter wizard without saving the data, click Cancel . For more information on filter options for the respective widget, refer to Widget filter options .
 Clone	To make a clone of a widget, click the clone icon.
 Configure	To configure the widget title and widget data refresh interval (in seconds); click the configure icon. To save the configured fields on the widget, click Apply . To clear the configured fields on the widget, click Clear . To close the configuration wizard without saving the data, click Cancel .
 Delete	To delete a cloned widget, click delete icon. In the confirmation dialog box, click Remove to remove the widget from the NetVault dashboard. NOTE: You can delete only a cloned widget from the NetVault dashboard.
 Resize	To expand a widget to view the data clearly, click the resize icon. To collapse a widget, click the resize icon.

Adding Widgets

You can add more widgets to the available dashboard. When you add all the widgets to the dashboard, no more widgets can be added.

To add widgets:

- 1 In the Navigation pane, click **Dashboard**.
- 2 In the default dashboard page, click the **ellipsis-horizontal**  icon and click **Add Widgets**.
- 3 Select the required widgets for the dashboard, and click **OK**.

Positioning Widgets

You can move a widget to a certain position on the dashboard by drag and drop action. If you re-position the dashboard and share with multiple users, these users can again re-positioned the shared dashboard. This re-positioning done by any shared user is displayed to rest of the shared users.

To position widgets:

- 1 In the Navigation pane, click **Dashboard**.
- 2 In the dashboard page, move the cursor on the widget name.
- 3 Drag the widget to the required position on the dashboard and drop.

Widget filter options

The following widget filter options are displayed when you click the filter icon on top of a widget. You can use one or more filters to display records that match the specified criteria.

Table 64. Widget filter options

Widget	Filter options and descriptions
Job Status	<p>Client: Use this option to filter jobs for a particular client.</p> <p>Plugin: Use this option to filter jobs performed using a particular plug-in.</p> <p>Job Type: Use this option to filter jobs by job type. You can select backup or restore job types.</p> <p>Group By: Use this option to group jobs by Day, Week, Month, Quarter, or Year.</p> <p>Record limit: Use this option to set the number of records to display on a widget. By default 7 records are displayed.</p> <p>Duration: Use this option to filter jobs that completed at a particular duration.</p> <p>From: To filter jobs from a specific date and time, do the following:</p> <ul style="list-style-type: none"> ▪ Type the start date, or click the button next to the box, and select the start date. ▪ Type the start time, or click the button next to the box, and select the start time. <p>To: To filter jobs up to a specific date and time, do the following:</p> <ul style="list-style-type: none"> ▪ Type the end date, or click the button next to the box, and select the end date. ▪ Type the end time, or click the button next to the box, and select the end time. <p>Select Policies: Use this option to filter the policy jobs by their policy name.</p>
Job Duration	<p>Client: Use this option to filter jobs for a particular client.</p> <p>Plugin: Use this option to filter jobs performed using a particular plug-in.</p> <p>Policy: Use this option to filter the policy jobs by their policy name.</p> <p>Job Type: Use this option to filter jobs by job type. You can select backup or restore job types.</p> <p>Record limit: Use this option to set the number of records to display on a widget. By default 10 records are displayed.</p> <p>Duration: Use this option to filter jobs that completed at a particular duration.</p> <p>From: To filter jobs from a specific date and time, do the following:</p> <ul style="list-style-type: none"> ▪ Type the start date, or click the button next to the box, and select the start date. ▪ Type the start time, or click the button next to the box, and select the start time. <p>To: To filter jobs up to a specific date and time, do the following:</p> <ul style="list-style-type: none"> ▪ Type the end date, or click the button next to the box, and select the end date. ▪ Type the end time, or click the button next to the box, and select the end time.
Media Space Usage	<p>Show space in: Use this option to show the media space in Bytes, KB, MB, GB, TB, or PB.</p> <p>Record limit: Use this option to set the number of records to display on a widget. By default 40 records are displayed.</p> <p>Select media groups: Use this option to filter the records on the basis of media group. By default, all media groups are selected.</p>
Client Status	<p>Client: Use this option to see the client status for a particular client. By default, all the clients are selected.</p>

Table 64. Widget filter options

Widget	Filter options and descriptions
Unusable Media	<p>Select reason: Use this option to set the reason for unusable media such as ASF, Bad, Foreign, Full, Offline, Other, Read Only, or Unusable. By default, all reasons are selected.</p> <p>Record limit: Use this option to set the number of records to display on a widget. By default 100 records are displayed.</p> <p>Select media groups: Use this option to filter the records on the basis of media group. By default, all media groups are selected.</p>
Device Status	<p>Device Type: Use this option to set the type of device as RAS Devices, Tape Drives, or Tape Libraries. By default, all the device types are selected.</p>
Policy Status	<p>Policy Status: Use this option to filter the policies by status as Complete, Failed, or Warning. By default, all statuses are selected.</p> <p>Policy State: Use this option to filter the policies by state as Active, Dormant, Quiesced, or Quiescing. By default, all states are selected.</p>
Event Stream	<p>Select Interval: Use this option to filter the events by time interval as 5 min, 10 min, 30 min, 1 Hr, 8 Hr, 12 Hr, or 24 Hr. By default, the time interval is set to 10 min.</p> <p>Select Event Type: Use this option to filter the events by type as Error Only, Key Events, or All Events. By default, Error Only event type is selected.</p>
Log Stream	<p>Select Interval: Use this option to filter the logs by time interval as 5 min, 10 min, 30 min, 1 Hr, 8 Hr, 12 Hr, or 24 Hr. By default, the time interval is set to 10 min.</p> <p>Select Level: Use this option to filter the logs by level as Background, Error, Information, Job Messages, server, or Warnings. By default, all levels are selected.</p> <p>Select Classes: Use this option to filter the logs by class as Database, devices, Jobs, Media, Plugins, Schedule, System, or UI. By default, all classes are selected.</p>
Running Jobs	<p>Client: Use this option to filter jobs for a particular client.</p> <p>Plugin: Use this option to filter jobs performed using a particular plug-in.</p> <p>Policy: Use this option to filter the policy jobs by the policy name.</p> <p>Job Type: Use this option to filter jobs by job type. You can select backup or restore job types.</p> <p>Record limit: Use this option to set the number of records to display on a widget. By default 10 records are displayed.</p> <p>Duration: Use this option to filter jobs that are completed at a particular duration.</p> <p>From: To filter jobs from a specific date and time:</p> <ul style="list-style-type: none"> • Type the start date, or click the button next to the box, and select the start date. • Type the start time, or click the button next to the box, and select the start time. <p>To: To filter jobs up to a specific date and time:</p> <ul style="list-style-type: none"> • Type the end date, or click the button next to the box, and select the end date. • Type the end time, or click the button next to the box, and select the end time.

i | **NOTE:** To select consecutive items, hold down the Shift key while clicking with the mouse button; to select non-consecutive items, hold down the Ctrl key while clicking with the mouse button. To hide an option in the chart or clear the selected option, click the corresponding option.

Managing jobs

- Viewing job activity
- Viewing job calendar
- Managing jobs
- Managing job definitions
- Viewing job history

Viewing job activity

You can monitor the progress and status of your jobs from the **Job Status** page. You can also use this page to manage your jobs and perform various job-related tasks.

To view job activity:

- 1 In the Navigation pane, click **Job Status**.
- 2 On the **Job Status** page, you can view the following information.

Figure 21. Job Status page

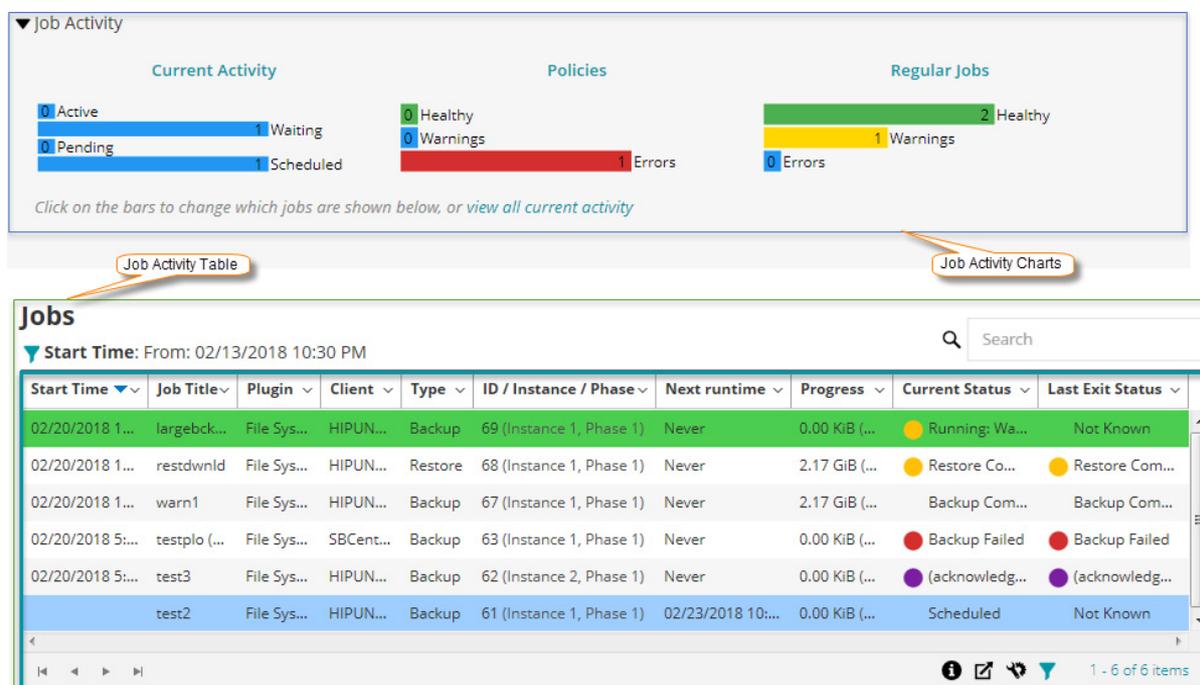


Table 65. Job Status page

Item	Description
Job activity charts	<p>This area shows the summary for current jobs, policy jobs, and regular jobs in the form of bar charts:</p> <ul style="list-style-type: none"> • Current Activity: The individual bars represent the number of jobs that are in active, waiting, pending, and scheduled states. • Policies: The individual bars represent the number of policy jobs that have completed successfully, completed with warnings, and failed. The Healthy count of policies shows the number of policy jobs that have completed successfully when all the jobs in the policy are successful. It shows 0 count if one or more jobs in the policy have warnings or errors. • Regular Jobs: The individual bars represent the number of regular jobs that have completed successfully, completed with warnings, and failed. <p>You can click a job activity to display the job details for that category in the job activity table. For example, you can click Policies in the Job Activity area to display the policy jobs in the job activity table.</p> <p>You can click a bar in the activity chart area to display the job details for that category in the activity table. For example, you can click the Active bar in the Current Activity area to display the jobs that are in progress.</p> <p>To return to the default view type for the activity table, click the View all current activity link.</p> <p>To hide the job activity statistics on Job Status page, follow these steps:</p> <ul style="list-style-type: none"> • In the header pane of NetVault WebUI, click the Settings icon and select Application Settings. Application Settings drawer is displayed on the right-side of the page. • In the Hide Job Statistics option, select Hide the Job Activity Statistics check box. • Click OK.
Job activity table	<p>By default, the table lists all current job activities (All activity category).</p> <p>For the All Activity category, you can view the following information: Start Time (also shows start date), Job Title, Plugin, Client, Type, ID/Instance/Phase (shows the instance and phase), Next runtime (date and time when the next instance is scheduled to run), Progress (current or average transfer rate), Current Status (current status of a job instance), and Last Exit Status (exist status of latest completed job instance).</p> <p>You can also add or remove following additional columns using the column menu: Policy Name, ID (shows only the job ID), Instance (shows only the job instance), Phase (shows only the job phase), Selection Set, Plugin Options Set, Schedule Set, Source Set, Target Set, Advanced Options Set, Duration (Displays the time in hours: minutes: seconds to complete the operation for backup/restore/reporting). For more information, see Sorting records in NetVault WebUI.</p> <p>The following job status light icons are displayed in the current status and last exit status columns based on the exit status:</p> <ul style="list-style-type: none"> • Red light: Aborted, Died, Failed, Did not run • Yellow light: Warnings, Stopped, and Waiting for media • Violet: Acknowledged <p>The following row background colors are displayed for a job based on current status:</p> <ul style="list-style-type: none"> • Green: Running • Blue: Scheduled <p>You can change the view type and also filter the records displayed in the table:</p> <p>Change the view type: To view the job details for a different job category, click the corresponding bar in the activity chart area. For example, to view the jobs that have completed with warnings, click the Warnings bar in the Regular Jobs area.</p>

- 3 By default, the table is sorted by Start Time (descending order).

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To set the filter options (For more information on table filter options for **Job Status** page, see [Table 7](#)), view the page size setting, column sort order, applied filters in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table.

For more information, see [Customizing tables in NetVault WebUI](#).

i | **NOTE:** Acknowledged jobs are not part of regular sorting. So, sorting results display normal jobs first and then acknowledged jobs.

- 4 You can use the **Search** option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string.
- 5 To perform a job-related task, click the corresponding button in the Operations pane if the Page Options are set to show the action options as buttons on bottom of page. By default, this option is displayed.

-OR-

To perform a job-related task using ‘Actions’ column in the Job Activity table, select a job (s) and click the corresponding context menu icon in the Actions column. Click the job-related action you want to apply.

Alternatively, click a link in the Navigation pane to open a different page.

Viewing job calendar

You can view your jobs for a month, a week, or a day on a calendar. It can be helpful to view your jobs in the calendar format to ensure that there are no scheduling conflicts. You may want to verify the calendar before you create a job. You can monitor the progress and status of your jobs from the **Job Calendar** page. You can also use this page to manage your jobs and perform various job-related tasks.

To view backup jobs on a calendar:

- 1 In the Navigation pane, click **Job Calendar**.
- 2 On the **Job Calendar** page, you can view the following information.

Table 66. Job Calendar page

Item	Description
Settings icon	<p>Click this icon to display the job calendar settings. 'Calendar settings' drawer is displayed on the right- side of the page showing the following options:</p> <ul style="list-style-type: none"> • Filter By Status: This area shows the following filter options based on the job status: <ul style="list-style-type: none"> ▪ Scheduled: To display the scheduled jobs, click this button. ▪ Active: To display the in-progress jobs, click this button. ▪ Success: To display the jobs that have completed successfully, click this button. ▪ Warnings: To display the jobs that have completed with warnings, click this button. ▪ Errors: To display the jobs that have failed, click this button. <p>By default, all the jobs are displayed on the Job Calendar page. To hide a job status, click the corresponding button. You can select one or more filters to display the corresponding jobs on job calendar.</p> <ul style="list-style-type: none"> • Timeslot Size: For Weekly and Daily view, you can set the time granularity for the jobs on calendar using Timeslot Size field. <p>To set the time granularity, select any of the following time slot sizes: 1 min, 5 mins, 15 mins, 20 mins, 30 mins, or 60 mins.</p> <p>By default, the time slot size is set to 30 mins.</p> <p>You can select the preferred filters and persist the setting over browser sessions.</p>
Job Calendar	<p>This area displays the jobs on the calendar based on Filter By Status and Timeslot Size settings.</p> <p>You can view your jobs for a month, a week, or a day on calendar. By default, weekly view of the Job Calendar is displayed. To display monthly view, click 'month' or to display the daily view, click 'day'.</p> <p>Monthly view: Displays the jobs for the month. The height of the job is fixed in this view.</p> <p>Weekly view: Displays the jobs for a week. The height of the job is corresponding to the time taken to run the job.</p> <p>Daily view: Displays the jobs for a day. The height of the job is corresponding to the time taken to run the job.</p> <p>The Red color line on the Job Calendar shows the current time.</p> <p>Move the pointer over a job to see the job summary such as job title, start time, end time, and so on.</p>

- 3 To quickly locate the job (s), you can use the Search box at the upper right corner of the page. The job (s) on the calendar is filtered based on Job name or Job ID. The data in the calendar is filtered as you type the search text into the box. You can also include wildcards ("?" or "**") in the search filter string.
- 4 To perform a job-related task, click the job and then click the job-related action you want to apply.

Managing jobs

This section includes the following topics:

- [Running a job immediately](#)

- Aborting a job
- Stopping a job
- Restarting a job
- Placing a job on hold
- Resuming a job
- Determining the reason for “waiting for media” status
- Viewing log messages for a job
- Viewing and managing a job
- Monitoring job progress
- Clearing job errors and warnings
- Removing a job schedule

Running a job immediately

You can use the **Run Now** method run an existing job immediately.

To run a job immediately:

- 1 In the Navigation pane, click **Job Status** or **Manage Job Definitions**.

i | **NOTE:** If the job that you want to run was only saved and not submitted for scheduling, use the **Manage Job Definitions** link.

- 2 In the list of jobs, select the job, and click **Run Now**.
- 3 In the confirmation dialog box, click **OK**.

i | **NOTE:** If you select a Phase 2 job, only that phase is started; phase 1 is not run. If a backup job includes Phase 2 jobs, selecting Phase 1 runs both Phase 1 and Phase 2 jobs. Phase 1 runs immediately, and when it completes successfully, Phase 2 is scheduled to run as per the job definition.

When the job is successfully started, a message is displayed at the upper-right corner of the NetVault WebUI.

Aborting a job

You can cancel an active job from the **Job Status** page.

To abort a job:

- 1 In the Navigation pane, click **Job Status**.
- 2 In the list of jobs, select the job, and click **Abort**.
- 3 In the confirmation dialog box, click **OK**.

When the job is successfully aborted, a message is displayed at the upper-right corner of the NetVault WebUI.

Stopping a job

The Plug-in *for FileSystem* lets you stop a job at any point and resume it later from the same point. To have this ability, you must configure the job as restartable.

For more information about restartable jobs, see the *Quest NetVault Plug-in for FileSystem User's Guide*.

When you stop the job, the plug-in generates an index for all items that have been processed up to that point and writes the index to the backup media and NetVault Database. The job status is then set to **Job Stopped**. If the plug-in is writing a large backup index, the jobs status continues to be reported as "Writing to Media: Storing Backup Index" until the index is written. When you restart the job later, the plug-in runs an Incremental Backup job to back up the remaining files and folders.

i | **NOTE:** The **Stop** and **Restart** methods do not work if you select multiple jobs simultaneously.

To stop a job:

- 1 In the Navigation pane, click **Job Status**.
- 2 In the list of jobs, select the job, and click **Stop**.
- 3 In the confirmation dialog box, click **OK**.

Restarting a job

The **Restart** method lets you resume a backup from the point the job was stopped. To have this ability, you must configure the job as restartable.

The following plug-ins support this feature:

- **Plug-in for FileSystem:** The Plug-in *for FileSystem* allows you to stop a backup job at any point and resume it later from the same point. To have this ability, you must configure the backup option **Enable Restartable Backup** for the job. For more information about this option, see the *Quest NetVault Plug-in for FileSystem User's Guide*. When you stop the job, the plug-in generates an index for all items that have been processed up to that point and sets the job status to **Job Stopped**. When you restart the job later, the plug-in runs an Incremental Backup job to back up the remaining files and folders.
- **Plug-in for VMware:** The Plug-in *for VMware* allows you to restart a job to back up only those virtual machines that failed previously. To have this ability, you must configure the backup option **Enable Restartable backups** for the job. For more information about this option, see the *Quest NetVault Plug-in for VMware User's Guide*. When a restartable backup completes with some failed virtual machines, the plug-in generates an index for the completed virtual machines and sets the job status to **Job Stopped**. When you restart a job, the plug-in runs an Incremental Backup job to back up the failed virtual machines.

i | **NOTE:** The **Restart** method does not work if you select multiple jobs simultaneously.

To restart a job:

- 1 In the Navigation pane, click **Job Status**.
- 2 In the list of jobs, select the stopped job, and click **Restart**.

Placing a job on hold

To prevent a scheduled job from running, you can place the job on hold. When you place a job on hold, its schedule is disabled until you resume the job.

To place a job on hold:

- 1 In the Navigation pane, click **Job Status**.
- 2 In the list of jobs, select the job, and click **Hold Schedule**.
- 3 In the confirmation dialog box, click **OK**.

When the schedule is suspended, the job shows a status of **On hold** in the **Next runtime** column.

i **NOTE:** If you restart NetVault, a job on hold remains in the same state, but it is rescheduled to run at its next scheduled time. If the job cannot be rescheduled because it was scheduled to run once and this time has now elapsed, the job status is set to **Did not Run**. Warning messages are generated in the NetVault Logs that describe why the job did not run.

Resuming a job

When you place a job on hold, its schedule is disabled until you resume the job. You can resume the scheduling of the job from the **Job Status** page.

To resume a job that was placed on hold:

- 1 In the Navigation pane, click **Job Status**.
- 2 In the list of jobs, select the job, and click **Resume Schedule**.
- 3 In the confirmation dialog box, click **OK**.

When the schedule is resumed, the **Next runtime** column is updated to display the date and time when the next instance is scheduled to run.

Determining the reason for “waiting for media” status

When a job is in **Waiting for Media** state, it implies that the job is unable to initiate data transfer as the target drive or media item is unavailable. The **Diagnose Job** method allows you to determine the exact reason for this state.

A job may be in the **Waiting for Media** state due to the following reasons:

- The target media or device is in use by a different job.
- The target device is offline.
- The target media item is not loaded.
- The Reuse Media option is not selected for the job. So, the job is waiting for new media.
- No blank or reusable media item is available for the job.

To determine the exact reason for the “Waiting for Media” status:

- 1 In the Navigation pane, click **Job Status**.
- 2 In the list of jobs, select the job, and click **Diagnose Job**.
- 3 On the **Diagnose Media Request** page, you can view the following information:
 - **Media Request ID:** Media request ID.

- **Request Type:** Backup or restore.
- **Status:** Status of the media request.
- **Priority:** Media request priority.
- **Client:** NetVault Client on which the job is running.
- **Media:** Target media and group label (if specified).
- **Reuse Media:** The “**Reuse Media**” option is selected or not.
- **Required Space:** Amount of space required on the media to complete the current request.
- **Ensure backup is the first on the target media:** The “**Ensure This Backup is First on the Media**” option is selected or not.
- **Auto-Label:** The “**Label Blank Media Automatically**” option is selected or not.
- **Mark Read-Only After:** The “**Mark Media Read-Only**” option is selected or not.
- **Media Format:** Media format.
- **Unload on Completion:** Media to be unloaded after job completion or not.
- **Drives:** Drive on which the tape resides.
- **Force Local:** The “**Local Drives Only**” option is selected or not.
- **Network Compression:** The “**Network Compression**” option is selected or not.

The **Reasons** table lists the reasons why the specified device or media cannot be used for the job. The following list provides some examples:

- Not enough space.
- Currently unavailable.
- The 'force local drives' option is set. This type of device is considered network attached.

i | **NOTE:** We recommend that you perform the steps described in this section and generate a log dump file when logging a case with Quest Technical Support.

Viewing log messages for a job

To view log messages for a job:

- 1 To view log messages for the latest instance, do the following:
 - a In the Navigation pane, click **Job Status**.
 - b In the list of jobs, select the job, and click **View Logs**.
- 2 To view log messages for previous instances, do the following:
 - a In the Navigation pane, click **Manage Job Definitions** or **Job History**.
 - b In the list of available job definitions, select the job, and click **Manage Job**.
 - c On the **Manage Job** page, select the instance, and click **View Logs**.

Viewing and managing a job

The Manage Job page, which is accessible from the Job Status page, lets you view, edit, and remove an existing job, as well as create a job.

To view and manage a job:

- 1 In the Navigation pane, click **Job Status**.
- 2 On the **Job Status** page, select the job you want to view, and then click **Manage Job**.

The **Manage Job** page opens.

- If you selected a backup job, you can edit the backup job on this page.
- If you selected a restore job, click **Edit Job** and make your changes on the Edit Restore Job page.

On the **Manage Job** page, you can view the following information:

- **Job Details:** This area shows the following information: job name, job ID, job type (backup or restore), and name of the client.
- **Configuration:** This area lets you create, edit, and view the protection sets. It shows the selections, the plug-in options, schedule, source storage, target storage, and advanced options. To make changes to any of these details, click the corresponding edit icon.

On the detail drawer for your selection, you can edit and save your changes or you can clone the existing set. If you clone the set, you must enter a new set name in the text box at the bottom of the drawer.

i **NOTE:** You can use the Client links in the Job Details area to modify a set, however you cannot use these links to specify a new set for the job. To modify the job definition, see [Editing job definitions](#).
To edit a backup job, you can understand more about the sets by previewing the set information. To preview the set summary, move the pointer over the information icons in the **Job Details** area on **Manage Job** page.

- **Recent Instances:** This table lists the recent instances of the job. It shows the following information: Run time, duration, job size, instance ID, phase ID, and status (Succeeded, Failed, Aborted, and others).

- 3 To save your changes or submit the job for scheduling, click **Save** or click **Save & Submit**, respectively.

If a new name was provided, the **Save Job** dialog box displays.

- 4 In the Save Job dialog box, select the applicable action from the following options:
 - **Save changes creating a new job using the supplied name:** Select this option to save the changes to a new job.
 - **Save changes renaming this job to the supplied name:** Select this option to rename the existing job.
- 5 Click **OK**.

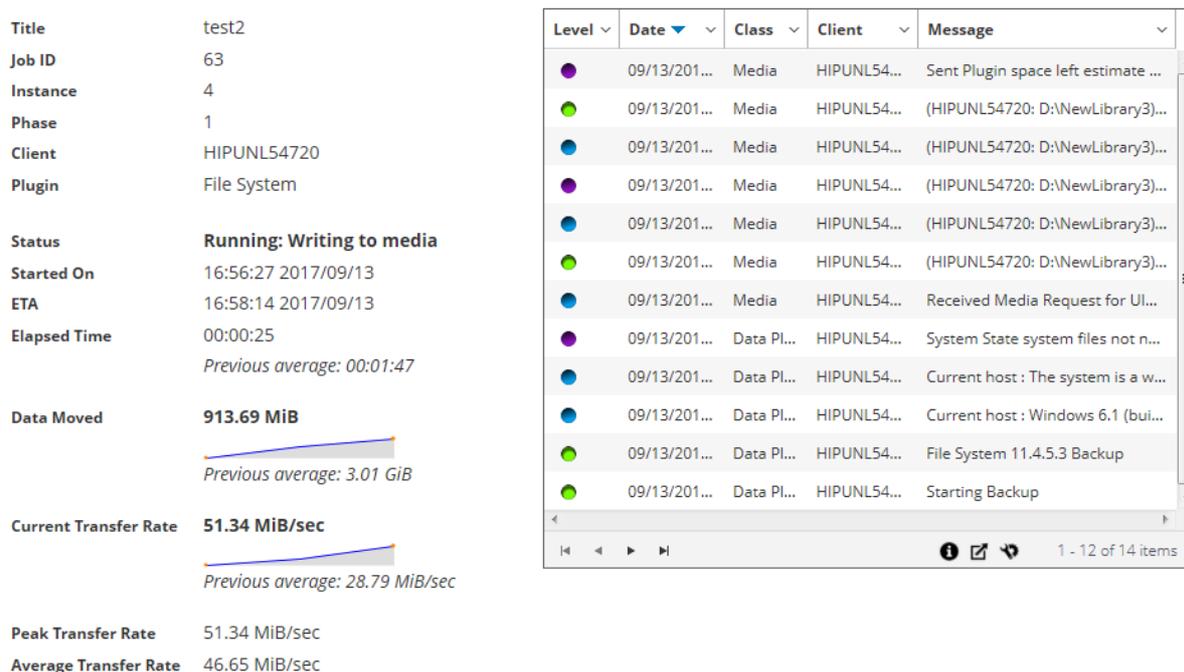
Monitoring job progress

You can monitor the progress of a job from the **Monitor Job** page. The page displays the job status, data transfer rate, log messages, and other job details.

To monitor the progress of a backup or restore job:

- 1 In the Navigation pane, click **Job Status**.
- 2 In the list of jobs, select the job, and click **Monitor**.
- 3 On the **Monitor Job** page, you can view the following information:

Figure 22. Monitor Job page



- **Job details:** This area displays the following details: job title, job ID, instance ID, phase number, client name, plug-in name, status, start time, expected completion time, time elapsed, size of data moved and data transfer chart, current transfer rate and data transfer chart, peak transfer rate, and average transfer rate. The fields: waiting for device, elapsed transfer, estimated transfer remaining, and estimated time remaining are also displayed when running a job using Plug-in *for NDMP*. The fields: waiting for device and elapsed transfer are also displayed when running a job using Plug-in *for SnapMirror to Tape*. However, the field 'expected completion time' is not displayed when running a job using Plug-in *for NDMP* and Plug-in *for SnapMirror to Tape*.
- **Job logs:** This area displays the log messages generated for the job. By default, the table is sorted by date and time. You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower-right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

- 4 To view the job definition or abort the job, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

Clearing job errors and warnings

To clear job errors or warnings:

- 1 In the Navigation pane, click **Job Status**.
- 2 Do one of the following:
 - To clear errors and warnings for all current jobs, set the view to “All activity,” and click **Acknowledge**.
 - To clear warnings for regular or policy jobs, click the “Warnings” bar in the Regular Jobs or Policies category, and click **Acknowledge**.
 - To clear errors for failed regular or policy jobs, click the “Errors” bar in the Regular Jobs or Policies category, and click **Acknowledge**.
 - To clear error or warning for a specific job, select the job in the table, and click **Acknowledge**.

Removing a job schedule

You can remove a job schedule if you do not want to run the job. When you remove the job schedule, the job definition is not deleted. You can resubmit the job definition to schedule it again.

To remove a job schedule:

- 1 In the Navigation pane, click **Job Status**.
- 2 In the list of jobs, select the job, and click **Remove Schedule**.
- 3 In the confirmation dialog box, click **OK**.

When the schedule is successfully deleted, a message is displayed at the upper-right corner of the NetVault WebUI.

Managing job definitions

This section includes the following topics:

- [Viewing job definitions](#)
- [Editing job definitions](#)
- [Deleting job definitions](#)
- [Adding a non-policy job to a policy](#)

Viewing job definitions

Job definitions are stored in the Scheduler Database. You can view the job definitions for all scheduled, saved, completed, and active jobs from the **Manage Job Definitions** page.

To view job definitions:

- 1 In the Navigation pane, click **Manage Job Definitions**.

On the **Manage Job Definitions** page, you can view the job definitions that are available in the NetVault Database. The page includes all scheduled and saved jobs.

Figure 23. Manage Job Definitions page

Manage Job Definitions

▼ All Search

Job Title	ID	Policy name	Type	Plugin	Client	Selection Set	Next Run Time
z	85	z	Backup	File System	LPT-SONALI-P	Test 1	Never
9	84		Backup	File System	LPT-SONALI-P	Test 1	Never
?	83		Backup	File System	LPT-SONALI-P	Test 1	Never
5?	82		Backup	File System	LPT-SONALI-P	Test 1	Never
1-2/3,4	81		Backup	File System	LPT-SONALI-P	Test 1	Never
test csv disk storage	80		Backup	File System	LPT-SONALI-P	Test 1	Never
test IE	79		Backup	Data Copy	LPT-SONALI-P	Test 21	Never
Restore_Saveset 16_5/16...	78		Restore	File System	LPT-SONALI-P	LPT-SONALI-P_Saveset_1...	Never
pq	77	pq	Backup	File System	LPT-SONALI-P	Test 1	Never

1 - 10 of 26 items

+ Add To Policy
Remove
Hold Schedule
Resume Schedule
Run Now
Manage Job

The table shows the following information:

- **Job Title:** Name or title of the job.
- **Job ID:** Job ID number (also shows the instance ID number and phase ID number).
- **Policy name:** Name of the policy if the job belongs to a backup policy.
- **Type:** Backup or restore.
- **Plugin:** Name of the plug-in used to perform the job.
- **Client:** Name of the NetVault Client for which the job was performed.
- **Selection Set:** Backup or Restore Selection Set.
- **Next Run Time:** Date and time when the next instance is scheduled to run.

By default, the table is sorted by Job ID number (descending order).

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

- 2 To set the filter options, view the page size setting, column sort order, or applied filters in the table; export the records; or edit table settings, click the icons at the lower- right corner of the table. For more information on table filter options for **Manage Job Definitions** page, see [Table 10](#).

For more information, see [Customizing tables in NetVault WebUI](#).

- 3 Optionally, use the **Search** option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string.
- 4 Select the job definition that you want to view, and click **Manage Job**.

On the **Manage Job** page, you can view the following information:

- **Job Details:** This area shows the following information: job name, job ID, job type (backup or restore), and name of the client.
- **Configuration:** This area lets you create, edit, and view the protection sets. It shows the selections, the plug-in options, schedule, source storage, target storage, and advanced options.

i **NOTE:** You can use the Client links in the Job Details area to modify a set, however you cannot use these links to specify a new set for the job. To modify the job definition, see [Editing job definitions](#). To edit a backup job, you can understand more about the sets by previewing the set information. To preview the set summary, move the pointer over the information icons in the **Job Details** area on **Manage Job** page.

- **Recent Instances:** This table lists the recent instances of the job. It shows the following information: Run time, duration, job size, instance ID, phase ID, and status (Succeeded, Failed, Aborted, and others).

5 To perform a job-related task, click the corresponding button in the Operations pane.

Alternatively, click a link in the Navigation pane to open a different page. However, the **Remove** button is disabled for a job associated with a policy. You can remove a job associated with a policy only from **Manage Policies** page. Also, if the job is associated with a policy, **Manage Job** page also lets you make changes to the job policy.

Editing job definitions

You can change one or more components (sets) for an existing job definition. When saving the changes, you can choose to overwrite the job or save the changes to a new job.

To edit job definitions:

- 1 In the Navigation pane, click **Manage Job Definitions**.
- 2 In the list of available job definitions, select the job, and click **Manage Job**.
- 3 Depending on the type of job, the backup or restore job wizard is started.
 - For a policy job, **Customize Policy Job** dialog box displayed.
 - To edit a policy job without changing its job type, click **No**. Edit Policy Job wizard is started.
 - To edit a policy job by changing its job type (for example, changing it to a normal backup or restore job), click **Yes**. Backup or restore job wizard is started.
- 4 Select an existing set or configure a new set for the component that you want to change. For more information, see [Creating backup jobs](#), [Creating restore jobs](#), or [Creating a policy](#).
- 5 To rename the job or to save the changes to a new job, type a name in the **Job Name** box.
- 6 Save or schedule the job, as required:
 - **Schedule the job:** To submit the job for scheduling, click **Save & Submit**.
 - **Save the definition without scheduling the job:** To save the job definition without scheduling it, click **Save**.
- 7 If a new name was provided in [Step 5](#), the **Save Job** dialog box is displayed. Select the applicable option in this dialog box:
 - **Save changes creating a new job using the supplied name:** Select this option to save the changes to a new job.
 - **Save changes renaming this job to the supplied name:** Select this option to rename the existing job.
- 8 Click **OK**.

Deleting job definitions

If a job definition is no longer required, you can remove it from the NetVault Database. If a job is active, you cannot delete its definition until you cancel the job. Similarly, if a job is schedule to run later, you cannot delete its definition until you cancel the job schedule.

To delete job definitions:

- 1 In the Navigation pane, click **Manage Job Definitions**.
- 2 In the list of available job definitions, select the jobs that you want to delete.

To select consecutive items, hold down the Shift key while clicking with the mouse button; to select non-consecutive items, hold down the Ctrl key while clicking with the mouse button. However, you can remove a job definition associated with a policy only from **Manage Policies** page.
- 3 Click **Remove**, and in the confirmation dialog box, click **OK**.

Adding a non-policy job to a policy

You can add a non-policy job to an existing policy.

To add a non-policy job to an existing policy:

- 1 In the Navigation pane, click **Manage Job Definitions**.
- 2 In the list of available job definitions, select the non-policy job, and click **Add to Policy**.
- 3 Select a policy on the **Select Policies** dialog box. Click **OK**.

i **NOTE:** Adding a non-policy job configured on a client which is associated to a selected policy adds the same job ID in the policy. However, adding a non-policy job configured on a client which is not associated to a selected policy adds a new job with same job name, job definition, and different job ID in the policy.

The following example describes adding a non-policy job (BackupJob) with Job ID (69) configured on a client (Client-D) which is not associated to a selected policy (Policy1) and includes two backup jobs on three clients (Client-A, Client-B, and Client-C):

- a In the list of available job definitions, select the non-policy job 'BackupJob' and click **Add to Policy**.
- b In the **Select Policies** dialog box, select the policy name 'Policy1'.
- c Click **OK**.
- d New backup jobs with the name 'BackupJob' and Job ID as 70, 71, and 72 for each client in the selected policy 'Policy1' will be added as shown below:

Job Name	Job ID	Client Name	Policy Name
BackupJob	69	Client-D	
BackupJob	70	Client-A	Policy1
BackupJob	71	Client-B	Policy1
BackupJob	72	Client-C	Policy1

Viewing job history

You can view the completed jobs from the **Job History** page. The page shows all the instances that have run for the jobs. When viewing job history, you can use one or more filters to display jobs that match the specified criteria.

To view job history:

- 1 In the Navigation pane, click **Job History**.
- 2 On the **Job History** page, you can view the job history records that are available in the NetVault Database.

i | **NOTE:** You can also access this page from the **Job Status** page. You can select a job on the **Job Status** page, and click **View history** to view all past instances for that job.

The table displays the following default information:

- **Job Title:** Name or title of the job.
- **End Time:** Job completion time.
- **Job ID:** Job ID number (also shows the instance ID number and phase ID number).
- **Job Type:** Backup or restore.
- **Plugin:** Name of the plug-in used to perform the job.
- **Selection Set:** Backup or Restore Selection Set.
- **Client:** Name of the NetVault Client for which the job was performed.
- **Status:** Status (Succeeded, Failed, Aborted, and others).

You can add or remove the following columns in the list using the column menu:

- **Duration:** Displays the time (hours: minutes: seconds) to complete the operation (back up, restore, and reporting)
- **Policy Name:** Displays the name of the backup policy.

The row background color in the table is displayed for a job based on run status. For more information, see [Job activity table](#).

- 3 By default, the table is sorted by End Time (descending order).

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To set the filter options (For more information on table filter options for **Job History** page, see [Table 12](#)), view the page size setting, column sort order, applied filters in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table.

For more information, see [Customizing tables in NetVault WebUI](#).

- 4 You can use the **Search** option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string.
- 5 The **Job History** page loads a maximum of 10,000 records initially. The total number of records fetched is displayed at the lower-right corner of the table.

You can click **Load more** to load the next set of records, if available. Each load action fetches a maximum of 10,000 records. This button is disabled if there are no more records to load.

- 6 To perform a job-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page. If the job is associated with a policy, click **Manage Job** to open the Manage Jobs page and make changes to the policy.

Monitoring logs

- [About NetVault logs](#)
- [Viewing log messages](#)
- [Downloading logs](#)
- [Exporting logs](#)
- [Manually purging the log messages](#)
- [Setting up a user-defined log event](#)
- [Removing a user-defined log event](#)
- [Searching the knowledge base](#)

About NetVault logs

All NetVault processes generate log messages to provide information about the status of various operations and report error conditions. These messages are stored in the NetVault Database. You can view the log messages from the **View Logs** page.

NetVault uses the Logging Daemon to manage and write the log messages to the database. The Logging Daemon runs on the NetVault Server.

Log messages contain information that can be used for diagnosing and troubleshooting problems. When you report a problem, Quest Technical Support might request you to provide the log dump file. You can use the **Download** or **Export** method available on the **View Logs** page to generate the log dump file. For more information, see [Downloading logs](#) or [Exporting logs](#).

- i** **NOTE:** NetVault also provides diagnostic tracing capabilities to collect detailed information about error conditions. This information is used for debugging and troubleshooting purposes.
- Unlike logging, tracing is disabled by default. When investigating a problem, Quest Technical Support may request you to enable tracing on the server or client machines.
- Diagnostic information is written to trace files. Each process generates its own trace file. By default, trace files are stored in the **trace-container** directory under the NetVault installation directory. For more information, see [Diagnostic tracing](#).

Log messages that are older than 30 days are automatically purged from the database. You can use the configuration settings for Logging Daemon to change the maximum age for log messages. You can also manually delete the log messages by using the **Purge Logs** method available on the **View Logs** page.

Viewing log messages

You can view the log messages from the **View Logs** page. From this page, you can perform various log-related tasks, like download or export log messages, set log events, or purge log messages.

To view the log messages:

- 1 In the Navigation pane, click **View Logs**.
- 2 On the **View Logs** page, you can view the log messages generated by various processes.

Figure 24. View Logs page

View Logs

▼ **Display Level:** Job Messages and above

Display Level: Job Messages and above [✓] [✕]

Severity	Date	Job ID	Job Instance	Class	Client	Message
	08/13/2018 16:24...	65	1	Jobs	HIPUNP57186	Finished job 65, p...
	08/13/2018 16:24...	65	1	Jobs	HIPUNP57186	Job Status: Restor...
	08/13/2018 16:24...	65	1	Data Plugin	HIPUNP57186	Job completed
	08/13/2018 16:24...	65	1	Media	HIPUNP57186	(HIPUNP57186: \...
	08/13/2018 16:24...	65	1	Media	HIPUNP57186	(HIPUNP57186: \...
	08/13/2018 16:24...	65	1	Media	HIPUNP57186	(HIPUNP57186: \...
	08/13/2018 16:24...	65	1	Data Plugin	HIPUNP57186	Building list of ite...
	08/13/2018 16:24...	65	1	Data Plugin	HIPUNP57186	Analyzing items t...

1 - 8 of 39 items

Load more ...

-

The table displays the following information:

- **Severity:** Severity level of the message.
Log messages are classified into the following severity levels: Background messages, Information messages, Job messages, Warning messages, Error messages, and Severe error messages.
- **Date:** Date and time when the log was generated.
- **Job ID:** Job identification number.
- **Class:** Type of operation that generated the logs.
Log classes include the following: System, Schedule, Jobs, Media, Database, Plugins, and UI.
- **Client:** Name of the client for which the log was generated.
- **Message:** Detailed log message or description.

The following table provides a brief description of the log severity levels.

Table 67. Log severity levels

Icon	Severity level	Description
	Background	General log messages.
	Information	Log messages related to media, scheduler, and system activities.
	Job message	Log messages related to backup, restore, and report jobs.
	Warning	Problems that might not have caused a job to fail.
	Error	Problems that might have caused a job to fail.
	Severe error	Critical problems that might have caused an operation to fail.

i **NOTE:** The icons for some log messages can contain an exclamation mark (for example, , , and others). This mark indicates that you can open the message to view the log context information. To view the context information, click anywhere in the **Severity** column. Alternatively, select the message, and click **More Info**. The dialog box that appears can include data transfer details, execution scripts, or other information. After reviewing the details, click **OK** to close the dialog box.

Setting the 'Display Level' to a particular severity level shows log messages for the selected severity and higher.

- 3 By default, the table is sorted by Date (newest to oldest).

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To set the filter options (For more information on table filter options for **View Logs** page, see [Table 8](#)), view the page size setting, column sort order, applied filters in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table.

For more information, see [Customizing tables in NetVault WebUI](#).

- 4 You can use the Search option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string. Type the search text into the box and click the tick icon or press Enter.

- 5 You can use the **Display Level** option to filter the messages based on their severity level. When you specify the severity level, the messages at that level and higher are displayed on the **View Logs** page.

By default, **Display Level** is set to Job Messages. With this setting, you can view Job Messages, Warnings, Errors, and Severe errors on the page. To specify the minimum severity level for messages that are displayed, select the level in the **Display Level** list.

- 6 The **View Logs** page loads a maximum of 10,000 records initially. The total number of records fetched is displayed at the lower-right corner of the table.

You can click **Load more** to load the next set of records, if available. Each load action fetches a maximum of 10,000 records. This button is disabled if there are no more records to load.

- 7 To stop or resume live updates, click the Pause or Resume Live Updates button.

- 8 To perform a log-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

Downloading logs

When investigating a reported problem, Quest Technical Support might request you to provide the log dump from the NetVault Server. You can generate the dump file by using the **Download** or **Export** method available on the **View Logs** page.

The **Download** method creates dump files on the local machine where the WebUI is running. This method is useful if you are accessing the server from a remote machine. The **Export** method creates the file on the NetVault Server. For more information about this method, see [Exporting logs](#).

The supported dump formats are text, binary, and database table dump. By default, the dump file includes all current logs. You can use the filter options to download log entries that meet the specified criteria. For instance, you can set filter options to display log entries for a specific period or a specific job ID number, and then download these entries to a dump file.

To download logs:

- 1 In the Navigation pane, click **View Logs**.
- 2 Click **Download**, and in the **Confirm Download Logs** dialog box, select the dump format:

- **Binary Log:** Select this option to export the logs to a binary file (“**.nlg**”).
- **Text Log:** Select this option to export the logs to a text file.
- **Database Table Dump:** Select this option to export the logs to a PostgreSQL dump (pg_dump) file.

i | **NOTE:** When you want to provide the log dump file to Quest Technical Support, use the Binary Log (.nlg) format.

- 3 Click **Download**, and in the confirmation dialog box, click **OK**.

In the File Download dialog box, select the Save option, and specify the location, if necessary. Make sure that you do not navigate away from the page until the browser starts downloading the file.

After the browser finishes downloading the file, you can find a compressed file in the download location.

Exporting logs

When investigating a reported problem, Quest Technical Support might request you to provide the log dump from the NetVault Server. You can generate the dump file by using the **Download** or **Export** method available on the **View Logs** page.

The **Download** method creates dump files on the local machine where the WebUI is running. This method is useful if you are accessing the server from a remote machine. For more information about this method, see [Downloading logs](#). The **Export** method creates the file on the NetVault Server.

The supported dump formats are text, binary, and database table dump. By default, the dump file includes all current logs. You can use the filter options to export log entries that meet the specified criteria. For instance, you can set filter options to display log entries for a specific period or a specific job ID number, and then export these entries to a dump file.

To export logs:

- 1 In the Navigation pane, click **View Logs**.
- 2 On the **View Logs** page, click **Export**.
- 3 In the **Export Logs** dialog box, configure the following options.

Table 68. Export logs

Option	Description
File Name	Type a filename for the log file. You can also select an existing file from the dump file list. If you select or specify an existing file, NetVault overwrites the file. Depending on the dump format selected, the dump files are created in the binary , text , or pgdump directories under <NetVault home>/logs/dumps (on Windows) or <NetVault home>/logs/dumps (on Linux). To create the dump file in a different location, specify the full path.
Binary Log	Select this option to export the logs to a binary file (.nlg).
Text Log	Select this option to export the logs to a text file (.txt).
Database Table Dump	Select this option to export the logs to a PostgreSQL dump (pg_dump) file (.dmp). NOTE: When you want to provide the log dump file to Quest Technical Support, use the Database Table Dump format.

- 4 Click **Export** to export the logs.

After the logs are successfully exported, a message is displayed at the upper-right corner of the NetVault WebUI.

Manually purging the log messages

Log messages that are older than 30 days are automatically purged from the database. You can also manually delete the log messages by using the **Purge Logs** method available on the **View Logs** page.

! | **NOTE:** To change the maximum age for log messages, see [Modifying the purge policy for log messages](#).

To manually purge the log messages:

- 1 In the Navigation pane, click **View Logs**.
- 2 On the **View Logs** page, click **Purge**.
- 3 In the **Purge Logs** dialog box, configure the following option:
 - **Delete log entries before:** Type or select the date and time to delete all log messages created before the specified time.
- 4 Click **Purge**.

After the logs are successfully deleted, a message is displayed at the upper-right corner of the NetVault WebUI.

Setting up a user-defined log event

To receive notifications when a log message is generated, you can set up a user-defined log event for that message. You can configure user-defined log events from the **View Logs** page. The user-defined log events are added to the **Log Daemon** event class.

To set up a log event:

- 1 In the Navigation pane, click **View Logs**.
- 2 In the logs table, select the applicable log message, and click **Set Event**.
- 3 In the **Set Log Event** dialog box, configure the following options.

Table 69. Set log event

Option	Description
Event Name	Specify a name for the log event.
Event Description	Provide a detailed description for the event.

- 4 Click **Set Event**.
After the event is successfully added, a message is displayed at the upper-right corner of the NetVault WebUI.
- 5 To receive a notification when the event occurs, set up a notification method for the event. For more information, see the [Setting up a global notification method](#).

Removing a user-defined log event

You can delete a user-defined log event if it is no longer required. This task can be performed from the **View Logs** page.

To remove a log event:

- 1 In the Navigation pane, click **View Logs**.
- 2 In the logs table, select the log message for which the event was set, and click **Set Event**.
- 3 In the **Set Log Event** dialog box, click **Remove**.

Searching the knowledge base

To know more about the errors or warning messages of the jobs and to get the solution, you can search the NetVault knowledge base articles.

To search the knowledge base:

- 1 In the Navigation pane, click **View Logs**.
- 2 In the logs table, select the log message with error or warning, and click **KB Search**.
- 3 In the **KB Search Results** dialog box, view the list of articles related to the job error or message.
- 4 For brief information of the article, click the add button to the left of the article.
- 5 Click **OK** to close the dialog box.

Managing storage devices

- [Monitoring device activity](#)
- [Managing disk-based storage devices in list view](#)
- [Managing disk-based storage devices in tree view](#)
- [Managing the Snapshot Array Manager](#)
- [Managing tape libraries in list view](#)
- [Managing tape libraries in tree view](#)
- [Managing tape drives in list view](#)
- [Managing tape drives in tree view](#)
- [Adding shared devices](#)

Monitoring device activity

You can use the **Device Activity** page to monitor data flows and data transfer rates for devices that are in use.

To view device activity:

- 1 In the Navigation pane, click **Device Activity**.
- 2 On the **Device Activity** page, you can view the following information.

Figure 25. Device Activity page



Table 70. Device Activity page

Item	Description
Device activity	<p>This area shows data transfers from clients to jobs on the left and jobs to storage devices on the right. Gray lines depict data flow; the thicker the line, the higher the rate of flow. Boxes represent clients, jobs, and devices; the taller the box, the higher the rate of flow.</p> <p>You can use the following settings available in the Detail area to modify the data flow view:</p> <ul style="list-style-type: none"> • Devices: Select this option to view data transfers from all jobs to a storage device. • Job groups: Select this option to view data transfers from a job group (plug-in or policy) to a storage device. The job group can be selected in the Grouping list. • Job: Select this option to view data transfers from client to jobs and from jobs to storage devices. • Focus and Un-focus: Use these buttons to show and hide data flow details for a specific device. Click the device box, and click Focus. To hide the details for that device, click Un-focus. • Pause and Play: Use these buttons to pause and resume data flow updates. To stop data flow updates, click Pause, and to resume updates, click Play. <p>The transfer rate shows “unknown” for backups running on older NetVault Clients.</p>
Summary	<p>This area shows the following information: total active transfers, active clients, active devices, current overall throughput, hottest transfer, and hottest target.</p> <p>You can click a box representing the client, job, or storage device in the Device Activity area to view information about that component.</p>

- 3 To perform a job-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

On the **Manage Devices** page, you can view the list of storage devices as a **List View** or **Tree View**. Select **Remember this selection** to save your choice. By default, **List View** is selected.

Managing disk-based storage devices in list view

This section includes the following topics:

- [Viewing disk device details](#)
- [Managing a Quest DR Series system, QoreStor, or Data Domain system](#)
- [Checking the status of a disk-based storage device](#)
- [Changing the status of a disk-based storage device](#)
- [Scanning a disk-based storage device](#)
- [Removing all savesets from a disk-based storage device](#)
- [Removing a disk-based storage device](#)

Viewing disk device details

To view disk device details:

- 1 In the Navigation pane, click **Manage Devices**.

You can view the list of storage devices added to the NetVault Server. The device status is indicated using the following icons.

Table 71. Device status icons

Icon	Description
	Device is online and available for use.
	Device is offline. NetVault is able to detect the device, but unable to access it for backup or restore jobs.
	Device is unavailable. NetVault is unable to detect the device.

- 2 To view the details of a particular device, click the corresponding **Manage Device** icon .

On the **<Type> Device Management** page, you can view the following information:

- **Device details:** The **Device details** area shows the following information:

DR Series systems and Data Domain systems:

- **Name:** The name of the storage device. The name is derived from the container name and the DNS name or IP address of the device.
- **Status:** The status of the device. AVAILABLE specifies that the device is available for backups and restores, while OFFLINE specifies that the device is unavailable and cannot be used for backups or restores.
- **Data Stored:** The total amount of data stored on the device.
- **Space Used:** The total space used by the NetVault Servers to which this device has been added.
- **Space Available:** The total disk space available on the storage device.

- **Deduplication ratio:** The Deduplication Ratio is calculated as follows:

$$\text{Deduplication ratio} = \frac{\text{Data Deduplicated}}{\text{Disk Used by Deduplicated Data}}$$

The **Deduplication ratio** is not displayed for DR Series systems.

NetVault SmartDisk devices:

- **Name:** The name of the storage device. The name is derived from the host name and IP address of the device.
- **Status:** The status of the device. AVAILABLE specifies that the device is available for backups and restores, while OFFLINE specifies that the device is unavailable and cannot be used for backups or restores.
- **Data Stored:** The total amount of data stored on the device.
- **Space Used:** The total space used by the NetVault Servers to which this device has been added.
- **Data Deduplicated:** The total amount of data that has been submitted for deduplication.
- **Data In Staging:** The amount of data stored in the Staging Store.
- **Space Available:** The total disk space available on the storage device.

This option shows the total disk space available across all configured Storage Volumes regardless of their deny/favor configuration. It does not include the Garbage Collection Reserve or Last Resort Threshold (LRT).

- **Deduplication ratio:** The Deduplication Ratio is calculated as follows:

$$\text{Deduplication ratio} = \frac{\text{Data Deduplicated}}{\text{Disk Used by Deduplicated Data}}$$

- **Staging Status:** The status of the Staging Store. It can be one of the following:
 - AVAILABLE: Staging Store is available for writing.
 - FULL: Staging Store is full; no more space is available for writing.
 - UNAVAILABLE: Disk index is unavailable.
 - UNLICENCED: NetVault SmartDisk is not licensed, the license has expired, or the license limit has exceeded.
 - NONE: The status is unknown.
- **Storage Status:** The status of the Storage space. It can be one of the following:
 - AVAILABLE: Storage is available for deduplication and rehydration.
 - FULL: Storage is full; no more space is available for data deduplication.
 - UNAVAILABLE: Chunk Index is unavailable; no data deduplication or rehydration is available.
 - UNLICENCED: NetVault SmartDisk is not licensed, the license has expired, or the license limit has exceeded.
 - NONE: The status is unknown.
- **Group:** The media group name. “None” indicates that the device is not added to any group. To add the device to a storage group, click the Edit icon, and in the **Edit Media Group** dialog box, specify the media group name. Click **Save** to save the details and close the dialog box.
- **Version:** The version number of the device.
- **License Type:** License type being used.
- **License Capacity:** License capacity.
- **License Expiration:** License validity period.

- **Dedupe Licensed:** If the deduplication option is licensed.
 - **Garbage Collection State:** The current phase of Garbage Collection.
 - **Deduplication Queue Length:** The number of elements or NetVault segments currently waiting to be deduplicated.
 - **Activity chart:** This area shows the activity graph if the device is in use by a backup or restore job.
 - **Job details:** This area shows the following information about the active jobs:
 - Data transfer rate.
 - Name or title of the job, job ID number, instance ID number, and phase ID number (1 or 2).
- 3 To perform a device-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

Managing a Quest DR Series system, QoreStor, or Data Domain system

Use the following guidelines to manage an appliance or software-defined storage device.

- 1 In the Navigation pane, click **Manage Devices**.

You can view the list of devices added to the NetVault Server. The device status is indicated using the following icons.

Table 72. Device status icons

Icon	Description
	Device is online and can be managed.
	Device is in maintenance mode and cannot be managed.
	Device is offline and cannot be managed.

- 2 On the **Manage Devices** page, click the **Manage Device** icon  for the device that you want to manage.

The name of the page that appears depends on the type of device that you select. For example, if you select a Quest DR Series system, the **Quest DR Device** page appears; if you select a QoreStor, the **Quest QoreStor Device** page appears.

Additional information is described in the following topics:

- [Viewing device details](#)
- [Viewing cleaner statistics](#)
- [Starting or stopping cleaner](#)
- [Managing device users](#)
- [Managing a container](#)
- [Managing storage groups](#)
- [Removing Quest DR Series system](#)
- [Removing QoreStor](#)

Viewing device details

On the **<type> Device** page, you can view the following information in the collapsible and expandable **Storage Utilization** section:

- **Device Machine:** Machine name of the device.
- **Device User:** User name of the device.
- **Device OS Version:** Device operation system version number.
- **Total Capacity:** Total storage capacity of the device.
- **Used Space:** The total disk space used by the storage device.
- **Available Space:** The total disk space available on the storage device.
- **API Version:** API version number of the RDA library.
- **Cleaner Status:** The cleaning status of the device. It can be one of the following:
 - **Done:** The cleaning of device is complete.
 - **Pending:** The cleaning of device is pending.
 - **Running:** The cleaning of device is in progress.
 - **Idle:** The cleaning of device is not started.
- **Cleaner Statistics:** To view the cleaner statistics, click [Viewing cleaner statistics](#).

Viewing cleaner statistics

- 1 On the **<type> Device** page, click **View Cleaner Statistics**.

On the **Cleaner Statistics** page, you can view the current and previous run statistics. The following information is displayed: File Processed, Bytes Processed, Bytes Reclaimed, Completion Time, Start Time, and End Time.

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

Starting or stopping cleaner

You can start the cleaning process only when the Cleaner Status of the device is Pending or Done and stop the cleaning process only when the Cleaner Status of the device is Running.

To start the cleaner:

- 1 On the **<type> Device** page, click **Start Cleaner**.
- 2 In the **Start Cleaner** dialog box, click **Start Cleaner**.

To stop the cleaner:

- 1 On the **<type> Device** page, click **Stop Cleaner**.
- 2 In the **Stop Cleaner** dialog box, click **Stop Cleaner**.

The Cleaner Status of the device is updated accordingly.

Managing device users

- 1 On the **<type> Device** page, click **Manage Users**.

The **Manage <type> Device Users** page displays the following information:

- **Device Name:** Name of the device.

- **Actions:** Options for user account-related tasks.
- **User Name:** User name of the user.

By default, the table is sorted by User Name (descending order).

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower-right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

You can use the **Search** option to filter the table data and find entries that contain the specified string. You can also include wildcards (“?” or “*”) in the search filter string.

- 2 Select the action that you want to perform.

You can create and manage the various users from the **<type> Device** page and configure the options described in the following sections:

- [Creating a device-related user account](#)
- [Setting or updating the user’s password](#)
- [Deleting a device-based user account](#)

Creating a device-related user account

- 1 On the **Manage <type> Device Users** page, click **Add User**.
- 2 In the **Add User** dialog box, configure the following options:

Table 73. Add User

Item	Description
User Name	Type the user name for the user account. The name must start with an alphabetical character and be between 6 to 31 character long.
New Password	Type a new password for the user account. A password can contain 8 to 16 characters.
Re-enter password	Retype the password for confirmation.

- 3 To save the details, click **Save**.

Setting or updating the user’s password

- 1 On the **Manage <type> Device Users** page, in the list of device-based users, locate the applicable user.
- 2 Click the corresponding  icon in the **Actions** column, and select **Change Password**.
- 3 In the **Change Password** dialog box, configure the following options.

Table 74. User password

Item	Description
Current Password	Type the current password for the user account.
New Password	Type a new password for the user account. A password can contain 8 to 16 characters.
Re-enter password	Retype the password for confirmation.

- 4 Click **Change Password** to save the details, and return to the **Manage <type> Device Users** page.

Deleting a device-based user account

- 1 On the **Manage <type> Device Users** page, in the list of device-based users, locate the applicable user.
- 2 Click the corresponding  icon in the **Actions** column, and select **Delete**.

- 3 In the **Delete User** dialog box, click **Delete**.

Managing a container

On the **<type> Device** page, you can view the list of containers or storage groups for the device. To save your choice, select **Remember this selection**. By default, **Containers** is selected.

For each container, the **<type> Device** page displays Storage Utilization information as well as the following information:

- **Actions:** Options for container-related tasks.
- **Container Name:** Name of the container.
- **Storage Group Name:** Name of the group to which the container is assigned.
- **Attached To NetVault:** **Yes** or **No** (If the container is attached to the NetVault or not).

By default, the table is sorted by Container Name (descending order).

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower-right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

You can use the **Search** option to filter the table data and quickly find entries that contain the specified string. You can also include wildcards ("?" or "**") in the search filter string.

You can create and manage a container from the **<type> Device** page and configure the options described in the following sections:

- [Creating a container](#)
- [Deleting a container](#)
- [Adding a container as a media](#)
- [Removing media](#)
- [Viewing container statistics](#)
- [Modifying DR Series system container settings](#)
- [Modifying Data Domain system settings](#)

Creating a container

- 1 On the **<type> Device** page, with the **Containers** option selected, click **Create Container**.
- 2 In the **Create Container** dialog box, select the storage group, and type the name of the container.
The container name must start with an alphabetical character.
- 3 Click **Save**.

Deleting a container

- 1 On the **<type> Device** page, locate the applicable container.
- 2 Click the corresponding  icon in the **Actions** column, and select **Delete**.
- 3 In the **Delete Container** dialog box, click **Delete**.

i | **NOTE:** Container deletion might take some time even after getting the message for successful deletion.

Adding a container as a media

You can add a container as a media only if the status in the field **Attached To NetVault** is **No**.

- 1 On the **<type> Device** page, locate the applicable container.

- 2 Click the corresponding  icon in the **Actions** column, and select **Add As A Media**.
- 3 On the **Add As A Media** dialog box, configure the following options:

Table 75. Add container as a media

Option	Description
Block size	The default block size is 512 KiB. The block size cannot be changed for Quest DR Series systems.
Stream Limit	The default value for this setting is 256 streams. This setting applies to all NetVault Servers to which the container is added. If the number of data streams exceeds the defined limit for the container, the Media Manager reports an error (“Device has too many streams”). You can set the soft stream limit to any value between 1 and 512. If the container is added to more than one NetVault Server, set the same soft stream limit on all servers.
Force Add	If the device is already added to another NetVault Server with the same name, select the Force Add check box. This option can be useful if you have performed a disaster recovery to rebuild the NetVault Server.

- 4 Click **Add As A Media**.

Removing media

You can remove a container as a media only if the status in the field **Attached To NetVault** is **Yes**.

- 1 On the **<type> Device** page, locate the applicable container.
- 2 Click the corresponding  icon in the **Actions** column, and select **Remove Media**.
- 3 In the **Remove Device** dialog box, click **Remove**.
- 4 If NetVault fails to remove the device, select the **Force Removal** check box in the confirmation dialog, and click **Remove**.

i | **NOTE:** You can use the **Force Removal** option to remove a device that is not in use. However, the device may still try to communicate with the NetVault Server.

Viewing container statistics

- 1 On the **<type> Device** page, locate the applicable container.
- 2 Click the corresponding  icon in the **Actions** column, and select **Statistics**.

In the **Statistics** dialog box, you can view the following information: Files Replicated Inbound, Bytes Processed Inbound, Bytes Replicated Inbound, Bytes Transferred Inbound, Errors Inbound, Files Replicated Outbound, Bytes Processed Outbound, Bytes Replicated Outbound, Bytes Transferred Outbound, Outbound Errors, Bytes Synthesized, Current Bytes, Bytes Transferred, Files Ingested, Files Aborted, Ingest Errors, Files Read, Bytes Read, and Read Errors.
- 3 To close the dialog box, click **OK**.

Modifying DR Series system container settings

You can modify the user password or stream limit for an existing container in the DR Series system.

- 1 In the Navigation pane, click **Manage Devices**.
 - 2 In the list of devices, click the icon  to locate the DR Device and to view the attached containers.
 - 3 Click the corresponding **Manage Device** icon .
- The **DR Series Device Management** page is displayed.

- 4 Click **Update**, and in the **Update Device Details** dialog box, modify the required settings.

Table 76. Modify DR Series system settings

Option	Description
Username	Specify a user account that can be used to log on to the device. Any user with RDA permission on the DR Series system can be used as username. NOTE: User management operations on DR Series system through NetVault can only be performed by backup_user .
Password	Type the password for the user account.
Stream Limit	The default value for this setting is 256 streams. This setting applies to all NetVault Servers to which the container is added. If the number of data streams exceeds the defined limit for the container, the Media Manager reports an error (“Device has too many streams”). You can set the soft stream limit to any value between 1 and 512. If the container is added to more than one NetVault Server, set the same soft stream limit on all servers.

- 5 Click **Update** to save the settings.

Modifying Data Domain system settings

You can modify the username, password, stream limit, or block size for an existing Data Domain system.

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate the device, and click the corresponding **Manage Device** icon .
- 3 Click **Update**, and in the **Update Device Details** dialog box, modify the required settings.

Table 77. Modify Data Domain system settings

Option	Description
Username	Specify a DD Boost user account that can be used to log on to the device for backups and restores. Verify that the user account is created on the Data Domain system before you add the device to the NetVault Server.
Password	Type the password for the user account.
Stream Limit	The default value for this setting is 32 streams. This setting applies to all NetVault Servers to which the container is added. If the number of data streams exceeds the defined limit for the container, the Media Manager reports an error (“Device has too many streams”). You can set the soft stream limit to any value between 1 and 256. If the container is added to more than one NetVault Server, set the same soft stream limit on all servers.
Block size	Type or select the block size for data transfers. The block size is specified in KiB. The default block size is 128 KiB.

- 4 Click **Update** to save the settings.

Modifying QoreStor system settings

You can modify the username, password, stream limit, or block size for an existing QoreStor system.

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate the device, and click the corresponding **Manage Device** icon .

- 3 Click **Update**, and in the **Update Device Details** dialog box, modify the required settings.

Table 78. Modify QoreStor system settings

Option	Description
Username	Specify a user account that can be used to log on to the device. Any user with RDA permission on the QoreStor system can be used as username. NOTE: User management operations on QoreStor system through NetVault can only be performed by backup_user.
Password	Type the password for the user account.
Stream Limit	The default value for this setting is 64 streams. This setting applies to all NetVault Servers to which the container is added. If the number of data streams exceeds the defined limit for the container, the Media Manager reports an error (“Device has too many streams”). You can set the soft stream limit to any value between 1 and 256. NOTE: You must consider target device capability before changing stream limit value for QoreStor. If the container is added to more than one NetVault Server, set the same soft stream limit on all servers.

- 4 Click **Update** to save the settings.

Managing storage groups

- 1 In the Navigation pane, click **Manage Devices**.
- 2 On the **Manage Devices** page, click the **Manage Device** icon  for the device whose storage groups you want to manage.
- 3 On the **<type> Device** page, select the **Storage Groups** option.

i | **TIP:** On the **<type> Device** page, you can view the list of containers or storage groups for the device. To save your choice, select **Remember this selection**. By default, **Containers** is selected.

The **<type> Device** page displays the following information:

- **Actions:** Options for the tasks related to the storage group.
- **Storage Group Name:** Name of the group.
- **Encryption Type:** Type of encryption (None, Static, Internal).
- **Compression Type:** Type of compression (Fast, Best).
- **Rotation Period:** Number of days for key rotation.
- **Container Count:** Number of containers in the storage device.
- **Dedupe Savings:** Percentage of deduplication savings per storage group.

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower-right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

You can use the Search option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “**”) in the search filter string.

- 4 Select the action that you want to perform.

You can create and manage a DR Series system and QoreStor storage groups from the **Manage <type> Device Storage Groups** page and configure the options described in the following sections:

- [Creating a storage group](#)
- [Modifying a storage group](#)
- [Viewing storage group statistics](#)
- [Deleting a storage group](#)

Creating a storage group

- 1 On the **<type> Device** page, with the **Storage Groups** option selected, click **Create Storage Group**.
- 2 On the **Create Storage Group** dialog box, configure the following options:

Table 79. Add storage group

Item	Description
Storage Group Name	Type the name for the storage group. The name must start with an alphabetical character.
Compression Type	Select one of the following compression levels for storage optimization: <ul style="list-style-type: none"> • Fast: Results in shorter backup time with less storage space saving. • Best: Results in longer backup time with more storage space saving.
Encryption Type	Select any of the following encryption types: <ul style="list-style-type: none"> • None: Default value. • Static: A global, fixed key is used to encrypt all data. • Internal: Content-encryption keys are generated and rotated at a specified time.
Passphrase	This field is enabled only when you select Internal and Static in the Encryption Type field. To encrypt content-encryption keys, type a passphrase. The passphrase string can contain a maximum of 255 characters and can contain alphanumeric and special characters. NOTE: The passphrase is mandatory for enabling encryption. If the passphrase is compromised or lost, you must change the passphrase so that the content-encryption keys do not become vulnerable.
Rotation Period	This field is enabled only when you select Internal in the Encryption Type field. Select the number of days for key rotation. The default value is 30 days. Generate the new key after the rotation period is expired. The valid range for this field is 7 to 25550 days.

- 3 To add the new storage group to the device, click **Save**.

Modifying a storage group

- 1 On the **<type> Device** page, locate the applicable storage group, click the corresponding  icon in the **Actions** column, and select **Edit**.
- 2 In the **Edit Storage Group** dialog box, update the options described in [Creating a storage group](#).
- 3 Click **Save**.

Viewing storage group statistics

- 1 On the **<type> Device** page, in the list of storage groups, locate the storage group, click the corresponding  icon in the **Actions** column, and select **Statistics**.

On the **Statistics** dialog box, you can view the following information: Physical Used, Current Bytes, Num Files, Bytes Post Dedupe, Bytes Post Compression, Bytes Post Encryption, Number of Inodes, Cleaner

Status, Encryption Status, Read Throughput, Write Throughput, Dedupe Savings, and Compression Savings.

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

Deleting a storage group

- 1 On the **<type> Device** page, in the list of storage groups, locate the storage group, click the corresponding  icon in the **Actions** column, and select **Delete**.
- 2 In the **Delete Storage Group** dialog box, click **Remove**.

i | **NOTE:** Storage group deletion might take some time even after getting the message for successful deletion.

Removing Quest DR Series system

To remove a Quest DR Device:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 To remove a particular DR device, click the corresponding **Remove** icon .
- 3 On the **Remove Quest DR Device** dialog box, click **Remove**.

Removing QoreStor

For information on removing QoreStor, see the *Quest QoreStor Installation Guide*.

Checking the status of a disk-based storage device

To check the status of an offline device:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate the device, and click the corresponding **Manage Device** icon .
- 3 Click **Check**, and then in the confirmation dialog box, click **Check** again.
If the device is operational, its status is changed to **“Available.”**

Changing the status of a disk-based storage device

To change the status to online or offline:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate the device, and then click the corresponding **Manage Device** icon .
- 3 If the device is offline, click **Online** to bring it back online.
- 4 If the device is online, click **Offline** to take it offline.

The **Offline** method marks the device as offline and makes the device unavailable to NetVault. This method does not physically take the system offline.

Scanning a disk-based storage device

You can use the **Scan** method to query all backups stored on a disk-based storage device and import those backups that are not indexed in the given NetVault Server's database.

To import backups to the database, the NetVault Server must have the same NetVault Machine Name as the original server that performed the backups. The amount of time it takes to scan the backups depends on the number of backups that need to be imported and the size of the backup indexes.

You can scan indexes that are generated with the same or previous versions of NetVault. You cannot scan indexes generated with a newer version of NetVault on an earlier version of the server if the servers do not use the same index version. If an index version is not supported, the index is not imported and a message is generated in the logs.

- IMPORTANT:** When a backup stored on a disk-based storage device (such as DR Series system, NetVault SmartDisk, or Data Domain system) is retired, that backup is deleted from the device. You cannot import the deleted backup by scanning the device.

To scan a disk-based storage device:

- 1 In the Navigation pane, click **Manage Devices**. In the list of devices, locate the device and click the corresponding **Manage Device** icon .

Alternatively, in the Navigation pane, click **Explore Storage**. Click **Explore Disk Storage**, and in the repository table, select the device.

- 2 On the **RAS Device Management** or **Explore Disk Storage** page, click **Scan**.

- 3 In the **Scan Device** dialog box, you can configure the following option:

- **Minimum life for imported backups:** This option specifies the minimum life for backups that are imported to the NetVault Database. This option only applies to the backups that are not available in the NetVault Database.

The default value for this option is seven days. You can change the default setting by modifying the Media Manager settings. For more information, see [Configuring general settings for Media Manager](#).

To change the minimum life setting for the current session, type or select a new value. The minimum life setting is specified in number of days.

Depending on the value set for this option, the retirement time for imported backups is modified as follows:

- If a backup is scheduled to retire before the specified period, its retirement time is set to the specified minimum life.
- If a backup is scheduled to retire after the specified period, its retirement time remains unchanged. For such backups, the backup life setting determines the retirement time.

Click **Scan** to start the scanning process and close the dialog box.

- NOTE:** If a backup stored on a DR Series system was performed using a non-standard block size, the scan process is unable to read the index for that backup from the device. To import indexes for such backups, you can configure an alternate index read block size in the **mediamgr.cfg** file. For more information, see [Configuring an alternate index read block size for Quest DR Series systems](#).

Removing all savesets from a disk-based storage device

The **Expire All** method lets you remove all savesets from a disk-based storage device. You can use this method before removing a device from NetVault if you no longer require the backups stored on the device.

To remove all savesets from a disk-based storage device:

- 1 In the Navigation pane, click **Manage Devices**. In the list of devices, locate the device and click the corresponding **Manage Device** icon .
Alternatively, in the Navigation pane, click **Explore Storage**. Click **Explore Disk Storage**, and in the repository table, select the device.
- 2 On the **RAS Device Management** or **Explore Disk Storage** page, click **Expire All**.
- 3 In the confirmation dialog box, provide the following information:
 - **Password:** Type the password for the NetVault Server. If no password is set for the NetVault Server, provide the system's root or administrator password.
 - **Enter 'EXPIRE' to confirm request:** To confirm, type **BLANK** (case-insensitive) in this box.
- 4 Click **OK**.

Removing a disk-based storage device

You can use the following procedure to remove a disk-based storage device that is no longer required.

When you remove a device, the backups stored on the device are not deleted. You can add the device to the same or different NetVault Server to use the backups. If you add the device to a different server, you must scan the device to use the backups stored on it.

To remove a disk-based storage device:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate the device and click the corresponding **Manage Device** icon .
- 3 On the **RAS Device Management**, click **Remove**. In the confirmation dialog box, click **Remove** again.
- 4 If NetVault fails to remove the device, select the **Force Removal** check box in the confirmation dialog, and click **Remove**.

i | **NOTE:** You can use the **Force Removal** option to remove a device that is not in use. However, the device may still try to communicate with the NetVault Server.

Managing disk-based storage devices in tree view

This section includes the following topics.

- [Viewing disk device details](#)
- [Checking the status of a disk-based storage device](#)
- [Changing the status of a disk-based storage device](#)
- [Setting storage properties for a disk-based storage device](#)
- [Scanning a disk-based storage device](#)
- [Removing a disk-based storage device](#)

Viewing disk device details

The **Manage Devices** page shows all the devices that have been added to the NetVault Server. The current status of the devices is depicted using the following light icons.

Table 80. Device status icons

Status Indicator	Description
Green light	Device is online and available for use.
Yellow light	Device is in use. NetVault is able to detect the device, but unable to access it for backup or restore jobs.
Red light	Device is currently offline. NetVault can detect the device but it cannot be accessed for backup or restore jobs.
Red cross	Device is unavailable (the SCSI cable is disconnected, the device is removed, or any other reason). NetVault cannot detect the device.

To view disk device details:

- 1 In the Navigation pane, click **Manage Devices**. Select **Tree View**.
- 2 To view the details of a particular device, click the corresponding device and then click **Status**.
- 3 The device status dialog box includes the following tabs.

DR Series systems and Data Domain systems:

▪ Activity

- **Name:** The name of the storage device. The name is derived from the container name and the DNS name or IP address of the device.
- **Status:** The status of the device. AVAILABLE specifies that the device is available for backups and restores, while OFFLINE specifies that the device is unavailable and cannot be used for backups or restores.
- **Stream count:** The number of elements or NetVault segments currently being sent to the DR Series system or Data Domain system.
- **Claims on device from this server:** The number of Media Manager batch tasks using the DR Series system or Data Domain system. A batch task is associated with each backup, restore, scan, or load index stream connection to the device. A single job can have more than one claim on a device. For example, a single multi-channel Data Copy or Exchange backup can make multiple simultaneous claims on a device.
- **Send (MB/sec):** The total megabytes per second (MBps) across all streams being sent out from the DR Series system or Data Domain system to NetVault (that is, when an NetVault restore job is being performed).
- **Receive (MB/sec):** The total MBps across all streams being sent to the DR Series system (that is, when a NetVault backup job is being performed).

i | **NOTE:** When you add a DR Series system to multiple NetVault Servers, the **Activity** tab will show the accumulated transfer statistics for all servers.

▪ Device

- **Name:** The name of the storage device. The name is derived from the container name and the DNS name or IP address of the device.
- **Details:** The DR OS version and the Client RDA API version.
- OR -
Data Domain OS and Client DD Boost API version.

NetVault SmartDisk devices:

▪ Activity

- **Name:** The name of the storage device. The name is derived from the host name and IP address of the device.
- **Status:** The status of the device. AVAILABLE specifies that the device is available for backups and restores, while OFFLINE specifies that the device is unavailable and cannot be used for backups or restores.
- **Stream count:** The number of elements or NetVault segments currently being streamed to the NetVault SmartDisk.
- **Claims on device from this server:** The number of Media Manager batch tasks using the NetVault SmartDisk device. A batch task is associated with each backup, restore, scan, or load index stream connection to the device. A single job can have more than one claim on a device. For example, a single multi-channel Data Copy or Exchange backup can make multiple simultaneous claims on a device.

- **Send (MB/sec):** The total megabytes per second (MBps) across all streams being sent out from the NetVault SmartDisk device to NetVault (that is, when a NetVault restore job is being performed).
- **Receive (MB/sec):** The total MBps across all streams being sent to the NetVault SmartDisk device (that is, when a NetVault backup job is being performed).
- **Items deduplicating:** The number of elements or NetVault segments currently being deduplicated.
- **Deduplication Queue Length:** The number of elements or NetVault segments currently waiting to be deduplicated.
- **Garbage Collection State:** The current phase of Garbage Collection.
- **Storage**
 - **Data protected:** The total amount of data currently being protected by the NetVault SmartDisk device.
 - **Data Deduplicated:** The total amount of data that has been submitted for deduplication.
 - **Disk used by deduplicated data:** The amount of disk space used by deduplicated data. It includes the space consumed by the Chunk Store, Chunk Index, and Manifests.
 - **Dedupe ratio:** The deduplication ratio is calculated as follows:

$$\text{Deduplication ratio} = \frac{\text{Data Deduplicated}}{\text{Disk Used by Deduplicated Data}}$$
 - **Data in staging:** The amount of data stored in the Staging Store.
 - **Space available:** The total disk space available on the storage device.

This option shows the total disk space available across all configured Storage Volumes regardless of their deny or favor configuration. It does not include the Garbage Collection Reserve or Last Resort Threshold (LRT).
 - **Staging status:** The status of the Staging Store. It can be one of the following:
 - AVAILABLE: Staging Store is available for writing.
 - FULL: Staging Store is full; no more space is available for writing.
 - UNAVAILABLE: Disk index is unavailable.
 - UNLICENCED: NetVault SmartDisk is not licensed, the license has expired, or the license limit has exceeded.
 - NONE: The status is unknown.
 - **Staging space available:** The total space available to the NetVault SmartDisk device across all configured Storage Volumes that do not deny Staging. It does not include the Garbage Collection Reserve or LRT.
 - **Storage Status:** The status of the Storage space. It can be one of the following:
 - AVAILABLE: Storage is available for deduplication and rehydration.
 - FULL: Storage is full; no more space is available for data deduplication.
 - UNAVAILABLE: Chunk Index is unavailable; no data deduplication or rehydration is available.
 - UNLICENCED: NetVault SmartDisk is not licensed, the license has expired, or the license limit has exceeded.
 - NONE: The status is unknown.
 - **Storage space available:** The total space available to the NetVault SmartDisk device across all configured Storage Volumes that do not deny Storage. It does not include the Garbage Collection Reserve or LRT.

- **License:** The amount of protected capacity that is currently licensed by the NetVault SmartDisk device.
- **Device**
 - **Name:** The name of the NetVault SmartDisk device, which is automatically derived from the host name and IP address.
 - **Machine ID:** The Machine ID for the NetVault SmartDisk instance that is required for obtaining a permanent NetVault SmartDisk license key
 - **Version:** The version number of the device.
 - **License Type:** License type being used.
 - **License Expiration:** License validity period.
 - **Dedupe Licensed:** If the deduplication option is licensed.

4 Click **OK**.

Checking the status of a disk-based storage device

To check the status of an offline device:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the device.
- 3 Click **Check**.

If the device is operational, its status is changed to “**Available**.”

Changing the status of a disk-based storage device

To change the status to online or offline:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the device.
- 3 If the device is offline, click **Online** to bring it back online.
- 4 If the device is online, click **Offline** to take it offline.

The **Offline** method marks the device as offline and makes the device unavailable to NetVault. This method does not physically take the system offline.

Setting storage properties for a disk-based storage device

To set the storage properties:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the device.

- 3 Click **Properties**.
- 4 In the **Storage Properties** dialog box, type the Group Label.
- 5 Click **OK**.

Scanning a disk-based storage device

You can use the **Scan** method to query all backups stored on a disk-based storage device and import those backups that are not indexed in the given NetVault Server's database.

To import backups to the database, the NetVault Server must have the same NetVault Machine Name as the original server that performed the backups. The amount of time it takes to scan the backups depends on the number of backups that need to be imported and the size of the backup indexes.

You can scan indexes that are generated with the same or previous versions of NetVault. You cannot scan indexes generated with a newer version of NetVault on an earlier version of the server if the servers do not use the same index version. If an index version is not supported, the index is not imported and a message is generated in the logs.

i | **IMPORTANT:** When a backup stored on a disk-based storage device (such as DR Series system, NetVault SmartDisk, or Data Domain system) is retired, that backup is deleted from the device. You cannot import the deleted backup by scanning the device.

To scan a disk-based storage device:

- 1 In the Navigation pane, click **Manage Devices**. In the list of devices, locate and click the device.
Alternatively, in the Navigation pane, click **Explore Storage**. Click **Explore Disk Storage**, and in the repository table, select the device.
- 2 Click **Scan**.
- 3 In the confirmation dialog box, click **OK**.

i | **NOTE:** If a backup stored on a DR Series system was performed using a non-standard block size, the scan process is unable to read the index for that backup from the device. To import indexes for such backups, you can configure an alternate index read block size in the **mediamgr.cfg** file. For more information, see [Configuring an alternate index read block size for Quest DR Series systems](#).

Removing a disk-based storage device

You can use the following procedure to remove a disk-based storage device that is no longer required.

When you remove a device, the backups stored on the device are not deleted. You can add the device to the same or different NetVault Server to use the backups. If you add the device to a different server, you must scan the device to use the backups stored on it.

To remove a disk-based storage device:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the device.
- 3 Click **Remove**. In the confirmation dialog box, click **OK**.
- 4 If NetVault fails to remove the device, select the **Force Removal** check box in the confirmation dialog, and click **OK**.

i | **NOTE:** You can use the **Force Removal** option to remove a device that is not in use. However, the device may still try to communicate with the NetVault Server.

Managing the Snapshot Array Manager

This section includes the following topics:

- [Modifying the Snapshot Array Manager settings](#)
- [Changing the user credentials for Snapshot Array Manager](#)
- [Changing the status of Snapshot Array Manager](#)
- [Removing the Snapshot Array Manager](#)

Modifying the Snapshot Array Manager settings

To modify the Snapshot Array Manager settings:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate the Snapshot Array Manager, and click the corresponding **Manage Device** icon .
- 3 On the **Snapshot Array Manager Management** page, modify the required settings.

Table 81. Modify Snapshot Array Manager settings

Option	Description
Device Display Name	Specify a display name for the Enterprise Manager.
Array Manager Type	The Array Manager Type cannot be modified after an array manager is added.
Network name/IP address	Displays the Fully Qualified Domain Name (FQDN) or IP address of the Enterprise Manager.
Network port	By default, the Enterprise Manager uses port number 3033 for receiving communication from all clients. If you have changed the Web Server port on the Enterprise Manager, specify the port number.
Username	Displays the username. To change the user account, use the Change Credentials button.
Status	Displays the status icon and the device status (online, offline, or unavailable). To change the status, click Mark Online or Mark Offline .

- 4 Click **Save** to save the settings.

Changing the user credentials for Snapshot Array Manager

To change the user credentials for Snapshot Array Manager:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate the Snapshot Array Manager, and click the corresponding Manage Device icon .
- 3 On the **Snapshot Array Manager Management** page, click **Change Credentials**, and on the **Change Snapshot Array Manager Credentials** page, provide the following details.

Table 82. Change credentials for Snapshot Array Manager

Option	Description
Username	Specify a user account that can be used to log on to the Enterprise Manager.
New password	Type a new password for the user account.
Confirm password	Re-type the password for confirmation.

- 4 Click **Save** to save the user credentials.

Changing the status of Snapshot Array Manager

To change the status to online or offline:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate the Snapshot Array Manager, and click the corresponding **Manage Device** icon .
- 3 On the **Snapshot Array Manager Management** page, click **Mark Online** or **Mark Offline**:
 - If the Snapshot Array Manager is offline, click **Mark Online** to bring it back online.
 - If the Snapshot Array Manager is online, click **Mark Offline** to take it offline.

The Offline method marks the Snapshot Array Manager as offline and makes it unavailable to NetVault. This method does not physically take the system offline.

Removing the Snapshot Array Manager

To remove the Snapshot Array Manager:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate the Snapshot Array Manager, and click the corresponding **Manage Device** icon .
- 3 Click **Remove**, and then in the confirmation dialog box, click **OK**.

Managing tape libraries in list view

This section includes the following topics:

- [Viewing tape library details](#)
- [Opening and closing library door](#)
- [Opening and closing entry/exit ports](#)
- [Unloading or importing tapes from entry/exit ports](#)
- [Exporting tapes to entry/exit ports](#)
- [Restarting ACSLS or NDMP Libraries](#)
- [Importing shadow tapes \(NetApp VTL\)](#)

- [Removing a tape library](#)

Viewing tape library details

To view tape library details:

- 1 In the Navigation pane, click **Manage Devices**.

On the **Manage Devices** page, you can view the list of storage devices added to the NetVault Server. The device status is indicated using the following icons.

Table 83. Device status icons

Icon	Description
	Device is online and available for use.
	Device is offline. NetVault is able to detect the device, but unable to access it for backup or restore jobs.
	Device is unavailable. NetVault is unable to detect the device.

- 2 To view the details of a particular tape library, click the corresponding Manage Library icon .

- 3 On the **Tape Library Management** page, you can view the following information:

- **Library details:** This area shows the following information:
 - **Name:** Library name.
 - **Vendor:** Vendor name.
 - **Product:** Library type.
 - **Drives:** Number of drives.
 - **Slots:** Number of slots.
 - **Client:** Client to which the device is attached.
 - **Status:** Device status (online or offline).
 - **Activity:** Device or door status.
 - **Door:** Door status (open or closed).
- **Drives table:** The Drives table lists all tape drives for the library. It shows the following information:
 - **Status icon:** Drive status icon.
 - **Bay:** Drive bay number.
 - **Name:** Drive name.
 - **Status:** Drive status (online or offline).
 - **Activity:** Idle or writing.
 - **Contents:** Label of the tape loaded in the drive. If the drive does not contain any tape, the column shows “Unloaded.”

- 4 To perform a library-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

Opening and closing library door

Before opening a library door, you must issue the **Open Door** command from the WebUI. NetVault puts a software lock on the library door to prevent anyone from opening the door without issuing the **Open Door** command. If you do not issue this command, NetVault does not know when tapes are added, removed, or rearranged, and it may attempt to load non-existent media.

To open a library door from WebUI:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to locate the library, and then click the corresponding Manage Library icon .
- 3 On the **Tape Library Management** page, click **Open Door**.
The library goes offline when you open the door.
- 4 To bring it back online, click **Close Door**.

Opening and closing entry/exit ports

To open or close an entry/exit port:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library, and then click the **Ports** link or the corresponding Manage Ports button to open the Port Browser.
- 3 To open an entry/exit port, select the port in the Ports table, and click **Open Port**.
When you issue the **Open Port** command, NetVault puts a software lock on library so that it knows that the port door is going to be opened.
- 4 To close the port after placing a tape, select **Close Port**. If you have placed a cleaning tape in the port, select **Close port with cleaning media**.
When you issue the **Close Port** command, NetVault knows that you have physically closed the port door, and it removes the lock on the library.

Unloading or importing tapes from entry/exit ports

To unload a tape from an entry/exit port:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library, and then click the **Ports** link or the corresponding Manage Ports button to open the Port Browser.
- 3 In the Ports table, select the slot that contains the tape, and click **Unload from Port**.
The tape is moved to a drive or slot:
 - If the media barcode is available in the NetVault Database, the tape is loaded to a free slot.
 - If the tape does not have a barcode or the barcode information is not available in the NetVault Database, the tape is loaded to a drive to read the header.

Exporting tapes to entry/exit ports

To export a tape to an entry/exit port:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library that contains the tape, and then click the **Slots** link or the corresponding Manage Slots button to open the Slot Browser.
- 3 In the Slots table, select the applicable slots, and click **Export**.

After sending a request to export the tapes to entry/exit ports, the WebUI opens the Port Browser page. In the Ports table, you can view the exported tapes.

- 4 In the Ports table, select any tape, and click **Open Port**.

When you issue the **Open Port** command, NetVault puts a software lock on library so that it knows that the port door is going to be opened.

- 5 Physically open the port door, and after removing the tape from the entry/exit port, close the port door. For more information about these procedures, see the device manual.

- 6 On the Slot Browser page, click **Close Door**.

When you issue the **Close Door** command, NetVault knows that you have physically closed the port door, and it removes the lock on the library.

Restarting ACSLS or NDMP Libraries

If an ACSLS or NDMP library encounters a network problem, use the following procedure to restart the library. The **Restart** method restarts the network and socket connections by removing them and adding the library again.

To restart an ACSLS or NDMP library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to locate the library, and then click the corresponding Manage Library icon .
- 3 On the **Tape Library Management** page, click **Restart**.
- 4 In the confirmation dialog box, click **OK**.

Importing shadow tapes (NetApp VTL)

The NetApp VTL Shadow Tape option allows you to quickly import a tape from the shadow tape pool whenever possible instead of obtaining the physical tape.

To use shadow tapes, you need to do the following:

- Configure the **Enable Shadow Tapes** option on the filer and library containing the virtual tapes. For more information about enabling shadow tapes, consult the relevant NetApp VTL documentation.
- Select barcodes as the default labels for virtual tapes in NetVault. NetVault requires this method of labeling media to be fully functional. You can use barcodes as default labels by selecting the **Use Barcodes as Labels** check box in the Media Manager settings dialog box. For more information, see [Configuring general settings for Media Manager](#).

With shadow tapes enabled, whenever a virtual tape is exported to a physical tape, the virtual tape is moved to the shadow tape pool. The shadow tape pool is invisible to the backup application and it is not listed as part of a virtual library, but it is available for quick access if the physical tape is later imported. It is also available for reading if the physical tape is stored off-site or is otherwise unavailable.

The NetApp VTL manages the space used by shadow tapes. It can delete a shadow tape if more space is required for the new backup data. The administrator can set a preferred retention time for shadow tapes. If the retention period has not expired, the NetApp VTL sends a notification before deleting the shadow tape.

Note the following:

- The media requests for shadow tapes can only be used for Restore or Duplication tasks as these tapes are converted to read-only virtual tapes.
- No permanent records are created in the NetVault Database for shadow tape media. The database only stores the details of actual media. The shadow tape attribute is associated with the media when they are imported to the library as shadow tapes. Therefore, you must export all shadow media before stopping or restarting NetVault. If you fail to export the media, they lose the shadow attribute and are converted to read-only items. For the same reason, you must export the shadow tapes before opening a library door.
- Error messages are displayed if you try to import media when nothing is available in the shadow tape pool or when the shadow tapes option is not supported on the device.

To import virtual tapes:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to locate the library, and then click the corresponding Manage Library icon .
- 3 On the **Tape Library Management** page, click **Import Media**.
- 4 In the **Media Barcode** list, select or type the barcode for the tapes that you want to import.
- 5 Click **Import**.

The requested tapes are imported to the medium changer from the shadow tape pool or physical library. When both shadow tapes and physical tapes are available, the shadow tapes are converted to read-only virtual tapes and imported to the entry/exit port. When only the physical tapes are available, virtual tapes are created from the physical tapes and imported to the entry/exit port.

Removing a tape library

You can use the following procedure to remove a tape library that is no longer required.

When you remove a tape library, it does not delete the media information from the NetVault Database. You can use the media on any other library that supports the media type. Scanning is not required if you use the media in the same NetVault domain. In a different NetVault Domain, you must scan the media to access the backups.

To remove a tape library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to locate the library, and then click the corresponding Manage Library icon .
- 3 On the **Tape Library Management** page, click **Remove**, and then in the confirmation dialog box, click **OK**.

Managing tape libraries in tree view

This section includes the following topics.

- [Viewing tape library details](#)
- [Modifying a library](#)
- [Changing the device view type](#)

- Opening and closing library door
- Opening and closing entry/exit ports
- Unloading or importing tapes from entry/exit ports
- Exporting tapes to entry/exit ports
- Restarting ACSLS or NDMP libraries
- Importing shadow tapes (NetApp VTL)
- Removing a tape library

Viewing tape library details

The **Manage Devices** page shows all the devices that have been added to the NetVault Server. The current status of the devices is depicted using the following light icons.

Table 84. Device status icons

Status Indicator	Description
Green light	Device is online and available for use.
Yellow light	Device is in use. NetVault is able to detect the device, but unable to access it for backup or restore jobs.
Red light	Device is currently offline. NetVault can detect the device but it cannot be accessed for backup or restore jobs.
Red cross	Device is unavailable (the SCSI cable is disconnected, the device is removed, or any other reason). NetVault cannot detect the device.

To view device status and performance statistics:

- 1 In the Navigation pane, click **Manage Devices**. Select **Tree View**.
- 2 To view the details of a particular device, click the corresponding library and then click **Status**.
- 3 The library status dialog box includes the following information under the Details tab:
 - **Name:** Library name.
 - **Machine:** Machine name.
 - **Status:** Device status (online or offline).
- 4 Click **OK** to close the Details dialog box.

Modifying a library

To modify a library, follow these steps:

- 1 In the Navigation pane, click **Manage Devices**. Select **Tree View**.
- 2 Click the applicable library, and click **Modify**.
- 3 **Edit Device** page appears. This page displays **Library Configuration** area and **Selected Drives** area. Click **Show Drives** tab to display **Choose Drives** area.
- 4 In the **Library Configuration** area, click the applicable library to display the library picture, type, and details in the **Selected Drives** area. The Selected Drives area shows the details of the first drive in the selected bay.
- 5 Click the library and then click **Configure**. Configure the parameters under the following tabs:

Option tab	Description
Configuration	For more information about these settings, see General settings for tape libraries .
Cleaning	For more information about these settings, see Drive cleaning settings .
Entry/Exit Ports	Select the desired options for entry/exit ports.
Mixed Media	For more information about these settings, see Mixed media settings .
SCSI Configuration	<p>The default timeout value for the different types of SCSI commands is set to zero, which correspond to the following intervals:</p> <ul style="list-style-type: none"> • Fast SCSI Commands – 300 seconds • Slow SCSI Commands – 900 seconds • Very Slow SCSI Commands – 3 hours <p>An error is logged if a command execution does not complete within the specified interval. You should not change the default timeout for any of the SCSI commands unless otherwise advised by Quest Support.</p>

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

7 Click **Save**.

Changing the device view type

NetVault provides two device view types:

- **Physical View** – This is the default view type. This view displays the actual structure of the library, including all drives, slots, and entry/exit ports. Media are shown at their current location in the library.
- **Logical View** – This view is centered around the actual media in a library. The device tree consists of two folders – Drives and Media. The libraries and drives are grouped as Drives, and the Media are grouped according to the Group Label.

To change the device view type, follow these steps:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 Click **Change View**.
- 3 NetVault will change the view from Physical to Logical or vice versa, and displays a message.

Opening and closing library door

Before opening a library door, you must issue the **Open Door** command from the WebUI. NetVault puts a software lock on the library door to prevent anyone from opening the door without issuing the **Open Door** command. If you do not issue this command, NetVault does not know when tapes are added, removed, or rearranged, and it may attempt to load non-existent media.

To open a library door from WebUI:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the applicable library. Click **Open Door**.
The library goes offline when you open the door.
- 3 To bring it back online, click the library and click **Close Door**.

Opening and closing entry/exit ports

To open or close entry/exit ports, follow these steps:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the applicable library. Click **Open Entry/Exit**.
- 3 To close the port after placing the media, click **Entry/Exit Closed**. If you have placed a cleaning media item in the port, click **Entry/Exit Closed with Cleaning Media** to move the media item to a cleaning slot in the library.

Unloading or importing tapes from entry/exit ports

To unload a tape from an entry/exit port:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the applicable library.
- 3 Select the port slot that contains the tape, and click **Unload**.

The tape is moved to a drive or slot:

- If the media barcode is available in the NetVault Database, the tape is loaded to a free slot.
- If the tape does not have a barcode or the barcode information is not available in the NetVault Database, the tape is loaded to a drive to read the header.

Exporting tapes to entry/exit ports

To export a tape to an entry/exit port:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the applicable slot. Click **Export**.

Restarting ACSLS or NDMP libraries

If an ACSLS or NDMP library encounters a network problem, use the following procedure to restart the library. The **Restart** method restarts the network and socket connections by removing them and adding the library again.

To restart an ACSLS or NDMP library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the applicable library. Click **Restart Library**.
- 3 In the confirmation dialog box, click **OK**.

Importing shadow tapes (NetApp VTL)

The NetApp VTL Shadow Tape option allows you to quickly import a tape from the shadow tape pool whenever possible instead of obtaining the physical tape.

To use shadow tapes, you need to do the following:

- Configure the **Enable Shadow Tapes** option on the filer and library containing the virtual tapes. For more information about enabling shadow tapes, consult the relevant NetApp VTL documentation.
- Select barcodes as the default labels for virtual tapes in NetVault. NetVault requires this method of labeling media to be fully functional. You can use barcodes as default labels by selecting the **Use Barcodes as Labels** check box in the Media Manager settings dialog box. For more information, see [Configuring general settings for Media Manager](#).

With shadow tapes enabled, whenever a virtual tape is exported to a physical tape, the virtual tape is moved to the shadow tape pool. The shadow tape pool is invisible to the backup application and it is not listed as part of a virtual library, but it is available for quick access if the physical tape is later imported. It is also available for reading if the physical tape is stored off-site or is otherwise unavailable.

The NetApp VTL manages the space used by shadow tapes. It can delete a shadow tape if more space is required for the new backup data. The administrator can set a preferred retention time for shadow tapes. If the retention period has not expired, the NetApp VTL sends a notification before deleting the shadow tape.

Note the following:

- The media requests for shadow tapes can only be used for Restore or Duplication tasks as these tapes are converted to read-only virtual tapes.
- No permanent records are created in the NetVault Database for shadow tape media. The database only stores the details of actual media. The shadow tape attribute is associated with the media when they are imported to the library as shadow tapes. Therefore, you must export all shadow media before stopping or restarting NetVault. If you fail to export the media, they lose the shadow attribute and are converted to read-only items. For the same reason, you must export the shadow tapes before opening a library door.
- Error messages are displayed if you try to import media when nothing is available in the shadow tape pool or when the shadow tapes option is not supported on the device.

To import virtual tapes:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the library containing the virtual tape.
- 3 Click **Import Media**. This displays the **Select Shadow Tape Media Barcode for Import** dialog box which provides a list of media barcodes for the tapes available in the shadow tape pool.
- 4 In the **Media Barcode** list, select or type the barcode for the tapes that you want to import.
- 5 Click **OK**.

The requested tapes are imported to the medium changer from the shadow tape pool or physical library. When both shadow tapes and physical tapes are available, the shadow tapes are converted to read-only virtual tapes and imported to the entry/exit port. When only the physical tapes are available, virtual tapes are created from the physical tapes and imported to the entry/exit port.

Removing a tape library

You can use the following procedure to remove a tape library that is no longer required.

When you remove a tape library, it does not delete the media information from the NetVault Database. You can use the media on any other library that supports the media type. Scanning is not required if you use the media in the same NetVault domain. In a different NetVault Domain, you must scan the media to access the backups.

To remove a tape library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the applicable library.
- 3 Click **Remove**, and then in the confirmation dialog box, click **OK**.

Managing tape drives in list view

This section includes the following topics:

- [Viewing tape drive details](#)
- [Configuring performance options for a tape drive](#)
- [Checking the status of a tape drive](#)
- [Changing the status of a tape drive](#)
- [Configuring cleaning slots](#)
- [Configuring the cleaning lives option](#)
- [Configuring automatic cleaning options for a tape drive](#)
- [Manually submitting a drive cleaning request](#)
- [Unloading a tape](#)
- [Loading a tape](#)
- [Removing a tape drive](#)

Viewing tape drive details

To view tape drive details:

- 1 In the Navigation pane, click **Manage Devices**.
On the **Manage Devices** page, you can view the list of storage devices added to the NetVault Server.
- 2 In the list of devices, click the icon  to open the library to list the available drives and slots. For each drive, the activity and status messages (for example, Idle, Writing, Loading media, and others) are displayed on the page.

The device status is indicated using the following icons.

Table 85. Device status icons

Icon	Description
	Device is online and available for use.
	Device is offline. NetVault is able to detect the device, but unable to access it for backup or restore jobs.
	Device is unavailable. NetVault is unable to detect the device.

- 3 To view the details of a particular tape drive, click the drive or the corresponding Manage Drive icon .
- 4 On the **Tape Drive Management** page, you can view the following information:
 - **Drive Info:** This area shows general information about the drive:

- **Name:** Drive name.
- **Vendor:** Vendor name.
- **Product:** Product name.
- **Client:** Client to which the device is attached.
- **Serial Number:** Serial number of the drive.
- **Status:** Online or offline.
- **Contents:** Label of the tape loaded in the drive. If the drive does not contain any tape, the column shows "Unloaded."
- **Offsite Location:** Offsite location of the tape, if specified.
- **Library:** Library name.
- **Bay:** Drive bay number.
- **Block Size:** Media block size.
- **Buffer Size:** Transfer buffer size.
- **Activity:** Idle or writing.
- **Statistics Info:** This area displays the drive usage statistics:
 - **Total Data Written:** Total amount of data written using the drive.
 - **Total Data Read:** Total amount of data read using the drive.
 - **Write Errors:** Number of write errors reported.
 - **Read Errors:** Number of read errors reported.
 - **Date Last Written:** Date on which the last write operation was performed.
 - **Date Last Read:** Date on which the last read operation was performed.
- **Cleaning Info:** This area displays the drive cleaning statistics:
 - **Date Last Cleaned:** Date on which the drive last cleaning operation was performed.
 - **Time Since Last Cleaning:** Time elapsed since the last automatic or manual cleaning operation.
 - **Times Cleaned:** Number of times the drive has been cleaned.
 - **Transferred Since Last Clean:** Amount of data read or written since the last cleaning operation.
 - **Usage Since Last Clean:** Duration of time the drive has been used for read or write operations since the last cleaning operation.
 - **Soft Errors Since Last Clean:** Number of read or write errors reported since the last cleaning operation.
- **Activity chart:** This area shows the activity graph if the device is being used by a backup or restore job.
- **Job details:** This area shows the following information for active jobs:
 - **Rate:** Data transfer rate.
 - **Job Title:** Name of the job.
 - **Job ID:** Job identification number.
 - **Job Instance:** Instance identification number.
 - **Job Phase:** Phase identification number (1 or 2).

5 To view the slot details, click the **Slots** link or the Manage Slots icon . On the Slot Browser page, you can view the following information:

- **Drives table:** The Drives table lists all tape drives for the library. It shows the following information:
 - **Status icon:** Drive status icon.
 - **Bay:** Drive bay number.
 - **Name:** Drive name.
 - **Status:** Drive status (online or offline)
 - **Activity:** Idle or writing.
 - **Contents:** Label of the tape loaded in the drive. If the drive does not contain any tape, the column shows “Unloaded.”
- **Slots table:** The Slots table lists all slots for the library. It shows the following information:
 - **Slot:** Slot number
 - **Status:** Tape is loaded or unloaded.
 - **Barcode:** Barcode of the tape in the slot.
 - **Media:** Media label of the tape in the slot.
 - **Media Group:** Media group label of the tape in the slot.
 - **Space Available:** Free space available in the tape.
 - **Cleaning lives status icon:** Displays the cleaning lives icon if the slot contains a cleaning tape.  indicates 5 or more cleaning lives remaining,  indicates 1-4 cleaning lives remaining, and  indicates 0 cleaning lives remaining.
 - **Cleaning Lives Left:** Displays the number of cleaning lives left.

6 To perform a device-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

Configuring performance options for a tape drive

To configure performance options for a tape drive:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library that contains the drive, and then click the drive or the corresponding Manage Drive icon .
- 3 On the **Tape Drive Management** page, click **Performance**.
- 4 In the **Drive Performance Options**, configure the following settings.

Table 86. Drive performance options

Option	Description
Drive Block Size	<p>This option specifies the block size used for read and write operations. The default value is 64KiB.</p> <p>NOTE: The drive block size should be the same size as the media block size for the VTL.</p> <p>You can change the media block size in increments of 1KiB, but many devices only accept a value in the multiples of 4KiB or 32KiB.</p> <p>NOTE: The changes to the media block size settings are only applied to blank media items. If you are reusing a media item, make it blank first for these changes to take effect.</p> <p>Increasing the block size can reduce the number of times a backup needs to read data and write it to media. However, large media block sizes do not always imply an overall faster backup. The maximum block size is limited by several factors, such as the OS, SCSI adapter, drive make, drive model, and drive type.</p> <p>On Linux and UNIX systems, you can increase the media block size for optimum performance.</p> <p>On Windows, you might be required to change the registry setting MaximumSGList to use block sizes larger than 64KB. Before changing this setting, check that the SCSI bus is only used by the tape devices. If other devices also use the SCSI bus, this registry change might prevent them from working. If you want to apply these changes only to a specific channel on the HBA, consult the hardware vendor.</p> <p>To change the registry setting on Windows, follow these steps:</p> <ol style="list-style-type: none"> 1 Start the Registry Editor. 2 Open the key <code>[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\<HBA Vendor>\Parameters</code> (where <HBA Name> is specific to your SCSI card — for example, QL2200 for a Qlogic 2200 card). 3 Create the Parameters key, if it is not present. 4 Under Parameters, create the Device key, if it is not present. 5 Under the Device key, add the DWORD registry value MaximumSGList, if it is not present. 6 Calculate the hexadecimal value of MaximumSGList: <ul style="list-style-type: none"> On 32-bit systems: $\text{MaximumSGList} = (\text{Maximum Block Size} / 4\text{KiB}) + 1$ <p>For example, if the block size is set to 256KiB, the value for this key is:</p> $(256\text{KiB}/4\text{KiB}) + 1 = 65$ <p>The decimal value is 65 and the hexadecimal value is 0x41.</p> <p>You can set the block size to any value from 64KiB through 1012KiB. The maximum value 255 is internally converted to 257 to make a block size of 1 MiB (1024 KiB).</p> On 64-bit systems: <p>On 64-bit systems, the default OS page size is 8KiB. The formula for calculating MaximumSGList is:</p> $\text{MaximumSGList} = (\text{Maximum Block Size} / 8\text{KiB}) + 1$ <p>Thus, the maximum value of 255 corresponds to a maximum media block size of 2MiB.</p> 7 Reboot the system to apply the changes.

Table 86. Drive performance options

Option	Description
Drive Transfer Buffer Size	<p>The transfer buffer or the shared memory is allocated in blocks of 32KiB. The default value is 8193KiB. For more information about changing the default value, see Changing the default Drive Transfer Buffer Size option.</p> <p>Increasing the transfer buffer size can improve backup performance. To calculate the buffer size, use the following formula:</p> $(<\text{Total number of buffers}> \times 32\text{KiB}) + 1 \text{ byte}$ <p>On Linux and UNIX systems, you require sufficient RAM and large Shared Memory segment. Before increasing the transfer buffer size, check the following settings on these platforms:</p> <ul style="list-style-type: none"> • Maximum size of a shared memory segment (SHMMAX) • Minimum size of shared memory segment (SHMMIN) • Maximum number of shared memory identifiers in the system (SHMMNI) • Maximum number of shared memory segments a user process can attach (SHMSEG) • Maximum number of semaphore identifiers in the system (SEMMNI) • Maximum number of semaphores in a set (SEMMSL) • Maximum number of semaphores in the system (SEMMS) • Maximum number of operations per semop call (SEMOPM) • Semaphore maximum value (SEMVMX) <p>The total allowed shared memory is determined by the formula $\text{SHMMAX} * \text{SHMSEG}$. These values are often limited by the ulimit setting, and the command <code>ulimit -a</code> can be used to view these system settings.</p> <p>On Windows, you require at least 2GB RAM and large virtual memory. You might also have to change the MaximumSGList setting on the SCSI card.</p> <p>For examples, see the Optimal transfer buffer size.</p>

- 5 To enable software compression for virtual tape drives, configure the following options.

Table 87. Software compression options for virtual tape drives

Option	Description
Software Data Compression	To perform software compression, select this check box. The data is compressed when it is transferred to the device during backup.

Table 87. Software compression options for virtual tape drives

Option	Description
Compression Threshold	<p>The value set for this option determines the minimum level of compression that must be achieved when data is compressed during a backup. For example, if you set the value to 80 percent, one of the following occurs:</p> <ul style="list-style-type: none"> • If the compressed data size is less than 80 percent of the original data size, the data is backed up in its compressed form. • If the compressed data size is more than 80 percent of the original data size, the data is backed up in its uncompressed form. <p>If you specify 80 percent, a file size of a 100MB must be \leq 80MB after compression however, the file size can be $>$80MB (like 81MB, 82MB, etc.), as the compression is done at block level.</p> <p>When the specified compression level is not achieved for any block, in that case NetVault backs up that block in its uncompressed form. Where, some blocks may be compressed and some blocks may remain in uncompressed form. Hence, in some case, the file size after compression can be greater than 80 percent.</p> <p>The extent that data can be compressed depends on the data contents. Also the encrypted data cannot be compressed. With some files, compression may actually result in a file that is larger than the original uncompressed file.</p>
Compression Data Blocks	<p>Type or select the number of data blocks per compression unit. The default block size is 8KiB.</p>

- 6 Click **OK** to save the settings and close the dialog box.

Optimal transfer buffer size

The following table provides examples of the optimal values that can be used for some drive types.

Table 88. Optimal transfer buffer size of different drive types

Drive type	Optimal transfer buffer size (in KiB)
Fast Modern Tape Drives For example, LTO series, SDLT, and SAIT	65537 (64MiB + 1KiB)
Medium Speed Tape Drives For example, DLT8000, DLT7000, and AIT-3	32769 (32MiB + 1KiB)
Older Professional Tape Drives For example, DLT2000, DLT4000, and AIT-2	16385 (16MiB + 1KiB)
Older Low Capacity Low-end Drives For example, EXB-8505, AIT-1, and DAT	8193 (8MiB + 1KiB)

Checking the status of a tape drive

To check the status of an offline tape drive:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library that contains the drive, and then click the drive or the corresponding Manage Drive icon .
- 3 On the **Tape Drive Management** page, click **Check**, and then in the confirmation dialog box, click **Check** again.

If the device is operational, its status is changed to “**Available.**”

Changing the status of a tape drive

To change the drive status to online or offline:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library that contains the drive, and then click the drive or the corresponding Manage Drive icon .
- 3 On the **Tape Drive Management** page, click **Online** or **Offline**:
 - If the device is offline, click **Online** to bring it back online.
 - If the device is online, click **Offline** to mark it offline.

This operation does not physically take the device offline; it just makes the device unavailable for use in NetVault.

Configuring cleaning slots

Before loading cleaning tapes, you must configure the slots that can hold the cleaning media. You can configure cleaning slots from the Slot Browser page.

i | **NOTE:** If cleaning slots are not defined, cleaning tapes without barcodes are loaded every time the library is initialized.

To configure cleaning slots for a library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library, and then click the **Slots** link or the corresponding Manage Slots icon  to open the Slot Browser.
- 3 In the Slots table, select the slot that you want to use, and click **Set Slot**.

The slot must be empty.
- 4 In the **Slot Settings** dialog box, select the **Set as cleaning slot** check box to reserve the slot for cleaning tapes. In the **Cleaning Slot** list, type or select the slot number.
- 5 Click **OK** to save the settings and close the dialog box.
- 6 After the dialog box is closed, a message stating “Library device needs to be restarted” is displayed.

Access the **Manage Devices** page, and restart the library.

When you open the Slot Browser again, the slot status is set to **CLN Slot**.

Configuring the cleaning lives option

NetVault lets you specify how many times a cleaning tape can be used for drive cleaning operations. You can configure the cleaning lives property from the Slot Browser page.

i | **NOTE:** To configure the cleaning lives option, a cleaning tape must be placed in the designated slot.

To configure the cleaning lives for a tape:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library, and then click the **Slots** link or the corresponding Manage Slots icon  to open the Slot Browser.
- 3 In the Slots table, select the slot that contains the cleaning tape, and click **Set Cleaning Life**.
- 4 In the **Set Cleaning Life** dialog box, configure the following option:
 - **Set media lives:** Type or select the number of times the tape can be used for cleaning a drive. The default value is 0.
- 5 Click **OK** to save the settings and close the dialog box.

Configuring automatic cleaning options for a tape drive

You can set up automatic cleaning routines for tape drives from the **Tape Drive Management** page. NetVault lets you use the following options to define a cleaning routine: number of days, amount of data transferred, hours of use, and number of read or write errors.

To configure automatic cleaning options for a tape drive:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library that contains the drive, and then click the drive or the corresponding Manage Drive icon .
- 3 On the **Tape Drive Management** page, click **Drive Cleaning Options**.
- 4 In the **Drive Cleaning Options** dialog box, configure the following options.

Table 89. Drive Cleaning Options

Option	Description
Days	To perform drive cleaning after every x days, select the Apply check box to the left, and type or select the value.
Data Transferred	To perform drive cleaning after every x GiB of data transfer, select the Apply check box to the left, and type or select the value.
Hours of Use	To perform drive cleaning after every x hour, select the Apply check box to the left, and type or select the value.
Soft Read/Write Errors	To perform drive cleaning after every x*100 soft read/write errors, select the Apply check box to the left, and type or select the value.

- 5 Click **OK** to save the settings and close the dialog box.

Manually submitting a drive cleaning request

You can manually submit a drive cleaning request from the **Tape Drive Management** page.

To manually submit a drive cleaning job:

- 1 In the Navigation pane, click **Manage Devices**.

- 2 In the list of devices, click the icon  to open the library that contains the drive, and then click the drive or the corresponding Manage Drive icon .
- 3 On the **Tape Drive Management** page, click **Clean Drive**.
After the task is completed, a message is displayed.

Unloading a tape

To unload a tape:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library, and then click the drive or the corresponding Manage Drive icon .
- 3 Click **Unload**.
After unload request is sent successfully, a message is displayed.
In a library, the tape is moved to an available slot, while in a standalone drive the tape is ejected.

Loading a tape

To load a tape:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library, and then click **Slots** to open the Slot Browser.
- 3 In the list of slots, select the slot that contains the tape, and click **Load**.
After the load request is sent successfully, a message is displayed.
The tape is loaded in an available drive.

Removing a tape drive

You can use the following procedure to remove a tape drive that is no longer required.

When you remove a tape drive, the information about the tape that is loaded is not deleted from the NetVault Database. You can use the tape on any other drive that supports the media type. Scanning is not required if you use the tape in the same NetVault domain. In a different NetVault Domain, you must scan the tape to access the backups.

To remove a tape drive:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to open the library that contains the drive, and then click the drive or the corresponding Manage Drive icon .
- 3 On the **Tape Drive Management** page, click **Remove**, and then in the confirmation dialog box, click **OK**.

Managing tape drives in tree view

This section includes the following topics:

- [Viewing tape drive details](#)
- [Configuring physical tape drives](#)
- [Checking the status of a tape drive](#)
- [Changing the status of a tape drive](#)
- [Configuring cleaning slots](#)
- [Configuring the cleaning lives option](#)
- [Configuring automatic cleaning options for a tape drive](#)
- [Manually submitting a drive cleaning request](#)
- [Loading a tape](#)
- [Marking media for re-use](#)
- [Removing a tape drive](#)

Viewing tape drive details

The **Manage Devices** page shows all the devices that have been added to the NetVault Server. The current status of the devices is depicted using the following light icons.

Table 90. Device status icons

Status Indicator	Description
Green light	Device is online and available for use.
Yellow light	Device is in use. NetVault is able to detect the device, but unable to access it for backup or restore jobs.
Red light	Device is currently offline. NetVault can detect the device but it cannot be accessed for backup or restore jobs.
Red cross	Device is unavailable (the SCSI cable is disconnected, the device is removed, or any other reason). NetVault cannot detect the device.

To view tape drive details:

- 1 In the Navigation pane, click **Manage Devices**. Select **Tree View**.
- 2 To view the details of a particular device, click the corresponding device and then click **Status**.
- 3 The device status dialog box includes the information under the following tabs:
 - **Details:** This area shows general information about the drive:
 - **Name:** Drive name.
 - **Machine:** Machine name.
 - **Drive status:** Online or offline.
 - **Media status:** Unloaded or empty.
 - **Drive Statistics:** This area shows the drive usage statistics:
 - **Total Data Written:** Total amount of data written using the drive.
 - **Total Data Read:** Total amount of data read using the drive.

- **Write Errors:** Number of write errors reported.
 - **Read Errors:** Number of read errors reported.
 - **Date Last Written:** Date on which the last write operation was performed.
 - **Date Last Read:** Date on which the last read operation was performed.
 - **Library Info**
 - **Name:** Library name.
 - **Physical slot position:** Physical slot number.
 - **Logical slot position:** Logical slot number.
 - **Cleaning:** This area shows the drive cleaning statistics:
 - **Date Last Cleaned:** Date on which the drive last cleaning operation was performed.
 - **Time Since Last Cleaning:** Time elapsed since the last automatic or manual cleaning operation.
 - **Times Cleaned:** Number of times the drive has been cleaned.
 - **Data Transferred Since Last Clean:** Amount of data read or written since the last cleaning operation.
 - **Time In Use Since Last Clean:** Duration of time the drive has been used for read or write operations since the last cleaning operation.
 - **Soft Errors Since Last Clean:** Number of read or write errors reported since the last cleaning operation.
 - **Statistics**
 - **Total Data Written:** Total amount of data written using the drive.
 - **Total Data Read:** Total amount of data read using the drive.
 - **Total Write Errors:** Total number of write errors reported.
 - **Total Read Errors:** Total number of read errors reported.
 - **Time Writing:** Time taken for writing operation.
 - **Time Reading:** Time taken for reading operation.
- 4 Click **OK** to close the Details dialog box.
- 5 The slot status dialog box includes the information under the following tabs:
- **Slot**
 - **Library Info**
 - **Name:** Library name.
 - **Physical slot position:** Physical position of the slot.
 - **Logical slot position:** Logical position of the slot.
 - **Locked by:** This option is displayed if a media is present in the slot or the slot is reserved.
 - **Media:** This section is similar to the Media section for the drives described earlier in this section.
- 6 Click **OK** to close the Details dialog box.

Configuring physical tape drives

To configure options for a tape drive:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the drive to configure.
- 3 Click **Configure**.
- 4 In the **Edit Drive** dialog box, configure the settings in the following tabs:
 - **NDMP Configuration:** The NDMP parameters are available here. The default settings for any parameter on this tab must not be changed unless otherwise advised by Quest Support. For more information about these settings, see [NDMP settings](#).
 - **Configuration:** Set the required parameters here. The default settings for any parameter on this tab must not be changed unless otherwise advised by Quest Support. For more information about these settings, see [General settings for tape drives](#).
 - **Performance:** For more information about these settings, see [Drive performance settings](#).
 - **Statistics:** For more information about these settings, see [Statistics collection settings](#).
 - **SCSI Configuration:** The default timeout value for the different types of SCSI commands is set to zero, which correspond to the following intervals:
 - Fast SCSI Commands – 300 seconds
 - Slow SCSI Commands – 900 seconds
 - Very Slow SCSI Commands – 3 hours

An error is logged if a command execution does not complete within the specified interval. You should not change the default timeout for any of the SCSI commands unless otherwise advised by Quest Support.
 - **Generic Cleaning:** To configure generic cleaning settings for a tape drive, select this option. For more information about these settings, see [Generic cleaning settings](#).
- 5 Click **OK** to save the settings and close the dialog box.

Optimal transfer buffer size

The following table provides examples of the optimal values that can be used for some drive types.

Table 91. Optimal transfer buffer size of different drive types

Drive type	Optimal transfer buffer size (in KiB)
Fast Modern Tape Drives For example, LTO series, SDLT, and SAIT	65537 (64MiB + 1KiB)
Medium Speed Tape Drives For example, DLT8000, DLT7000, and AIT-3	32769 (32MiB + 1KiB)
Older Professional Tape Drives For example, DLT2000, DLT4000, and AIT-2	16385 (16MiB + 1KiB)
Older Low Capacity Low-end Drives For example, EXB-8505, AIT-1, and DAT	8193 (8MiB + 1KiB)

Checking the status of a tape drive

To check the status of an offline tape drive:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the drive.
- 3 Click **Check**, and then in the confirmation dialog box, click **OK**.
If the device is operational, its status is changed to “**Available**.”

Changing the status of a tape drive

To change the drive status to online or offline:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the drive.
- 3 Click **Online** or **Offline**:
 - If the device is offline, click **Online** to bring it back online.
 - If the device is online, click **Offline** to mark it offline.

This operation does not physically take the device offline; it just makes the device unavailable for use in NetVault.

Configuring cleaning slots

Before loading cleaning tapes, you must configure the slots that can hold the cleaning media. You can configure cleaning slots from Edit Device page. For more information, see [Modifying a library](#).

i | **NOTE:** If cleaning slots are not defined, cleaning tapes without barcodes are loaded every time the library is initialized.

Configuring the cleaning lives option

NetVault lets you specify how many times a cleaning tape can be used for drive cleaning operations. You can configure the cleaning lives property from the **Manage Devices** page.

i | **NOTE:** To configure the cleaning lives option, a cleaning tape must be placed in the designated slot.

To configure the cleaning lives for a tape:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the slot.
- 3 Click **Life**.
- 4 In the **Lives** box, enter or select the number of times the tape can be used for cleaning a drive.
- 5 Click **OK** to save the settings and close the dialog box.

Configuring automatic cleaning options for a tape drive

You can set up automatic cleaning routines for tape drives from the **Manage Devices** page. NetVault lets you use the following options to define a cleaning routine: number of days, amount of data transferred, hours of use, and number of read or write errors.

To configure automatic cleaning options for a tape drive:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the drive.
- 3 Click **Clean Properties**.
- 4 In the **Cleaning Interval** dialog box, configure the following options.

Table 92. Drive Cleaning Options

Option	Description
Days	To perform drive cleaning after every x days, select the Apply check box to the left, and type or select the value.
Data Transferred	To perform drive cleaning after every x GiB of data transfer, select the Apply check box to the left, and type or select the value.
Hours of Use	To perform drive cleaning after every x hour, select the Apply check box to the left, and type or select the value.
Soft Read/Write Errors	To perform drive cleaning after every x*100 soft read/write errors, select the Apply check box to the left, and type or select the value.

- 5 Click **OK** to save the settings and close the dialog box.

Manually submitting a drive cleaning request

You can manually submit a drive cleaning request from the **Tape Drive Management** page.

To manually submit a drive cleaning job:

- 1 In the Navigation pane, click **Manage Devices**.
 - 2 In the list of devices, locate and click the drive.
 - 3 Click **Clean**.
- After the task is completed, a message is displayed.

Unloading a tape

To unload a tape:

- 1 In the Navigation pane, click **Manage Devices**.
 - 2 In the list of devices, locate and click the drive.
 - 3 Click **Unload**.
- After unload request is sent successfully, a message is displayed.

In a library, the tape is moved to an available slot, while in a standalone drive the tape is ejected.

Loading a tape

To load a tape:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of slots, select the slot that contains the tape, and click **Load**.

After the load request is sent successfully, a message is displayed.

The tape is loaded in an available drive.

Marking media for re-use

To manually mark a piece of media for re-use:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of drive/slot, click the drive/slot in which the piece of media resides, and click **Re-use**.
- 3 In the confirmation dialog, click **OK**.
- 4 When a piece of media is manually marked for re-use, NetVault will retain its media label and group associations. To reuse such media, you must set the **Reuse Media** option on the Target tab to either **Any** or **With Same Group Label as Target Media**. NetVault will overwrite the existing data on the media when they are re-used.

i | **NOTE:** The media re-use option does not work for the NetVault SmartDisk devices. The NetVault SmartDisk devices use a process called Garbage Collection which finds and removes the chunks that are no longer in use from the chunk store and reclaims the disk space.

Removing a tape drive

You can use the following procedure to remove a tape drive that is no longer required.

When you remove a tape drive, the information about the tape that is loaded is not deleted from the NetVault Database. You can use the tape on any other drive that supports the media type. Scanning is not required if you use the tape in the same NetVault domain. In a different NetVault Domain, you must scan the tape to access the backups.

To remove a tape drive:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the library that contains the drive.
- 3 Click **Modify**.
- 4 On the **Edit Device** page, click the drive to be removed.
- 5 In the Selected Drives area, click **Remove Drive**.

-OR-

Click the applicable drive and then click **Remove**.

- 6 Click **Save**.

Adding shared devices

Adding shared drives can only be performed using tree view of the **Manage Devices** page.

Drive sharing requires an infrastructure where multiple machines can establish direct paths to the devices, such as in a switched Fiber Channel environment. In such environments, you can share the devices with multiple SmartClients for better resource utilization through direct access. The shared drives can be controlled by multiple machines (NetVault Server, Client, or Filers), but the library arm changer remains under the control of one NetVault Client. Each device that is to be shared by more than one NetVault Client uses a Dynamically Shared Device (DSD) license.

The following sections describe the procedure for adding shared devices in NetVault domain. Before you proceed, make sure all the target NetVault Clients are added to the NetVault Server.

IMPORTANT: Do not execute any changes for a drive while it is actively being used. All jobs that use the drive must be inactive before any changes are implemented.

This section includes the following topics:

- [Adding shared drives to non- shared library using the semi- automatic method](#)
- [Adding shared drives to non-shared library manually](#)
- [Sharing standalone drives](#)

Adding shared drives to non- shared library using the semi- automatic method

To add the shared drives to a non-shared library using the semi-automatic method:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the applicable library.
- 3 Click **Modify**.
- 4 On the **Edit Device** page, locate, and click the applicable library.
- 5 Click **Scan for shared drives**.

i | **NOTE:** At this time NetVault probes all NetVault Clients in the clients list for devices that match the serial numbers of those that have already been added.

- 6 Click **Save**.

The library will be added and listed in the Library Configuration area on the **Edit Device** page. The status will change to online when the library initialization process completes.

Adding shared drives to non-shared library manually

To manually add the shared drives to a non-shared library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the applicable library.
- 3 Click **Modify**.

- 4 On the **Edit Device** page, click **Show Drives** to view the NetVault Server or SmartClient node to which the device is attached.
- 5 Click the applicable drive type, and then click **Open** to display the available drives.

i | **NOTE:** When selecting the drives to be shared, ensure that the drives are assigned to the appropriate bay (data transfer element address). To obtain the correct number for each drive, refer to the relevant Library Operations or User's Guide.

- 6 Click the drive to be shared, and click **Add Shared**.

i | **NOTE:** You cannot select the drive that is already present in the bay under Library Configuration area.

-OR-

Click the drive to be shared, and click **Select**.

i | **NOTE:** An error message appears if you select the drive that is already added to bay. Remove the drive before adding to different bay.

-OR-

Click and drag the drive to the desired bay in the Library Configuration area. Add drive to bay dialog box appears. Click **Select** or **Add Shared**.

-OR-

If the desired bay is empty. Click and drag the drive to the Selected Drives area. Add drive to bay dialog box appears. Click **Select** or **Add Shared**.

- 7 Increase or change the bay number.
- 8 Repeat steps 6 and 7 for each drive that will be shared.
- 9 Click **Save**.

The library will be added and listed in the Library Configuration area on the **Edit Device** page. The status will change to online when the library initialization process completes.

Sharing standalone drives

To add the standalone drives as shared:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the applicable library.
- 3 Click **Modify**.
- 4 On the **Edit Device** page, click **Show Drives** to view the NetVault Server or SmartClient node to which the device is attached.
- 5 Click the applicable drive type, and then click **Open** to display the available drives.
- 6 Click the drive, and click **Add Shared**. This displays the drive picture, type, and details in the Selected Drives area.

i | **NOTE:** You cannot select the drive that is already present in the bay under Library Configuration area.

-OR-

Click the drive to be shared, and click **Select**.

i | **NOTE:** An error message appears if you select the drive that is already added to the bay. Remove the drive before adding to different bay.

-OR-

Click and drag the drive to the desired bay in the Library Configuration area. Add drive to bay dialog box appears. Click **Select** or **Add Shared**.

-OR-

If the desired bay is empty. Click and drag the drive to the Selected Drives area. Add drive to bay dialog box appears. Click **Select** or **Add Shared**.

7 Configure the applicable drive parameters. For more information, see [Configuring physical tape drives](#).

8 Click **Save**.

The drive will added and listed in the Library Configuration area on the **Edit Device** page. The status will change to online when the drive initialization process completes.13.0

Managing storage media

- Viewing storage summary
- Viewing disk storage details
- Viewing tape storage details
- Managing tape storage media
- Managing savesets

Viewing storage summary

To view storage summary:

- 1 In the Navigation pane, click **Explore Storage**.
- 2 On the **Explore Storage** page, you can view the following information.

Figure 26. Explore Storage page

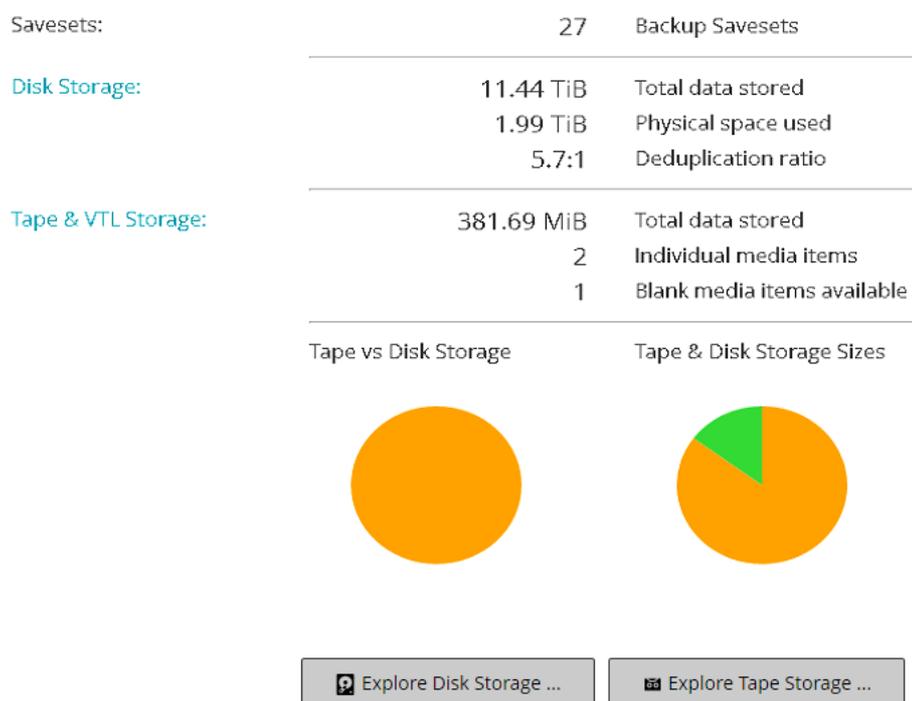


Table 93. Explore Storage page

Item	Description
Saveset	This area shows the total backup savesets stored in disk-based and tape-storage devices.
Disk Storage	This area shows the following information: <ul style="list-style-type: none"> • Total data stored: Total data stored in disk-based backup devices (total data size without deduplication). • Physical space used: Total physical space used by backups stored in disk-based backup devices (total data size with deduplication). • Deduplication ratio: The ratio of size of data before deduplication to the size of data after deduplication.
Tape & VTL Storage	This area shows the following information: <ul style="list-style-type: none"> • Total data stored: Total data stored in physical and virtual tape media. • Individual media items: Number of media items used by the backups stored in physical and virtual tape media • Blank media items available: Number of blank media items available on the added physical and virtual tape devices.
Tape vs Disk Storage (Pie chart)	The pie chart shows the total data stored in tape- and disk-based devices. <ul style="list-style-type: none"> ■ Total data stored in disk-based backup devices. ■ Total data stored in physical and virtual tape media.
Tape & Disk Storage Sizes (Pie chart)	The pie chart shows the total data stored in tape-based devices, total data stored in disk-based devices, and physical space used by backups stored in disk-based devices. <ul style="list-style-type: none"> ■ Total data stored in disk-based backup devices (total data size without deduplication). ■ Total physical space used by backups stored in disk-based backup devices (total data size with deduplication). ■ Total data stored in physical and virtual tape media.

- 3 To explore disk or tape storage repository, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

Viewing disk storage details

To view disk storage details:

- 1 On the **Explore Storage** page, click **Explore Disk Storage**.
- 2 On the **Explore Disk Storage** page, you can view the repository table.

The page displays the following information:

- **Repository Name:** The name of the device or storage container.
- **Type:** Device type (for example, NetVault SmartDisk, Quest DR Series system, or Data Domain system).
- **Record Count:** Number of data records stored in the device.
- **Saveset Count:** Number of savesets stored in the device.

- **Space Free:** Amount of space available.
- **Space Used:** Amount of space used.
- **Deduplication ratio:** The ratio of size of data before deduplication to the size of data after deduplication.

3 By default, the table is sorted by repository name.

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

4 You can use the Search option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string.

5 To view the details of a particular repository, select the item in the repository table, and click **Explore Repository**.

6 On the **Explore Disk Storage Repository** page, you can view the following information:

- **Repository summary:** This area shows the following information:
 - **Repository name:** The name of the device or storage container.
 - **Data stored:** Total data stored in the device (total data size without deduplication).
 - **Physical space used:** Total physical space used by backups stored in the device (total data size with deduplication).
 - **Space available:** Free space available in device.
 - **Deduplication ratio:** The ratio of size of data before deduplication to the size of data after deduplication.
 - **Storage data type:** Pie chart showing the percentage of storage used by various plug-ins.
- **Saveset table:** The Saveset table lists all backups stored in the repository. It shows the saveset creation date, saveset name, saveset size, job ID number, instance number, and phase number.

By default, the table is sorted by saveset date. You can sort the table by one or more columns, as required. You can also use the Search option to filter the table data and quickly find entries that contain the specified string in any column values.

To filter the saveset list, click **Filter Options**, and set the filter criteria:

- To view savesets created for a particular client, click the **Client** list, and select the client in the list.

To view savesets created during a particular period, click the **Saveset Date** list, and select one of the following options: Last 24 hours, Last Week Last Month, Last 6 Months, Last Year, or Any

The **Explore Disk Storage Repository** page loads a maximum of 5000 records initially. The total number of records fetched is displayed at the lower-right corner of the table.

You can click **Load more** to load the next set of records, if available. Each load action fetches a maximum of 5000 records. This button is disabled if there are no more records to load.

7 To perform a storage-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page

Viewing tape storage details

To view tape storage details:

- 1 On the **Explore Storage** page, click **Explore Tape Storage**.
- 2 On the **Explore Tape Storage** page, you can view the media table.

The page displays the following information:

- **Label:** Media label.
- **Group:** Media group label.
- **Barcode:** Media barcode.
- **Library:** Name of the library.
- **Record Count:** Number of data records stored in the tape.
- **Saveset Count:** Number of savesets stored in the tape.
- **Space Free:** Amount of free space available.

i | **NOTE:** The amount of free space available on a tape is not calculated in NetVault. This information is obtained from the tape drive and displayed on the **Explore Tape Storage** page.

- **Space Used:** Amount of space used.
- **Online:** Tape is online (●) or offline (●).

- 3 By default, the table is sorted by media label.

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

- 4 You can use the Search option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string.
- 5 To view the details of a particular tape, select the item in the list, and click **Explore Media**.
- 6 On the **Explore Tape Media Item** page, you can view the following information:

- **Tape summary:** This area shows the following information:
 - **Barcode:** Media barcode.
 - **Label:** Media label.
 - **Group:** Media group label.
 - **Library:** Name of the library where the tape resides.
 - **Offsite Location:** Off-site location, if specified.
 - **Media Type:** Disk file (virtual tape) or physical tape.
 - **Data Stored:** Total data stored in the tape.
 - **Space Available:** Amount of free space available.

i | **NOTE:** The amount of free space available on a tape is not calculated in NetVault. This information is obtained from the tape drive and displayed on the **Explore Tape Storage** page.

- **Reuse Policy:** Tape can be reused or not.
- **Date Last Written:** Date on which the last write operation was performed.

- **Date Last Read:** Date on which the read operation was performed.
- **Times Reused:** The number of times the tape has been reused.
- **Read Errors:** Number of read errors.
- **Write Errors:** Number of write errors.
- **Usable:** Tape is usable or unusable.
- **Read-Only:** Tape is marked read-only or not.
- **Media Online:** Tape is online or offline.
- **Media Usage:** Pie chart showing the amount of storage used by various savesets.
- **Saveset table:** The Saveset table lists all backups stored in the repository. It shows the saveset creation date, saveset name, saveset size, job ID number, instance number, and phase number.

By default, the table is sorted by saveset date. You can sort the table by one or more columns, as required. You can also use the Search option to filter the table data and quickly find entries that contain the specified string in any column values.

- 7 To perform a storage-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page

Managing tape storage media

This section includes the following topics:

- [Labeling tape media](#)
- [Blanking tape media](#)
- [Scanning tape media](#)
- [Marking a tape as unusable](#)
- [Marking a tape as read-only](#)
- [Marking a tape for reuse](#)
- [Removing an offline tape](#)

Labeling tape media

Each piece of media, whether a tape cartridge or a virtual tape in a VTL, uses a label for identification. Media items can be labeled using media barcodes, system-generated strings, or user-defined strings. You can manually assign labels to blank media or automatically label the media during a backup.

By default, NetVault assigns a system-generated string to a blank media item. You can modify the configuration settings to use media barcodes as the default labels. For more information, see [Configuring general settings for Media Manager](#).

The following sections describe the procedures that you can use to assign media labels to blank tapes:

- [Labeling multiple tapes in a library in list view](#)
- [Labeling multiple tapes in a library in tree view](#)
- [Labeling a single tape in list view](#)
- [Labeling a single tape in tree view](#)

Labeling multiple tapes in a library in list view

To label multiple tapes in a library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to locate the library, and then click the corresponding Manage Library icon .
- 3 On the **Tape Library Management** page, click **Media Label**, and configure the following options.

Table 94. Tape media labeling

Option	Description
Type of Media	<p>Select the type of media that you want to label. The available options are:</p> <ul style="list-style-type: none"> • Blank: Select this check box to label any blank, non-labeled piece of media in a device that is accessible to the NetVault Server. • Other: Select this check box to label media types that do not belong to any category listed here. • NetVault 5: Select this check box to label any piece of media that was used to back up data with NetVault 5.x. • Reusable: Select this check box to label reusable media items.
Type of Label	<p>Select the type of media label. The available options are:</p> <ul style="list-style-type: none"> • Barcode: To use media barcodes as the media labels, select this option. • Machine and Date: To use a system-generated string as the media label, select this option. This string consists of the NetVault Server Name, the current date, and a seed number. • User Defined: To assign user-defined label, select this option, and provide the following details: <ul style="list-style-type: none"> – Label: Specify the string that you want to use as the media label. A label can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. NetVault does not support a “%” character in the string. There is no length restriction on media and group labels. However, the combined display range for the media label, barcode and group label is 100 characters. Therefore, a maximum of 40 to 50 characters is recommended for the media and group labels. – Seed: To identify individual media items, a sequential number is added to the user-defined string. The option defines the initial value for the sequence. This value is increased by one for each item. The default value for this option is one.
Group Label	<p>To add the media items to a group, select the group label in the list. If you want to create a group label, type the string.</p> <p>A label can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. Group labels are case-insensitive.</p>
All Media in List	<p>To label all media items in the selected library, select this check box.</p>
Media to Label	<p>To label specific media items, select the individual media items in the list.</p> <p>To select consecutive items, hold down the Shift key while clicking with the mouse button; to select non-consecutive items, hold down the Ctrl key while clicking with the mouse button.</p>

- 4 Click **OK** to save the settings.

Labeling multiple tapes in a library in tree view

To bulk label media items or media groups, follow these steps:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the target library.
- 3 Click **Bulk Label**.
- 4 In the **Bulk Label Media** window, configure the following parameters:

i | **NOTE:** The bulk media labeling and grouping features are not applicable to NetVault SmartDisk Device.

- **Type of Media to Label:** Select the type of media to be labeled:
 - **Blank:** Select this check box to label any blank, non-labeled piece of media in a device that is currently accessible to the NetVault Server.
 - **Other:** Select this check box to bulk label media types that do not belong to any category listed here.
 - **NetVault 5:** Select this check box to label any piece of media that was used to back up data with NetVault 5.x.
 - **Reusable:** Select this check box to label a piece of media that is marked as reusable.
- **Type of Label:** Select the type of media label:
 - **Barcode:** To use media barcodes as the media labels, select this option.
 - **Machine and Date:** To generate a string using the NetVault Server Name, Current Date, and a Seed Number, select this option.
 - **User:** To assign a user-defined label, select User, and enter the following details:
 - **Label:** Enter the string to be used as media label. NetVault does not support a “%” character in the string. A media label can contain alphanumeric and non-alphanumeric characters, but it cannot include Non-English characters.

i | **NOTE:** There is no length restriction on media and group labels. However, the combined display range for the media label, barcode, and group label is 100 characters. Therefore, a maximum of 40-50 characters is recommended for the media and group labels.

- **Seed:** For the identification of individual media items, NetVault adds a sequential number to the user-defined string. The Seed parameter provides the initial value for the sequence. This value is incremented by one for each media item. The default value for Seed is one.
- **Group Label:** To add the media items to a group, select the group label in the list. If you want to create a new group label, enter the string. The group label can contain alphanumeric and non-alphanumeric characters, but it cannot include Non-English characters.

Group labels are case-insensitive. NetVault does not distinguish between uppercase and lowercase letters in a media group label while creating and using media groups.

- **Select Media Items:** Select the applicable method:
 - **Select Individual Media Items:** In the media list, select the individual media items for labeling. To select consecutive items, hold down the Shift key while clicking with the mouse button; to select nonconsecutive items, hold down the Ctrl key while clicking with the mouse button.
 - **Select All Media Items:** To label all media items in the list, select the **All Media in List** check box.
- **Enter ‘LABEL’ to Confirm Request:** To confirm bulk labeling, enter LABEL in the box. This string is case insensitive.

- 5 Click **OK**.

Labeling a single tape in list view

To label a single tape:

- 1 In the Navigation pane, click **Manage Devices**, and in the list of devices, click the icon  to open the library that contains the tape.
- 2 If the tape is loaded in a drive, click the drive or the corresponding Manage Drive icon .

Otherwise, click the **Slots** link or the corresponding Manage Slots icon  to open the Slot Browser. In the list of slots, select the slot that has the tape.

- 3 Click **Label**, and in the **Media Label** dialog box, configure the following options.

Table 95. Labeling a single tape (list view)

Option	Description
Label	Specify a label for the tape. A label can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. NetVault does not support a “%” character in the string. There is no length restriction on media and group labels. However, the combined display range for the media label, barcode and group label is 100 characters. Therefore, a maximum of 40 to 50 characters is recommended for the media and group labels.
Group Label	To add the tape to a group, select the group label in the list. If you want to create a group label, type the string. A label can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. Group labels are case-insensitive.
Offsite Location	Specify the off-site location for the tape.

- 4 Click **OK** to save the settings and close the dialog box.

Labeling a single tape in tree view

To label a single tape:

- 1 In the Navigation pane, click **Manage Devices**, and in the list of devices, open the library that contains the tape.
- 2 In the list of slots, click the slot that has the tape and then click **Properties**.
- 3 In the **Media Label** dialog box, configure the following options.

Table 96. Labeling a single tape (tree view)

Option	Description
Media Label	<p>Specify a label for the tape.</p> <p>A label can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. NetVault does not support a “%” character in the string.</p> <p>There is no length restriction on media and group labels. However, the combined display range for the media label, barcode and group label is 100 characters. Therefore, a maximum of 40 to 50 characters is recommended for the media and group labels.</p>
Group Label	<p>To add the tape to a group, select the group label in the list. If you want to create a group label, type the string.</p> <p>A label can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. Group labels are case-insensitive.</p>
Offsite Location	<p>Specify the off-site location for the tape.</p>
Tape Format	<p>Select MTF on Windows and CPIO on Linux/UNIX.</p>

- 4 Select **Read Only** or **Unusable** check box as per your requirement.
- 5 Click **OK** to save the settings and close the dialog box.

Blanking tape media

Blanking a tape deletes or erases the backup data residing on the tape. It removes the NetVault header from the tape and deletes the media label and removes any group association. Blanking a tape also removes the indexes for backups stored on the selected tapes from the NetVault Database.

After blanking, a tape becomes available to NetVault for storing future backups.

To purposely destroy the data that is stored on a tape, you must blank it from NetVault and have its data securely removed by tools that are designed for such purposes.

The following sections describe the procedures that you can use to blank tape media:

- [Blanking multiple tapes in a library in list view](#)
- [Blanking multiple tapes in a library in tree view](#)
- [Blanking a single tape in list view](#)
- [Blanking a single tape in tree view](#)

Blanking multiple tapes in a library in list view

To blank multiple tapes in a library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to locate the library, and then click the corresponding Manage Library icon .
- 3 On the **Tape Library Management** page, click **Bulk Blank**, and configure the following options.

Table 97. Bulk blank

Option	Description
All media in list	To blank all media items in the Media that can be blanked list, select this check box.
Media that can be blanked	To blank specific media items, select the items in this list. To select consecutive items, hold down the Shift key while clicking with the mouse button; to select non-consecutive items, hold down the Ctrl key while clicking with the mouse button.
Password	Type the password for the NetVault Server. If no password is set for the NetVault Server, provide the system's root or administrator password.
Enter 'BLANK' to confirm request	To confirm, type BLANK (case-insensitive) in this box.

- 4 Click **OK**.

Blanking multiple tapes in a library in tree view

To blank multiple tapes in a library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the library in which the media items reside.
- 3 Click **Bulk Blank**.
- 4 In the **Bulk Blank Media** dialog box, configure the following parameters:
 - **Select Media Items**
 - **Select Individual Media Items:** In the media list, select the individual media items for blanking. To select consecutive items, hold down the Shift key while clicking with the mouse button; to select non-consecutive items, hold down the Ctrl key while clicking with the mouse button.
 - **Select All Media Items:** To blank all media items in the list, select the **All Media in List** check box.
 - **Password:** Enter the password for the NetVault Server.
 - **Enter 'BLANK' to Confirm Request:** To confirm bulk blanking, enter BLANK in the box. This string is case-insensitive.
- 5 Click **OK**.

i | **NOTE:** This procedure cannot be used to blank an NetVault SmartDisk Device.

Blanking a single tape in list view

To blank a single tape:

- 1 To blank a tape from the **Manage Devices** page:
 - a In the Navigation pane, click **Manage Devices**, and in the list of devices, click the icon  to open the library that contains the tape.
 - b If the tape is loaded in a drive, click the drive or the corresponding Manage Drive icon .

Otherwise, click the **Slots** link or the corresponding Manage Slots icon  to open the Slot Browser. In the list of slots, select the slot that has the tape.

- 2 To blank a tape from the **Explore Storage** page:
 - a In the Navigation pane, click **Explore Storage**.
 - b Click **Explore Tape Storage**. In the list of media items, select the tape, and click **Explore Media**.
- 3 Click **Blank**, and then in the confirmation dialog box, click **OK**.

Blanking a single tape in tree view

To blank a single tape:

- 1 In the Navigation pane, click **Manage Devices**, click the slot that has the tape.
- 2 Click **Blank**, and then in the confirmation dialog box, click **OK**.

Scanning tape media

You can use the **Scan** method to query all backups stored on a tape and import those backups that are not indexed in the given NetVault Server's database. You can also use the **Scan** method to import "foreign" tapes into the NetVault Database.

To import and restore NetVault Database backups to an alternate backup server, the alternate NetVault Server must have the same NetVault Machine Name as the original server that performed the backups. The amount of time it takes to scan the tapes depends on the number of backups that need to be imported and the size of the backup indexes. The scanning process does not read the data on the tape; this process skips between the start and end of backups to read the index for each backup saveset.

You can scan indexes that are generated with the same or previous versions of NetVault. You cannot scan indexes generated with a newer version of NetVault on an earlier version of the server if the servers do not use the same index version. If an index version is not supported, the index is not imported and a message is generated in the logs.

When a tape is swapped between libraries, removed from the NetVault Server, or loaded on a device that is controlled by a different NetVault Server, the NetVault Server is unable to find any information about that tape in the NetVault Database. In this case, the first scan retrieves header information from the tape and adds this information to the NetVault Database and then marks the tape as "foreign". NetVault cannot import backups stored on a "foreign" tape until you scan the tape again.

The following sections describe the procedures that you can use to scan tape media:

- [Scanning all tapes in a library in list view](#)
- [Scanning all tapes in a library in tree view](#)
- [Scanning a single tape in list view](#)
- [Scanning a single tape in tree view](#)

Scanning all tapes in a library in list view

To scan all tapes in a library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, click the icon  to locate the library, and then click the corresponding Manage Library icon .

- 3 On the **Tape Library Management** page, click **Scan All**. (If NetVault fails to start the process, click **Force Scan**.)
- 4 In the **Scan Device** dialog box, you can configure the following option:

- **Minimum life for imported backups:** This option specifies the minimum life for backups that are imported to the NetVault Database. This option only applies to the backups that are not available in the NetVault Database.

The default value for this option is seven days. You can change the default setting by modifying the Media Manager settings. For more information, see [Configuring general settings for Media Manager](#).

To change the minimum life setting for the current session, type or select a new value. The minimum life setting is specified in number of days.

Depending on the value set for this option, the retirement time for imported backups is modified as follows:

- If a backup is already retired, its retirement time is set to the specified minimum life. If you specify zero, the retirement time for the retired saveset is set to one hour.
- If a backup is scheduled to retire before the specified period, its retirement time is set to the specified minimum life.
- If a backup is scheduled to retire after the specified period, its retirement time remains unchanged. For such backups, the backup life setting determines the retirement time.

Click **Scan** to start the scanning process and close the dialog box.

Scanning all tapes in a library in tree view

To scan all tapes in a library:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the library.
- 3 Click **Scan**. (If NetVault fails to start the process, click **Force Scan**.)
- 4 In the confirmation dialog box, click **OK**.

Scanning a single tape in list view

To scan a single tape:

- 1 Use one of the following methods to select the tape that you want to scan:
 - To select a tape from the **Manage Devices** page:
 - a In the Navigation pane, click **Manage Devices**, and in the list of devices, click the icon  to open the library that contains the tape.
 - b If the tape is loaded in a drive, click the drive or the corresponding Manage Drive icon .Otherwise, click the **Slots** link or the corresponding Manage Slots icon  to open the Slot Browser. In the list of slots, select the slot that has the tape.
- To select a tape from the **Explore Storage** page:
 - a In the Navigation pane, click **Explore Storage**.
 - b Click **Explore Tape Storage**. In the list of media items, select the tape, and click **Explore Media**.
- 2 Click **Scan**.

3 In the **Scan Device** dialog box, you can configure the following option:

- **Minimum life for imported backups:** This option specifies the minimum life for backups that are imported to the NetVault Database. This option only applies to the backups that are not available in the NetVault Database.

The default value for this option is seven days. You can change the default setting by modifying the Media Manager settings. For more information, see [Configuring general settings for Media Manager](#).

To change the minimum life setting for the current session, type or select a new value. The minimum life setting is specified in number of days.

Depending on the value set for this option, the retirement time for imported backups is modified as follows:

- If a backup is already retired, its retirement time is set to the specified minimum life. If you specify zero, the retirement time for the retired saveset is set to one hour.
- If a backup is scheduled to retire before the specified period, its retirement time is set to the specified minimum life.
- If a backup is scheduled to retire after the specified period, its retirement time remains unchanged. For such backups, the backup life setting determines the retirement time.

Click **Scan** to start the scanning process and close the dialog box.

Scanning a single tape in tree view

To scan a single tape:

- 1 In the Navigation pane, click **Manage Devices**.
- 2 In the list of devices, locate and click the slot that has the tape.
- 3 Click **Scan**. (If NetVault fails to start the process, click **Force Scan**.)

Marking a tape as unusable

If a piece of media is damaged or not suitable for use, you can mark it as “unusable” so that it is not selected for any job. You can set this property from the **Explore Tape Media Item** page.

To mark a tape as unusable:

- 1 In the Navigation pane, click **Explore Storage**.
- 2 Click **Explore Tape Storage**. In the list of media items, select the tape, and click **Explore Media**.
- 3 To mark the tape as unusable, click **Mark Unusable**.
- 4 To change this property, select the tape, and click **Mark Usable**.

Marking a tape as read-only

You can mark a tape as “read-only” to protect it from further writes.

There are two ways to enable this property for a tape:

- You can set the **Protect Media from Further Writes after Backup** option in the Target Set to enable write-protection for a tape after the backup completes. For more information, see [Configuring media sharing options](#).
- Alternatively, you can set the read-only property from the **Explore Tape Media Item** page. The instructions are provided in this section.

To mark a tape as read-only:

- 1 In the Navigation pane, click **Explore Storage**.
- 2 Click **Explore Tape Storage**. In the list of media items, select the tape, and click **Explore Media**.
- 3 To mark the tape as read-only, click **Mark Read Only**.
- 4 To change this property, select the tape, and click **Mark Writable**.

i | **IMPORTANT:** A piece of media can be marked as “read-only” to stop further writes if a SCSI error occurs during a write operation. When this error occurs, check for hardware errors. If no tape or media error is found, set the tape as “writable.”

Marking a tape for reuse

A piece of media is automatically marked for reuse when the last saveset stored on it is retired. You can also manually set this property from the **Explore Tape Media Item** page. NetVault overwrites any existing data on the tapes when they are reused.

When a piece of media is manually marked for re-use, NetVault retains its media label and group associations. To reuse such media, you must set the **Reuse Media** option in the Target Set to one of the following: **Any** or **With the same group label as target media**.

To manually mark a tape for reuse:

- 1 In the Navigation pane, click **Explore Storage**.
- 2 Click **Explore Tape Storage**. In the list of media items, select the tape, and click **Explore Media**.
- 3 Click **Reuse**, and then in the confirmation dialog box, click **OK**.

Removing an offline tape

You can use the following procedure to remove an offline tape.

When you remove an offline tape, the information about that tape is deleted from the NetVault Database; the backups stored on the tape are not deleted. To use the backups stored on the tape, you must scan the tape and import the media information into the NetVault Database.

To remove an offline tape:

- 1 In the Navigation pane, click **Explore Storage**.
- 2 Click **Explore Tape Storage**. In the list of media items, select the tape that you want to remove, and click **Explore Media**.

Only offline tapes can be removed from NetVault.

- 3 Click **Remove**, and then in the confirmation dialog box, click **OK**.

Managing savesets

This section includes the following topics:

- [Viewing saveset details](#)
- [Configuring saveset expiry options](#)
- [Expiring all savesets on disk-based storage](#)
- [Deleting savesets from disk-based storage devices](#)

- [Deleting savesets from tape-based storage devices](#)

Viewing saveset details

To view the details of a saveset:

- 1 In the Navigation pane, click **Explore Storage**.
- 2 If the saveset is stored in a disk-based storage device, click **Explore Disk Storage**. In the repository table, select the device, and click **Explore Repository**.
- 3 If the saveset is stored in a physical or virtual tape, click **Explore Tape Storage**. In the list of media items, select the tape, and click **Explore Media**.
- 4 In the saveset list, select the target saveset, and click **Examine Saveset**.
- 5 On the Saveset Information page, you can view the following details:
 - **Job:** Job ID number and instance ID number
 - **Title:** Job name
 - **Tag:** Tag assigned to the saveset
 - **Server:** Name of the NetVault Server
 - **Client:** Name of the NetVault Client for which the saveset is generated.
 - **Plugin:** Name of the plug-in used to create the saveset.
 - **Date:** Saveset creation date.
 - **Expires At:** Saveset expiration date and time.
 - **Incremental:** Incremental Backup or not.
 - **Archive:** Archive option was selected or not.
 - **Size:** Saveset size.

- 6 To view the list of media items, click **Media list**.

In the dialog box that appears, 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.

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

Configuring saveset expiry options

NetVault supports generation- and time-based retirement methods for backups. You can specify these options in the Backup Advanced Options Set. If the backups are stored on tape media, you can also use the **Change Expiry** method to set or change the expiry date or generation count later.

This section describes how to set the retirement age or maximum generation count for one or more savesets on physical or virtual tape media by using the **Change Expiry** method. For more information about the **Backup Life** option, see [Setting backup retirement options](#).

If a backup has any dependent backups, you can choose to do the following:

- Defer the removal of savesets until all dependent backups are ready for retirement

- Force removal of savesets according to their schedule

For more information about backup retirement methods and retirement rules, see [Backup retirement](#).

To configure saveset expiry options:

- 1 In the Navigation pane, click **Explore Storage**.
- 2 Click **Explore Tape Storage**.
- 3 In the list of media items, select the tape, and click **Explore Media**.
- 4 In the saveset list, select one or more target savesets, and click **Change Expiry**.
- 5 On the **Media Management - Change Saveset Expiry Date** drawer, configure the options described in the following table.

Table 98. Change saveset expiry period

Option	Description
Change Expiry Date	<p>To configure time-based retirement, select this check box, and do one of the following:</p> <ul style="list-style-type: none"> • Select the On option, and type or select the date and time in the respective boxes. • Alternatively, select the Never option to retain the backup indefinitely. <p>NOTE: In time-based retirement, the time component (HH:MM) does not represent the actual retirement time. It only represents the time due for backup retirement. The actual time of retirement is determined by the interval at which Media Manager scans the Media Database to identify the backups that it needs to retire. The default interval between two scans is 60 minutes. Thus, if the retirement time is set to 10:20, the backup is actually retired at 11:00. You can change this default setting in the mediamgr.cfg file. For more information, see Configuring default interval for backup retirement scans.</p>
Change Generation Cycle	<p>To configure generation-based retirement, select this check box, and do one of the following:</p> <ul style="list-style-type: none"> • Select the Discard after option, and in the associated box, type or select the number of Full Backups. • Alternatively, select the Never option to retain the backup indefinitely.
Force Expiry	<p>By default, if a backup has any dependent backups, its retirement is deferred until all dependent backups are ready for retirement.</p> <p>You can select this check box to retire a backup according to its retirement schedule. Forcing this behavior can cause early retirement of dependent Incremental and Differential Backups.</p> <p>To apply this rule globally to all backups, you can modify the Media Manager settings. For more information, see Configuring retirement rules for dependent backups. If the Retirement Timing Control option for Media Manager is set to Force Always, the Force Expiry option is used, regardless of state of this check box.</p>

- c Click **Apply** to save the settings.

i IMPORTANT:

- If you set both **Change Expiry Date** and **Change Generation Cycle** options for a Full Backup, the backup is retired only when both conditions are met. For example, if you set the **Discard After Full Backup Count** option to four Full Backups and the **Discard After** option to 30 days, the backup is retired after four Full Backup counts and 30 days.
- When a backup stored on a disk-based storage device (such as Quest DR Series system, NetVault SmartDisk, or Data Domain System) is retired, that backup is deleted from the device. You cannot import the deleted backup by scanning the device.

To configure saveset expiry option from Create Restore Job page

- 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.
- 2 Select the saveset to change the expiry date and time.
- 3 Click the set expiry date and time, against **Expiry** field.
- 4 Configure the options available in Media Management - Change Saveset Expiry Date page. See, [Table 98](#).

Expiring all savesets on disk-based storage

This task describes how to expire all savesets that are stored on disk-based storage.

! | **CAUTION:** Expiring all savesets removes ALL data stored on the selected device.

To expire all savesets on disk-based storage:

- 1 In the Navigation pane, click **Explore Storage**.
- 2 Click **Explore Disk Storage**.
- 3 In the repository table, select the device, and click **Explore Repository**.
- 4 On the **Explore Disk Storage Repository** page, select the saveset, and then click **Expire All**.
- 5 On the **Expire All Savesets** modal, enter the password for the device, and then enter the word EXPIRE to confirm the action.
- 6 Click **OK**.

Deleting savesets from disk-based storage devices

Deleting a saveset from a disk-based storage device involves removing its index from the NetVault Database and deleting that backup from the device. You cannot import the deleted backup by scanning the device.

If a backup has any dependent backups, you can choose to do the following:

- Defer the removal of savesets until all dependent backups are ready for retirement
- Force immediate removal of the savesets

i | **IMPORTANT:** Regardless of your selection here, if the **Retirement Timing Control** option for Media Manager is set to **Force Always**, the savesets are removed immediately. For more information about this global setting, see [Configuring retirement rules for dependent backups](#).

To delete savesets from disk-based storage devices:

- 1 In the Navigation pane, click **Explore Storage**.
- 2 On the **Explore Storage** page, click **Explore Disk Storage**. In the repository table, select the device, and click **Explore Repository**.
- 3 To delete multiple savesets:
 - a In the saveset list, select the savesets that you want to delete.
To select consecutive items, hold down the Shift key while clicking with the mouse button; to select non-consecutive items, hold down the Ctrl key while clicking with the mouse button.
 - b Click **Remove Savesets**.

- c In the **Remove Saveset** dialog box, select the **Remove all duplicates associated with selected saveset(s)** check box if you also want to remove the Duplicates of the selected savesets.
 - d Under **Select removal timing option**, specify whether you want to force immediate removal of the savesets or defer the removal until all dependent savesets have reached their retirement date:
 - **Mark for removal:** Use this option to mark the selected savesets for removal. If no dependent Incremental or Differential Backups exist for a saveset, the saveset is removed immediately. If any dependent Incremental or Differential Backups exist for a saveset, the saveset is removed only after all dependent backups have reached their retirement date.
 - **Force immediate removal:** Use this option to remove the selected savesets immediately. The dependent Incremental and Differential Backups that exist for the selected savesets are also removed immediately, even if these backups have not reached their retirement date.
 - e Click **OK**.
- 4 To examine a saveset before deleting it:
- a In the saveset list, select the saveset that you want to delete, and click **Examine Saveset**.
 - b Click **Remove**.
 - c In the **Remove Saveset** dialog box, select the **Remove all duplicates associated with this saveset** check box if you also want to remove the Duplicates of the selected saveset.
 - d Under **Select removal timing option**, specify whether you want to force immediate removal of the saveset or defer the removal until all dependent savesets have reached their retirement date:
 - **Mark for removal:** Use this option to mark the selected saveset for removal. If no dependent Incremental or Differential Backups exist for the saveset, the saveset is removed immediately. If any dependent Incremental or Differential Backups exist for the saveset, the saveset is removed only after all dependent backups have reached their retirement date.
 - **Force immediate removal:** Use this option to remove the selected saveset immediately. The dependent Incremental and Differential Backups that exist for the selected saveset are also removed immediately, even if these backups have not reached their retirement date.
 - e Click **OK**.

Deleting savesets from tape-based storage devices

Deleting a saveset from a tape-based storage device essentially involves removing its index from the NetVault Database. You can still scan the media to import the backup index into the NetVault Database, and use the backup.

If a backup has any dependent backups, you can choose do the following:

- Defer the removal of savesets until all dependent backups are ready for retirement
- Force immediate removal of the savesets

i | **IMPORTANT:** Regardless of your selection here, if the **Retirement Timing Control** option for Media Manager is set to **Force Always**, the savesets are removed immediately. For more information about this global setting, see [Configuring retirement rules for dependent backups](#).

To delete savesets from tape-based storage devices:

- 1 In the Navigation pane, click **Explore Storage**.
- 2 On the **Explore Storage** page, click **Explore Tape Storage**. In the list of media items, select the tape, and click **Explore Media**.
- 3 In the saveset list, select one or more savesets that you want to delete, and click **Remove Savesets**.

- 4 In the **Remove Saveset** dialog box, select the **Remove all duplicates associated with this saveset** check box if you also want to remove the Duplicates of the selected saveset.
- 5 Under **Select removal timing option**, specify whether you want to force immediate removal of the saveset or defer the removal until all dependent savesets have reached their retirement date:
 - **Mark for removal:** Use this option to mark the selected saveset for removal. If no dependent Incremental or Differential Backups exist for the saveset, the saveset is removed immediately. If any dependent Incremental or Differential Backups exist for the saveset, the saveset is removed only after all dependent backups have reached their retirement date.
 - **Force immediate removal:** Use this option to remove the selected saveset immediately. The dependent Incremental and Differential Backups that exist for the selected saveset are also removed immediately, even if these backups have not reached their retirement date.
- 6 Click **OK**.

Managing user and group accounts

- [About user accounts](#)
- [About user and group privileges and presets](#)
- [Creating a user account](#)
- [Modifying a user account](#)
- [Deleting a user account](#)
- [Setting a user policy](#)
- [Using Secure Mode](#)
- [Using Presets](#)
- [User privileges](#)
- [Predefined Presets](#)
- [Integrating Active Directory with NetVault](#)
- [Using Azure Active Directory as an identity provider](#)

About user accounts

The NetVault administrator can set up different user accounts, and assign appropriate privileges, group memberships, and quota to these accounts based on the functions the users need to perform.

For example, the administrator can create a user account for Database Administrator and grant this account the privileges to perform backups and restores of the database systems. Similarly, the administrator can create a user account to manage clients or storage devices.

There are two predefined user accounts in NetVault:

- **admin:** The administrator account for NetVault.
- **default:** A standard user account that can be used to perform various operations in NetVault.

These user accounts are used to perform various tasks from the NetVault WebUI and CLI. The **admin** and **default** user accounts are assigned all privileges in NetVault. By default, no password is assigned to the **admin** and **default** user accounts. To prevent unauthorized access to the NetVault Server, you can assign a secure password to these user accounts.

i | **NOTE:** Only the administrator account can create and manage user accounts in NetVault. Non-administrator user accounts can only change or reset their own passwords.

If your environment uses Active Directory (AD) or Microsoft Azure Active Directory (Azure AD), you can also integrate AD and Azure AD users with NetVault, as well as manage user groups. For more information, see [Integrating Active Directory with NetVault](#) and [Using Azure Active Directory as an identity provider](#).

About user and group privileges and presets

A privilege is a permission to perform a specific task in NetVault.

To accomplish any task in NetVault, a user must have the required NetVault privileges. If a user does not have the privileges, the request is denied. In NetVault, the Auditor Daemon (nvavp) authorizes the user requests based on the privileges granted to the user or user groups. For information about the types of user and user group privileges in NetVault, see [User privileges](#).

A Preset is a set of user privileges. Presets facilitate the task of assigning user privileges. Instead of assigning the individual privileges to each user account, the administrator can assign a Preset. When a Preset is assigned to a user, the user gets all the permissions included in that set.

NetVault includes several predefined Presets, which contain all the necessary privileges for specific user roles and user groups. The predefined Presets include the following: Administrator, Backup Administrator, Backup Operator, Media Operator, Monitor, and Helpdesk. For more information about the privileges included in these sets, see [Predefined Presets](#).

The administrator can also create user-defined presets for any additional user roles that are required in the backup environment. For more information, see [Using Presets](#).

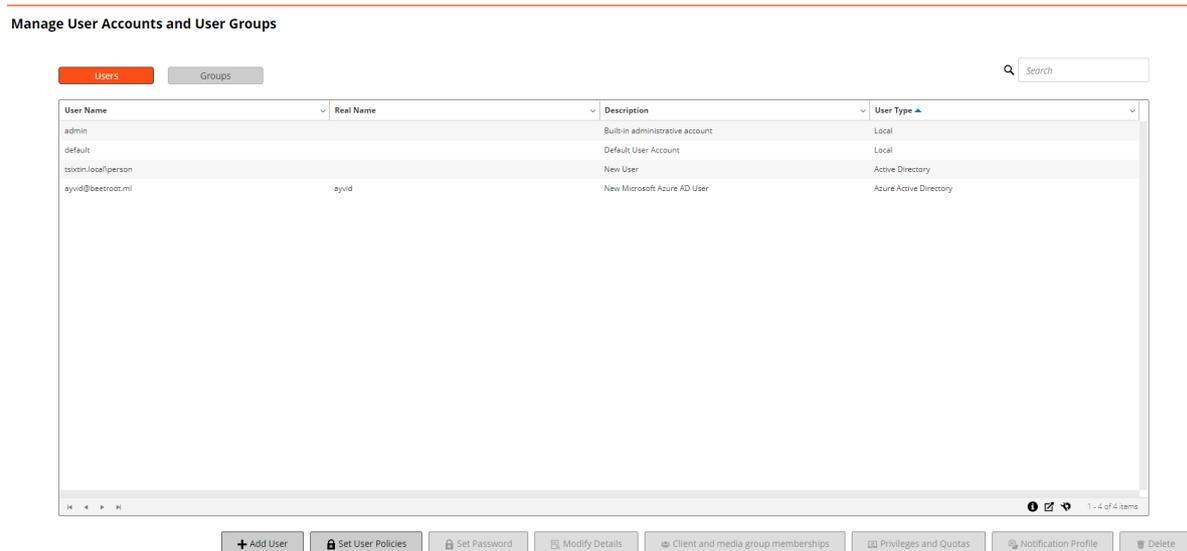
Creating a user account

User accounts can be created from the **Manage User Accounts and User Groups** page. Only the administrator can create and manage user accounts in NetVault.

To create a user account:

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, click **Add User**.

Figure 27. Manage User Accounts and User Groups page



NetVault creates a user account and assigns the default name **New User** (or “New User [n]” if the default name is in use) to this user.

By default, the table is sorted by User Name.

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower-right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

To quickly find entries containing specific field values or text in any column, you can use the Search box at the upper-right corner of the table. The data in the table is filtered as you type the search text into the box. You can also include wildcards (“?” or “*”) in the search filter string.

i | **NOTE:** The NetVault Server does not delete the user account if you exit without completing the user addition procedure. If you do not require the user account, you must manually delete it from the **Manage User Accounts and User Groups** page. For more information, see [Deleting a user account](#).

- 3 Configure the options described in the following sections:
 - [Setting user password](#)
 - [Setting a user policy](#)
 - [Configuring user details](#)
 - [Configuring memberships for a user account or user group](#)
 - [Granting privileges and quota to a user account](#)
 - [Setting up user notification profile](#)
- 4 Click **Done** to save the user details.

Setting user password

i | **IMPORTANT:** You can only assign or update the password for a local NetVault user. Password management for AD users and Azure AD users is managed by Active Directory and Microsoft Azure Active Directory, respectively.

To set or change password for a user account:

- 1 On the **Manage User Accounts and User Groups** page, click **Set Password**.
- 2 To remove the existing password and reset it to blank, select the **Reset password to blank** check box.
— or —

To set or change the user password, configure the following options.

Table 99. User password

Item	Description
Current Password	Type the current password for the user account. Leave it blank if no password is set for the account.
New Password	Type a new password for the user account. A password can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. A password can contain a maximum of 100 characters.
Confirm Password	Re-type the password for confirmation.

- 3 Click **Apply** to save the user details, and return to the user settings page.

Configuring user details

i | **IMPORTANT:** Do *not* use the @ symbol in the **User Name** field. NetVaultNetVault assumes that a name that contains @ is a domain user and automatically tries to authenticate it with Active Directory when the user logs on.

To configure the user name, contact information, and other details:

- 1 On the **Manage User Accounts and User Groups**, click **Modify Details**.
- 2 Configure the following options.

Table 100. User details

Item	Description
Identification	<p>Under Identification, provide the following details:</p> <ul style="list-style-type: none">• User Name: In User Name, type a unique name for the user account. You can assign a name based on the user group, role, or actual name. A user name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction, but a maximum of 20 characters is recommended on all platforms. The following characters are not supported in user names: " / \ : ; * ? < > ^ <p>NOTE: Do <i>not</i> use the @ symbol in the User Name field.</p> <ul style="list-style-type: none">• Real Name: In Real Name, specify the actual name of the user.
Contact Information	<p>Under Contact Information, provide the following details:</p> <ul style="list-style-type: none">• Email-1: Use this box to specify the primary email addresses for the user account.• Email-2: Use this box to specify any additional email address for the user account.• Email-3: Use this box to specify any additional email address for the user account.• Telephone: Use this box to specify the telephone number for the user account.• Cellular: Use this box to specify the mobile phone number for the user account.• Pager: Use this box to specify the pager number for the user account. <p>NOTE: The email ID configured in the Email-1 box is used for email notifications if you set up a notification profile for the user account. For more information, see Setting up user notification profile.</p>
Other Details	<p>Under Other Details, provide the following details:</p> <ul style="list-style-type: none">• Workstation: Use this box to specify the workstation name.• Description: Use this box to specify the workstation description.• Location: Use this box to specify the workstation location.• Password never expires: By default, this option is selected. To apply the password policy settings for a user account, clear this option. <p>NOTE: You can use the Password never expires option for the user accounts, if password for the selected user does not need to expire.</p>

- 3 Click **Apply** to save the user details, and return to the **Manage User Accounts and User Groups** page.

Configuring memberships for a user account or user group

To configure memberships for a user account or user group:

- 1 On the **Manage User Accounts and User Groups** page, click **Group Memberships**.
For local users, select the **Users** tab, and for groups, select the **Groups** tab.
- 2 On the **Editing Group Memberships for User** page, add or remove the client group, media group, or storage tier memberships using the following options:
 - To grant access to specific groups, select the groups in the **Not a Member Of** list, and click **Join**. The selected client groups are moved to the **Member of** list.
(To select consecutive items, hold down the Shift key while clicking with the mouse button; to select non-consecutive items, hold down the Ctrl key while clicking with the mouse button.)
 - To remove a group, select the group in the **Member of** list, and click **Leave**.
- 3 To save the group membership information for the user, and return to the user settings page, click **Apply**.

Granting privileges and quota to a user account

To grant user privileges and job and media quota for a user account:

- 1 On the **Manage User Accounts and User Groups** page, click **Privileges and Quotas**.
For local users, select the **Users** tab, and for groups, select the **Groups** tab.
- 2 Grant appropriate user privileges and job and media quota.

Table 101. User privileges and quotas

Item	Description
User Privileges	<p>To grant or revoke user privileges, do the following:</p> <ul style="list-style-type: none"> • Grant all privileges: To grant all privileges to a user account, select the User is granted ALL privileges check box. • Assign presets: To assign a predefined or user-defined preset, select the Preset in the Privileges Presets list, and click Load. A predefined Preset can be identified by a trailing ◆ symbol in its name. <p>When you assign a Preset to a user, the user is granted all the permissions included in that set. The selected privileges are moved to the Granted list.</p> <p>You can assign only one Preset. If you load a new Preset, the Granted list is overwritten with the new set of privileges.</p> <p>For more information about creating, modifying, and deleting Presets, see Using Presets.</p>
Media Quota	<p>To set the media quota for a user, configure the applicable option:</p> <ul style="list-style-type: none"> • Infinite: This option is selected by default. To allow access to infinite amount of media, use this option. • Up to: To set the quota limit, select this option. Type or select the amount of media that is available to the user. The usage quota is specified in terabytes. <p>When the quota limit is reached, the jobs submitted by the user fail. The media usage amount is calculated from the existing records in the Media Database. When a saveset is retired, the amount of media used by it is added to the available pool.</p>
Job Quota	<p>To set the job quota for a user, configure the applicable option:</p> <ul style="list-style-type: none"> • Infinite: This option is selected by default. To allow access to infinite number of jobs, use this option. • Up to: To set the quota limit, select this option. Type or enter the maximum number of jobs the user can perform. <p>When the quota limit is reached, the user is not allowed to submit any other job. The number of jobs performed by a user is derived from the existing records in the Scheduler Database. For any deleted job definitions, a user can submit an equal number of new jobs.</p>

- 3 Click **Apply** to save the group membership information for the user, and return to the **Manage User Accounts and User Groups** page.

Granting add and remove permissions

To grant a user add and remove permissions for a specific group:

- 1 Add the user to a group that contains at least two clients.
For more information, see [Configuring memberships for a user account or user group](#).
- 2 In the Navigation pane, click **Users and Groups**.
- 3 On **Manage User and Group Accounts** page, click Group Memberships.

- 4 On the **Editing Group Memberships for User** page, in the **Not a Member Of** list, click **Join** for all client groups.
- 5 Click **Apply**.
- 6 On the **Manage User and Group Accounts** page, click **Privileges and Quotas**.
- 7 In the **Add** list, select the **Clients — Add/remove clients** privilege.
- 8 Click **Add**.

You can now use this user to add a client to the group.

Setting up user notification profile

To set up notification profile for a user account:

- 1 On **Manage User Accounts and User Groups**, click **Notification Profile**.
For local users, select the **Users** tab, and for groups, select the **Groups** tab.
- 2 In the events table, open the event class and event type, and select the notification method that you want to use.

Table 102. User notification profile

Item	Description
E-mail	Select this method to send an email notification to the user when the event occurs. The user notifications are delivered to the email ID configured in the E-mail-1 box on the User Details page.
Windows Pop-up Message	Select this method to display pop-up messages to the user when the event occurs. This method is only supported on Windows-based clients. The pop-up messages are not displayed if a firewall or any other tool is configured to block such messages. The pop-up message notification method is not available in the recent versions of Windows.

- 3 Click **Apply** to save the user details, and return to the **Manage User Accounts and User Groups** page.

Modifying a user account

To modify a user account:

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, select the user account, and click **Modify Details**.
- 3 Modify the user settings, as necessary. For more information, see, [Configuring user details](#).
- 4 Click **Apply** to save the user details.

Deleting a user account

To delete a user account:

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, select the user account, and click **Delete**.

i | **NOTE:** You can select multiple user accounts to delete at one time.

3 In the confirmation dialog box, click **Delete** to remove the user account from the NetVault Server.

i | **NOTE:** The **admin** and **default** user accounts cannot be deleted.

Setting a user policy

The user policy defines the maximum age for a user password and applies globally to all *local* NetVault users. It also helps you to enable secure mode for domain users.

To set a user policy:

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, click **Set User Policies**.
- 3 On the **Set User Policies** page, provide the following information:
 - Select or clear the **Secure Mode** check box to enable or disable the secure mode. For more information see, [Using Secure Mode](#).
 - Select the **Passwords expire after** check box, and type or select the maximum number of days a password can be used before the user is required to change it.
 - Select the **Display Reminder** check box, and specify how many days in advance users are notified to change their password. The message is displayed every time the user logs on to NetVault.
- 4 Click **Apply** to save the user policy.

Using Secure Mode

Secure Mode enables only the domain users having their associated group(s) already added in NetVault Server to log in or to be added in NetVault Server. The privileges and other access is managed through the groups itself, and domain users cannot be assigned privileges directly while Secure Mode is enabled. This imposes a restriction on the domain users in secure mode as their associated group(s) need to be added into NetVault Server before the user logs in or is added to NetVault. The privileges and other access are updated as per the associated groups present in NetVault Server on every login of the domain user.

To use the Secure Mode feature:

- 1 A domain user must login to NetVault Server using domain credentials.
- 2 NetVault administrator must provide required privileges to this domain user and then this domain user adds the domain groups to the NetVault Server.
- 3 Any user with **Users- Administer user accounts** privilege is able to enable or disable **Secure Mode**.

Enabling or Disabling Secure Mode

To enable or disable secure mode:

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, click **Set User Policies**.

- 3 On the **Set User Policies** page, select or clear the **Secure Mode** check box to enable or disable the secure mode.
- 4 Click **Apply** to save the settings.

Using Presets

This section includes the following topics:

- [Creating a Preset](#)
- [Modifying a Preset](#)
- [Deleting a Preset](#)

Creating a Preset

You can create a user-defined Preset by selecting the necessary privileges for a user role and providing a unique name to the set.

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On **Manager User and Group Accounts** page, click **Privileges and Quotas**.
- 3 In the **Denied** list, select the privileges that you want to include, and click **Add**.
- 4 Under **Privileges Presets**, click **Save As**.
- 5 In the **Save Privileges Preset** dialog box, provide a unique name for the Preset, and click **Save**.

The Preset is added to the Privileges Presets list. You can assign this Preset to any other user.

Modifying a Preset

You can modify a Preset by adding or removing the user privileges. Any changes made to the Presets do not affect the permissions granted to the existing users.

To modify a Preset:

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, click **Privileges and Quotas**.
- 3 In the **Privileges Presets** list, select the Preset that you want to modify, and click **Load**.
The selected privileges are moved to the **Granted** list.
- 4 To grant additional privileges, select the privileges in the **Denied** list, and click **Add**.
The selected privileges are moved to the **Granted** list.
- 5 To revoke any privileges, select the privileges in the **Granted** list, and click **Remove**.
- 6 Under **Privileges Presets**, click **Save As**.
- 7 In the **Save Privileges Preset** dialog box, select the Preset name, and click **Save** to overwrite the set.

Deleting a Preset

You can delete a Preset if it is no longer required.

i | **NOTE:** You can delete a predefined Preset, but when you restart the NetVault Server, the set is recreated. The default settings are restored for the predefined presets when these sets are recreated on server restart.

To delete a Preset:

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, click **Privileges and Quotas**.
- 3 In the **Privileges Presets** list, select the Preset that you want to delete, and click **Delete**.
- 4 In the confirmation dialog box, click **OK**.

User privileges

The following table provides a brief description of the types of user privileges in NetVault.

Table 103. User privileges in NetVault

Privilege	Description
Clients — Add/remove clients	Permission to add and remove NetVault Clients.
Clients — Add/remove virtual/clustered clients	Permission to add and remove virtual clients.
Clients — Administer client groups	Permission to create, modify, and delete client groups.
Clients — Configure client	Permission to configure clients.
Clients — Get client properties	Permission to view client properties.
Clients — Set firewall relationship	Permission to set firewall relationship between the NetVault Server and Client.
Device — Add or Update Snapshot Array	Permission to add or modify Snapshot Array.
Device — Manage Snapshot Array	Permission to manage snapshots using Explore Snapshot feature
Devices — Add libraries	Permission to add tape libraries to the NetVault Server.
Devices — Add Random Access Store	Permission to add disk-based backup devices.
Devices — Add simple drives	Permission to add standalone drives to the NetVault Server.
Devices — Clean drives	Permission to run the Clean command for a drive.
Devices — Manage devices	Permission to perform device management tasks.
Devices — Open and close entry/exit ports	Permission to issue commands to open or close the entry/exit ports.
Devices — Open and close library doors	Permission to issue commands to open or close the library doors.
Devices — Perform device checks	Permission to run check off-line devices.
Devices — Reconfigure devices	Permission to reconfigure added devices.
Devices — Remove devices	Permission to remove devices from the NetVault Server.
Devices — Set drive cleaning properties	Permission to set the drive cleaning options.
Devices — Update Random Access Store	Permission to modify disk-based backup devices.
Jobs — Abort jobs	Permission to abort active jobs.
Jobs — Acknowledge policy errors	Permission to acknowledge policy errors and remove the error flags.
Jobs — Administer backup/restore sets	Permission to create, modify, and delete NetVault Sets.
Jobs — Administer policies	Permission to create and manage policies.
Jobs — Delete job	Permission to delete NetVault jobs.
Jobs — Delete scheduled phase	Permission to delete scheduled jobs.
Jobs — Hold job	Permission to place jobs on hold.

Table 103. User privileges in NetVault

Privilege	Description
Jobs — Jobs owned by this user may run	Permission to submit and run NetVault jobs. NOTE: This user privilege lets you submit or run a job, but it does not let you create or modify jobs. To create backup and restore jobs, you require the following user privileges: <ul style="list-style-type: none"> Jobs — Submit/update backup Jobs Jobs — Submit/update restore Jobs
Jobs — Quiesce policy	Permission to place backup policies in a quiesced state.
Jobs — Restart job	Permission to restart Plug-in <i>for FileSystem</i> backups.
Jobs — Resume job	Permission to resume jobs placed on hold.
Jobs — Run predefined jobs instantly	Permission to issue the Run Now command.
Jobs — Stop Job	Permission to stop active Plug-in <i>for FileSystem</i> backups.
Jobs — Submit/update backup jobs	Permission to create and modify backup jobs. NOTE: This user privilege lets you create and modify backup jobs, but it does not let you run a backup job. To run a job, you require the user privilege Jobs — Jobs owned by this user may run.
Jobs — Submit/update restore jobs	Permission to create and modify restore jobs. NOTE: This user privilege lets you create and modify restore jobs, but it does not let you run a restore job. To run a restore job, you require the user privilege Jobs — Jobs owned by this user may run.
Jobs — View backup jobs	Permission to view backup job definitions.
Jobs — View backup/restore sets	Permission to view set definitions.
Jobs — View policies	Permission to view policy definitions.
Jobs — View restore jobs	Permission to view restore job definitions.
Media — Blank ANSI media	Permission to blank ANSI media.
Media — Blank bad media	Permission to blank a bad media item.
Media — Blank media	Permission to blank media.
Media — Blank non-NetVault media	Permission to blank non-NetVault media.
Media — Export media	Permission to export media to entry/exit port.
Media — Get media or device item status	Permission to view device and media status.
Media — Import media	Permission to import NetApp VTL's shadow tapes.
Media — List media	Permission to view media list.
Media — Load/unload media from drives	Permission to load and unload media from drives.
Media — Manage media requests	Permission to change media request priority and place media request on hold.
Media — Mark media for re-use	Permission to manually mark media for reuse.
Media — Modify backup expiry data	Permission to modify the retirement period for backups.
Media — Remove media	Permission to remove media information from the NetVault Database.
Media — Scan foreign media	Permission to scan foreign media items and import backups residing on them.
Media — View and diagnose media requests	Permission to view and diagnose Media requests.
Media — View backup expiry data	Permission to view the retirement period for backups.
Media — View media properties	Permission to view media properties.
Media — Write media labels	Permission to perform individual and bulk media labeling operations.

Table 103. User privileges in NetVault

Privilege	Description
Reports — Modify/edit report Jobs and components	Permission to create and modify user-defined reports and report components.
Reports — View and run reports	Permission to generate and view reports.
StorageTier — Administer Storage Tier	Permission to create, modify, and delete storage tier.
StorageTier — List Storage Tier	Permission to view storage tier list.
System — Change global notification profile	Permission to set up global notification profile for NetVault.
System — Dump log entries	Permission to generate log dumps in text, binary, or database table dump formats.
System — Install/remove software packages	Permission to install and remove NetVault plug-ins.
System — Install license	Permission to install product license keys.
System — Manage dashboards	Permission to manage NetVault dashboard.
System — Permitted to use CLI tools	Permission to use the CLI utilities.
System — Purge log entries	Permission to delete log messages from the NetVault Database.
System — View dashboards	Permission to view NetVault dashboard.
Users — Administer user accounts	Permission to create, modify, or delete user accounts.
Users — Reset password for user	Permission to reset password for user accounts.
Server- Manage Server Settings	Permission to manage server settings.
Server - View Events	Permission to view events.
Server - View Logs	Permission to view logs.
Service - Administer and restart	Permission to administer and restart the NetVault services.

Predefined Presets

- **Administrator:** This role allows the user to administer NetVault and perform all functions in NetVault.
- **Backup Administrator:** This role allows the user to administer backup and restore jobs and perform device- and media-related functions.
- **Backup Operator:** This role allows the user to run and manage backup jobs. This role also allows the user to administer restore jobs.
- **Helpdesk:** This role allows the user to perform restore jobs.
- **Media Operator:** This role allows the user to perform device- and media-related functions.
- **Monitor:** This role allows the user to monitor clients, devices, jobs, and storage media.

The following table lists the privileges included in the predefined Presets.

Table 104. Predefined Presets

Privilege	Administrator	Backup Administrator	Backup Operator	Helpdesk	Media Operator	Monitor
Clients — Add/remove clients	X					
Clients — Add/remove virtual/clustered clients	X					

Table 104. Predefined Presets

Privilege	Administrator	Backup Administrator	Backup Operator	Helpdesk	Media Operator	Monitor
Clients — Administer client groups	X					
Clients — Configure client	X					
Clients — Get client properties	X	X	X	X		X
Clients — Set firewall relationship	X	X				
Device — Add or Update Snapshot Array	X	X				
Devices — Add libraries	X	X				
Devices — Add Random Access Store	X	X				
Devices — Add simple drives	X	X				
Devices — Clean drives	X	X			X	
Devices — Manage devices	X	X				
Devices — Open and close entry/exit ports	X	X			X	
Devices — Open and close library doors	X	X			X	
Devices — Perform device checks	X	X			X	X
Devices — Reconfigure devices	X	X				
Devices — Remove devices	X	X				
Devices — Set drive cleaning properties	X	X			X	
Devices — Update Random Access Store	X	X				
Jobs — Abort jobs	X	X	X			
Jobs — Acknowledge policy errors	X	X	X			
Jobs — Administer backup/restore sets	X	X	X	X		
Jobs — Administer policies	X	X				
Jobs — Delete job	X	X				
Jobs — Delete scheduled phase	X	X	X			
Jobs — Hold job	X	X	X			
Jobs — Jobs owned by this user may run	X	X	X	X		
Jobs — Quiesce policy	X	X	X			
Jobs — Restart job	X	X	X			
Jobs — Resume job	X	X	X			

Table 104. Predefined Presets

Privilege	Administrator	Backup Administrator	Backup Operator	Helpdesk	Media Operator	Monitor
Jobs — Run predefined jobs instantly	X	X	X	X		
Jobs — Stop Job	X	X	X			
Jobs — Submit/update backup jobs	X	X	X			
Jobs — Submit/update restore jobs	X	X	X	X		
Jobs — View backup jobs	X	X	X	X		X
Jobs — View backup/restore sets	X	X	X	X		X
Jobs — View policies	X	X	X	X		X
Jobs — View restore jobs	X	X	X	X		X
Media — Blank ANSI media	X	X			X	
Media — Blank bad media	X	X			X	
Media — Blank media	X	X				
Media — Blank non-NetVault media	X	X				
Media — Export media	X	X			X	
Media — Get media or device item status	X	X			X	
Media — Import media	X	X			X	
Media — List media	X	X	X	X	X	X
Media — Load/unload media from drives	X	X	X		X	
Media — Manage media requests	X	X			X	
Media — Mark media for re-use	X	X			X	
Media — Modify backup expiry data	X	X			X	
Media — Remove media	X	X			X	
Media — Scan foreign media	X	X			X	
Media — View and diagnose media requests	X	X	X		X	X
Media — View backup expiry data	X	X			X	
Media — View media properties	X	X	X		X	X
Media — Write media labels	X	X			X	
Reports — Modify/edit report Jobs and components	X	X				X
Reports — View and run reports	X	X				X

Table 104. Predefined Presets

Privilege	Administrator	Backup Administrator	Backup Operator	Helpdesk	Media Operator	Monitor
Storage Tier — Administer storage tier	X	X			X	
Storage Tier — List storage tier	X	X	X		X	X
System — Change global notification profile	X	X				
System — Dump log entries	X	X	X			
System — Install/remove software packages	X					
System — Install license	X					
System — Manage dashboards	X					
System — Permitted to use CLI tools	X	X				
System — Purge log entries	X	X				
System — View dashboards	X	X	X	X	X	X
Users — Administer user accounts	X					
Users — Reset password for user	X					

Integrating Active Directory with NetVault

NetVault installed on a Windows or Linux machine can communicate with Microsoft Active Directory (AD). Additionally, NetVault installed on a Linux machine can communicate with OpenLDAP Directory Services and Samba Active Directory. Integrating AD with NetVault enables role-based access control in NetVault. It lets users log in to NetVault using their AD credentials. It also lets AD users manage NetVault users.

i | NOTE: NetVault Server supports Secure Light Directory Access Protocol (LDAPS) protocol for both Microsoft Active Directory (AD) and OpenLDAP Directory.

i | IMPORTANT: During Active Directory integration with NetVault, the Active Directory service must always be available. Additionally, the NetVault Server should be a member of the domain.

Adding an AD group to NetVault and specifying privileges, levels of access, and notifications for the group ensures that members of the group have the corresponding access from the first time that they log on.

- [Considerations for Linux-based NetVault](#)
- [Managing AD users](#)
- [Managing AD groups](#)

Considerations for Linux-based NetVault

Before you integrate Active Directory (AD) with a Linux-based installation of NetVault, consider the following conditions and prerequisites:

- To let the Linux server communicate with AD, integrate NetVault with the domain controller using Samba Winbind or System Security Services Daemon (SSSD).
- Host name of the Linux machine should be set to its FQDN. (For example, *mymachine.mydomain.com*)
- Enable the use of fully qualified domain names on the Linux machine on which NetVault is installed. Short names are not permitted.
 - **Samba Winbind:** Under the `[global]` section, `smb.conf` must have the following entry:

```
winbind use default domain = no
```
 - **SSSD:** under the `[domain/<domain name>]` section, `sssd.conf` must have the following entry:

```
use_fully_qualified_names = True
```
- On the Linux machine where NetVault is installed, to retrieve all domain controller groups, enable enumeration of group databases.
 - **Samba Winbind:** Under the `[global]` section, `smb.conf` must have the following entry:

```
winbind enum groups = yes
```
 - **SSSD:** Under the `[domain/<domain name>]`, to retrieve domain controller groups and to fetch groups from child and trusted domains, respectively, the `sssd.conf` must have the following entries:

```
enumerate = True
subdomain_enumerate = all
```
- While logging in or adding a domain user, using the DNS suffix with the domain name is recommended; for example, `domain.local\user` or `username@dns.local`.

Managing AD users

The following list identifies what various users can and cannot do within NetVault as it pertains to Active Directory (AD):

- When an AD user logs in using an AD login, such as a domain name followed by the user name or a name that contains the `@` symbol, NetVault authenticates the login against Active Directory. This behavior lets an AD user log in to NetVault even without a local NetVault account.
 - If authentication succeeds and this is the first time that the user is logging in, NetVault takes one of the following actions:
 - If a corresponding AD group exists in NetVault, NetVault adds the user and assigns the user the privileges, access levels, and notification settings that are defined for the corresponding AD group.
 - If the user does not have a NetVault account, NetVault adds the user but excludes all privileges, access levels, and notification settings. When the local NetVault Administrator logs in, the administrator can view the AD user entry and assign the applicable privileges.
 - NetVault automatically fetches the respective user details from the domain controller and stores it locally in the NetVault database. In case of secure mode, the details are fetched from domain controller and overwritten in the NetVault database on every successful login.
 - If authentication fails, NetVault displays an error message and prevents the user from logging in.
- All NetVault administrators can view all local and AD NetVault users.
- An AD NetVault user with administrative privileges can add authenticated AD users; a local NetVault Administrator cannot.

- An AD NetVault user with administrative privileges can rename an existing NetVault user to an authenticated AD NetVault user; a local NetVault Administrator cannot. You must ensure that the password fields remains empty to maintain the integrity of the user's AD password.
- Once the AD user is added in NetVault, it cannot be renamed to any other user name.
- No one can change an AD user's password or set the password to use the **Password never expires** option.
- A local NetVault Administrator or an AD NetVault user with administrative privileges can change user-based information that is accessed by using the **Modify Details** option as this information is stored in the NetVault Database. However, if the secure mode is enabled, the NetVault administrator and other users are allowed to edit or change only **E-mail 2** and **E-mail 3** of an AD NetVault user.

i **NOTE:** On Windows based NetVault Server, if the Workstation attribute is set for a user in the AD database, then NetVault fetches that information and stores it against the Workstation in the user details. If the user has multiple workstations, they display as separated by commas. If you want the AD user with the Workstation attribute set in AD to log in to the NetVault Server, then the user must specify the NetVault Server in the AD Workstations list.

On Linux-based NetVault Server, the workstation field cannot be fetched, hence the user is unable to login. We recommended that workstation should not be set in AD for Linux Server.

An administrator can perform the following tasks in NetVault:

- [Adding an AD user](#)
- [Deleting an AD user](#)

Adding an AD user

The preferred method of adding an AD user is for the applicable user to log in and let NetVault complete the authentication process. However, if you have the applicable privileges, at a minimum, you must have the **Users — Administer user accounts** privilege, you can manually add the user, where it will also consider the privileges defined to selected user before renaming to AD user.

When you add an AD user to the NetVault Server, NetVault automatically fetches the respective user information from AD and populates it in the NetVault user details. However, on Linux-based NetVault Server, user details are fetched only after the first login and not if the AD user is added manually. On windows based NetVault Server, when you add users manually, the privileges defined for the AD group to which the user belongs are automatically assigned. If the user has an existing local NetVault account, the privileges defined for that account are taken into consideration by NetVault before redefining the local user as an AD user. For more information on adding, updating, and deleting user accounts, see [Creating a user account](#), [Modifying a user account](#), and [Deleting a user account](#).

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, click **Add User**.
- 3 Select the user on **Manage User Accounts and User Groups** page and click **Modify Details**.
- 4 Enter the AD user name in the `<domain>\<name>` (For example, `domain\username`) or `<name>@<domain>` (For example, `username@domain.com`) format, and click **Apply**.

Deleting an AD user

A local NetVault Administrator or an AD NetVault user with administrative privileges can remove an AD user from NetVault. This does *not* affect the user's account in Active Directory; it only removes the user from the NetVault Database.

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, select the user account that you want to remove, and click **Delete**.

i | **NOTE:** You can select multiple user accounts to delete at one time.

- 3 In the confirmation dialog box, click **Delete**.

Managing AD groups

The following list identifies what various users can and cannot do within NetVault as it pertains to AD user groups:

- A NetVault Administrator can search, list, update, and delete existing NetVault AD groups. However, this administrator cannot list AD groups that have not been added from Active Directory.
- An AD NetVault user with administrative privileges can search, list, update, and delete existing NetVault AD groups. This administrator can also list AD groups that have not been added from Active Directory yet and can add AD groups to NetVault.
- NetVault users who do not have administrative privileges cannot perform any tasks related to AD groups in NetVault.
- While renaming NetVault user to AD user all the privileges of NetVault user will be assigned to new AD user along with the privileges of NetVault AD groups.
- As stated previously, a new AD user can log in to NetVault, and NetVault automatically authenticates the login. If authentication succeeds, NetVault assigns the user the privileges that are defined for the corresponding AD group.

i | **IMPORTANT:** If you manually add an AD user instead of using the automated authentication process, the corresponding privileges that are defined for the user's AD group are automatically assigned.

- If the AD user has an existing local NetVault account, all privileges from that account are assigned to the user when the account is redefined as an AD user account. The privileges defined for the corresponding AD group are also assigned to the redefined account.

An administrator can perform the following tasks in NetVault:

- [Adding an AD group to NetVault Server](#)
- [Viewing a list of AD groups that exist in NetVault](#)
- [Modifying description of added AD group in NetVault](#)
- [Deleting an AD group that was added to NetVault](#)

Adding an AD group to NetVault Server

An AD NetVault user with administrative privileges must add AD groups to the NetVault Server before they can be viewed, updated, or deleted by a local NetVault Administrator.

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, select **Groups** tab, a list of existing groups added in NetVault Server are displayed.
- 3 To add an AD group to the NetVault Database, the AD user must have the **Users — Administer user accounts** privilege, click **Add Group**.

Select **All Groups** to see a list of all the available AD groups that you can add to NetVault Server.

To search the group name, type in **Enter Group Name** search field partially or completely, the search results displays only matching groups. The **Search Group by Name** option is a default selection to add a group in NetVault.

i | **NOTE:** On Linux based NetVault Server, complete group name along with the domain FQDN must be provided. (For example, group_name@FQDN of domain i.e. grp1@testdomain.com)

- 4 Select the applicable group, enter an optional description in the text box, and click **Save Group**.
Once the group is added to the NetVault Database, you can define the privileges, levels of access, and notifications for the group within NetVault.
- 5 To update membership information, click **Client and media group memberships**.
On the **Editing User Group Memberships** page, update the applicable information.
For more information, see [Configuring memberships for a user account or user group](#).
- 6 To update privilege- and quota-based information, click **Privileges and Quotas**.
On the **Editing Privilege and Quota Settings for User Group** page, update the applicable information.
For more information, see [Granting privileges and quota to a user account](#).
- 7 To update notification information, click **Notification Profile**.
On the **Editing Notification Profile for User Group** page, update the applicable information.
For more information, see [Setting up user notification profile](#).

Viewing a list of AD groups that exist in NetVault

Any NetVault administrator can list the AD groups that exist in the NetVault Database.

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page select **Groups** tab, a list of all the groups added in NetVault Server are displayed.

Modifying description of added AD group in NetVault

Any NetVault administrator can modify or edit the AD group description that exist in the NetVault Database.

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page select **Groups** tab, select the applicable group to modify or edit description, and click **Modify Details**.
- 3 Modify the AD group description, as necessary, and click **Apply** to save the description.

Deleting an AD group that was added to NetVault

Any NetVault administrator can delete the AD groups that exist in the NetVault Database.

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, select the applicable group to remove in **Groups** tab, and click **Delete**.

i | **NOTE:** You can select multiple groups to delete at one time.

- 3 In the confirmation dialog box, click **Delete**.

Using Azure Active Directory as an identity provider

The following sections detail how to use and manage a Microsoft Azure Active Directory (Azure AD) account as a NetVault user account.

i | **NOTE:** When using Azure AD as an identity provider, Quest recommends that you use Google Chrome or Mozilla Firefox as your browser.

- [Setting up an Azure AD login with NetVault](#)
- [Managing an Azure AD user](#)
- [Managing Azure AD groups](#)

Setting up an Azure AD login with NetVault

Beginning with release 12.3, NetVault includes the option for users to log in using their Microsoft Azure Active Directory (Azure AD) credentials. Before you can use this option, complete the following prerequisites:

- 1 [Registering NetVault on the Azure portal](#)
- 2 [Setting authentication for Azure AD users](#)
- 3 [Restricting an application to a set of users](#)
- 4 [Configuring Azure AD in NetVault](#)

Registering NetVault on the Azure portal

Before a user can use Microsoft Azure Active Directory (Azure AD) credentials to log in to NetVault, NetVault must be registered on the Azure portal.

To register NetVault on the Azure portal:

- 1 Sign in to the Azure portal.
- 2 In the left-navigation pane of the Azure portal, click **Azure Active Directory**.
- 3 Under **Manage**, click **App registrations**.
- 4 On the **App registrations** page, click **New registration**.
- 5 On the Create page, enter the following details about the NetVault registration:

Table 105. Details for registering an application in Azure

Detail	Description
Name	Enter a name for the application; for example, NetVault.
Application type	Select Web .
Redirect URI	For Web App/API applications, provide the base URL for the app; for example, http://localhost:8443 if the web app is running on your local machine. Users would use this URL to sign in to a web client application. NOTE: After signing out from the NetVault Server WebUI and signing in again from the same page results in the error "The reply url specified in the request does not match the reply urls configured for the application." To avoid this issue, sign in from the base URL you provided during registration.

- 6 Click **Register**.
Azure AD assigns a unique Application ID to your application for that tenant.

Setting authentication for Azure AD users

Before a user can use their Azure AD credentials to log in to NetVault, you must set up authentication in the Azure AD user interface using one of the following procedures.

To set authentication for Azure AD users with App registrations:

- 1 In the Azure AD user interface, navigate to **App registrations**.
- 2 Select **NetVault** as the application that you registered.
- 3 Navigate to **Authentication**.
- 4 Under Implicit Grant, select **Access tokens and ID tokens**.
- 5 Click **Save**.

To set authentication for Azure AD users with App registrations:

- 1 In the Azure AD user interface, navigate to **App registrations**.
- 2 Select **NetVault** as the application that you registered.
- 3 Click **Manifest**.
- 4 Set `"oauth2AllowImplicitFlow"` and `"oauth2AllowImplicitFlow"` to **true**.
- 5 Click **Save**.

Restricting an application to a set of users

Azure AD lets you restrict access to applications to specific sets of users or security groups. To view the procedure for enabling the app to enable user assignments, search the Microsoft Azure Active Directory documentation. If the user assignment requirement is set to Yes, you must assign the Azure AD users to the application in the Azure portal before they can log in to NetVault.

To assign a user to NetVault in the Azure portal:

- 1 Navigate to Enterprise applications.
- 2 Select NetVault as the application that you registered.
- 3 Under **Users and groups**, click **Add user**.
- 4 Select the appropriate user.
- 5 Click **Assign**.

Configuring Azure AD in NetVault

To configure Azure AD in NetVault:

- 1 Log in to the NetVault Server as an administrator.
- 2 From the Navigation pane in the NetVault WebUI, click **Change Settings**.
- 3 On the **NetVault Server Settings** page, under **User management**, click **Identity Provider**.
- 4 In the Identity Provider window, enter the **Application ID** and **Directory ID**.

; | **NOTE:** To retrieve the Directory ID for the application, see the Azure AD properties.

- 5 Click **Apply**.

Managing an Azure AD user

The following list identifies what various users can and cannot do within NetVault as it pertains to Azure AD:

- When an Azure AD user logs in using Sign In with Azure AD, NetVault authenticates the login against Azure Active Directory. This behavior lets an Azure AD user log in to NetVault without needing a local NetVault account.
- If authentication succeeds and this is the first time that the user is logging in, NetVault takes one of the following actions:
 - If a corresponding Azure AD group exists in NetVault, NetVault adds the user and assigns the user the privileges, access levels, and notification settings that are defined for the corresponding Azure AD group.
 - If the user does not have a NetVault account, NetVault adds the user but excludes all privileges, access levels, and notification settings. When the local NetVault administrator logs in, the administrator can view the Azure AD user entry and assign the applicable privileges.
 - NetVault automatically fetches the respective user details from the domain controller and stores them locally in the NetVault Database. When in secure mode, NetVault fetches the details from the domain controller and overwrites them in the NetVault Database on every successful login.
- If authentication fails, NetVault displays an error message and prevents the user from logging in.
- All NetVault administrators can view all local, AD, and Azure AD NetVault users.
- An Azure AD NetVault user with administrative privileges can add authenticated Azure AD users; a local NetVault administrator cannot.
- An Azure AD NetVault user with administrative privileges can rename an existing local NetVault user to an authenticated Azure AD NetVault user; a local NetVault administrator cannot.

i | **NOTE:** To maintain the integrity of the user's Azure AD password, ensure that the password field remains empty.

- After you add the Azure AD user in NetVault, you cannot rename the user to any other user name.
- No one can change the password of an Azure AD user or set the password to use the **Password never expires** option.
- A local, AD, or Azure AD NetVault user with administrative privileges can change user-based information that is accessed by using the Modify Details option as this information is stored in the NetVault Database. However, if the secure mode is enabled, the NetVault administrator and other users are allowed to edit or change only E-mail 2 and E-mail 3 of an Azure AD NetVault user.

The following procedures detail ways of managing a Microsoft Azure Active Directory (Azure AD) user account.

- [Adding an Azure AD user](#)
- [Deleting an AD user](#)

Adding an Azure AD user

As with Active Directory (AD), the preferred method of adding an Azure AD user is for the applicable user to log in and let NetVault complete the authentication process. However, if you have the applicable privileges, at a minimum, you must have the **Users — Administer user accounts** privilege, you can manually add the user, where it will also consider the privileges defined to the selected user before renaming it to an Azure AD user.

When you add an Azure AD user to the NetVault Server, NetVault automatically fetches the respective user information from Azure AD and populates it in the NetVault user details. When you add users manually, NetVault automatically assigns the user the privileges that are defined for the Azure AD group to which the user belongs.

To add an Azure AD user account to NetVault:

- 1 In the Navigation pane, click **Users and Groups**.

- 2 On the **Manage User Accounts and User Groups** page, click **Add User**.
- 3 Select the user on **Manage User Accounts and User Groups** page and click **Modify Details**.
- 4 Enter the Azure AD user name; for example, username@domain.com.
- 5 Click **Apply**.

Deleting an Azure AD user

A local NetVault Administrator or NetVault user with administrative privileges can remove an Azure AD user from NetVault. This does *not* affect the user's account in Azure AD; it only removes the user from the NetVault Database.

To delete an Azure AD user from the NetVault Database:

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, select the user account that you want to remove, and click **Delete**.

 **NOTE:** You can select multiple user accounts to delete at one time.
- 3 In the confirmation dialog box, click **Delete**.

Managing Azure AD groups

An administrator can perform the following tasks for Microsoft Azure Active Directory (Azure AD) groups.

- [Adding an Azure AD group to NetVault Server](#)
- [Viewing a list of Azure AD groups that exist in NetVault](#)
- [Modifying the description of an Azure AD group in NetVault](#)
- [Deleting an Azure AD group from NetVault](#)

Adding an Azure AD group to NetVault Server

An Azure AD NetVault user with administrative privileges must add Azure AD groups to the NetVault Server before they can be viewed, updated, or deleted by a local NetVault Administrator.

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, select **Groups** tab, a list of existing groups added in NetVault Server are displayed.
- 3 To add an Azure AD group to the NetVault Database, the Azure AD user must have the **Users — Administer user accounts** privilege, click **Add Group**.

Select **All Groups** to see a list of all the available Azure AD groups that you can add to NetVault Server.

To search the group name, type in **Enter Group Name** search field partially or completely, the search results displays only matching groups. The **Search Group by Name** option is a default selection to add a group in NetVault.

- 4 Select the applicable group, enter an optional description in the text box, and click **Save Group**.

Once the group is added to the NetVault Database, you can define the privileges, levels of access, and notifications for the group within NetVault.
- 5 To update membership information, click **Client and media group memberships**.

On the **Editing User Group Memberships** page, update the applicable information.

For more information, see [Configuring memberships for a user account or user group](#).

- 6 To update privilege- and quota-based information, click **Privileges and Quotas**.
On the **Editing Privilege and Quota Settings for User Group** page, update the applicable information.
For more information, see [Granting privileges and quota to a user account](#).
- 7 To update notification information, click **Notification Profile**.
On the **Editing Notification Profile for User Group** page, update the applicable information.
For more information, see [Setting up user notification profile](#).

Viewing a list of Azure AD groups that exist in NetVault

Any NetVault administrator can list the Azure AD groups that exist in the NetVault Database.

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page select **Groups** tab, a list of all the groups added in NetVault Server are displayed.

Modifying the description of an Azure AD group in NetVault

Any NetVault administrator can modify or edit the Azure AD group description that exist in the NetVault Database.

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page select **Groups** tab, select the applicable group to modify or edit description, and click **Modify Details**.
- 3 Modify the Azure AD group description, as necessary, and click **Apply** to save the description.

Deleting an Azure AD group from NetVault

Any NetVault administrator can delete the Azure AD groups that exist in the NetVault Database.

- 1 In the Navigation pane, click **Users and Groups**.
- 2 On the **Manage User Accounts and User Groups** page, select the applicable group to remove in **Groups** tab, and click **Delete**.

i | **NOTE:** You can select multiple groups to delete at one time.

- 3 In the confirmation dialog box, click **Delete**.

Monitoring events and configuring notifications

- [About NetVault events](#)
- [About notification methods](#)
- [Event classes](#)
- [Events types](#)
- [Using global notification methods](#)
- [Using custom notification methods](#)
- [Removing a user-defined job event](#)
- [Removing a user-defined report event](#)

About NetVault events

An event can be described as any significant occurrence in the NetVault system. It can indicate a problem that requires your response or a condition that you want to be notified about. Events can be predefined or user-defined.

Events are recorded in event logs. You can view event log messages from the **View Events** page. Event logs can be used to track activities or respond to problems or errors reported by the system.

Predefined events

NetVault includes a set of predefined events, which are associated with various job-related and non-job-related occurrences in the system. The predefined events are organized into different categories or classes. Each category or class includes one or more event types. For information about predefined event classes and event types, see [Event classes](#) and [Events types](#).

User-defined events

You can create user-defined events in NetVault to monitor specific events associated with backup jobs, restore jobs, policies, report jobs, and log messages. For example, you can create a user-defined event for a backup job and receive notification if the job completes successfully, completes with warnings, or fails. For more information, see the relevant sections in this guide.

About notification methods

Different methods available in NetVault for sending notifications when an event occurs. You can use global notification methods, user notification profiles, or custom notification methods.

Global notification methods

The global notification methods include SysOp Email, Simple Network Management Protocol (SNMP) Traps, Print Report, and Run a Job. For more information about these methods, see [Using global notification methods](#).

User notification profile

You can set up a user notification profile to send email notifications to a NetVault user when an event occurs. For more information, see [Setting up user notification profile](#).

Custom notification method

Custom Notification Methods are created using external scripts, and typically used to perform some task when an event occurs in the NetVault Domain. For more information, see [Using custom notification methods](#).

Event classes

The predefined events in NetVault are organized into different categories or classes:

- Audit
- Catalog
- DR Appliance
- Device
- Job
- Licensing
- Log Daemon
- Machines
- Media
- Media Database
- NetVault Time
- Policy
- Scheduler Database
- Stats Collection

Events types

The following table provides a brief description of the predefined event types available in various event classes.

Table 106. Events types

Event class	Event type	Description
Audit	Failed to Update Audit File	Occurs when NetVault is unable to update the audit logs.
	Update Session Map Failed	Occurs when NetVault is unable to update the session owner mapping file.
Catalog	Catalog Upgrade Available	Occurs when a Catalog Search package upgrade is available.
DR Appliance	Container Created	Occurs when a container is created on a Quest DR Series system.
	Container Deleted	Occurs when a container is deleted from a Quest DR Series system.
Device	Check	Occurs when a request is received to check a device.
	Check Serial Number of Drive	Occurs when a request is received to check the serial number of a drive.
	Check Serial Number of All Drives	Occurs when a request is received to check the serial numbers of all drives.
	Close Door	Occurs when a request is received to close a library door
	Close EEPort	Occurs when a request is received to close an entry/exit port.
	Close EEPort Clean	Occurs when a request is received to close an entry/exit port containing cleaning media.

Table 106. Events types (Continued)

Event class	Event type	Description
	Device Forced Offline	Occurs when the offline command is selected for a device.
	DR Devices Marked Orphan	Occurs when a DR device is marked as orphan.
	Device Gone Offline	Occurs when a device becomes offline.
	Drive Unavailable	Occurs when a drive becomes offline.
	Library Gone Offline	Occurs when a library becomes offline.
	Library Scan Completed	Occurs when a media scan request is completed.
	Map	Occurs when an ACSLS drive is mapped.
	No Suitable Drive	Occurs when NetVault is unable to find a suitable drive to run a backup or restore job.
	Open Door	Occurs when a request is received to open a library door.
	Open EEPort	Occurs when a request is received to open an entry/exit port.
	Reconfigure Device	Occurs when a device is modified.
	Remove	Occurs when a device is removed.
	Restart Library	Occurs when a library is restarted.
	Synchronize Silo Media	Occurs when silo media items are synchronized.
	Unmap	Occurs when an ACSLS drive is un-mapped.
	Update Serial Number of Drive	Occurs when a request is received to update the serial number of a drive.
	Update Serial Number of all Drives	Occurs when a request is received to update the serial number of all drives.
Job	All Job Retries Failed	Occurs when all retry attempts for a job have failed.
	Job Abort Requested	Occurs when a request is received to abort a job.
	Job Aborted	Occurs when a job is aborted.
	Job Completed Successfully	Occurs when a job completes successfully.
	Job Completed with Warnings	Occurs when a job completes with warnings.
	Job Created	Occurs when a job is created.
	Job Deleted	Occurs when a job is deleted.
	Job Died	Occurs when a job terminates unexpectedly.
	Job Failed	Occurs when a job fails.
	Job Finished	Occurs when a job is completed.
	Job Hold	Occurs when a job is put on hold.
	Job Modified	Occurs when a job is modified.
	Job Phase Starting	Occurs when phase 1 or phase 2 starts for a job.
	Job Resume	Occurs when a job that was placed on hold is resumed.
	Job Retry Scheduled	Occurs when a job is rescheduled after a failed attempt.
	Job Run Now	Occurs when a job is submitted to run immediately.
	Job Scheduled	Occurs when a job is submitted.
	Job Stop Requested	Occurs when a request is received to stop a job.
	Job Stopped	Occurs when a job is stopped.
	Scheduled Phase Deleted	Occurs when a scheduled phase is deleted for a job.
	Set Created	Occurs when a set is created.
	Set Deleted	Occurs when a set is deleted.
	Set Modified	Occurs when a set is modified.

Table 106. Events types (Continued)

Event class	Event type	Description
Licensing	License Exceeded	Occurs when usage exceeds the available licenses.
	License Expiring	Occurs when the evaluation license is nearing expiry or has expired. This event occurs when the validity period for the evaluation license is less than or equal to seven days.
Log Daemon	Home Drive Becoming Full	Occurs when disk space usage reaches warning or critical threshold.
Machines	Client Added	Occurs when a NetVault Client is added to the domain.
	Client Down	Occurs when a NetVault Client becomes offline.
	Client Group Created	Occurs when a client group is created.
	Client Group Deleted	Occurs when a client group is deleted.
	Client Group Modified	Occurs when a client group is modified.
	Client Group Renamed	Occurs when a client group is renamed.
	Client Removed	Occurs when a client is removed.
	Virtual Client Added	Occurs when a cluster-aware plug-in is installed on the server. For more information about virtual clients and cluster-aware plug-ins, see Working with client clusters .
	Virtual Client Removed	Occurs when a cluster-aware plug-in is removed. For more information about virtual clients and cluster-aware plug-ins, see Working with client clusters .
Media	Blank	Occurs when a request is received to blank a media item.
	Delete Group	Occurs when a media group is deleted.
	Delete Job Instance	Occurs when a job instance is deleted.
	Export	Occurs when a piece of media is exported to an entry/exit port.
	Import	Occurs when a request is received to import a piece of media.
	Import Clean	Occurs when a request is received to import a cleaning tape.
	Label	Occurs when a label is assigned to a piece of media.
	Load	Occurs when a piece of media is loaded into a drive.
	Media Blanked	Occurs when a piece of media is blanked.
	Media Deleted	Occurs when a piece of media is removed from the NetVault Database.
	Media Full	Occurs when a piece of media becomes full.
	Media Labeled	Occurs when a media label request is completed.
	Media Loaded	Occurs when a media load request is completed.
	Media Marked Bad	Occurs when a piece of media is marked bad.
	Media Request Change Priority	Occurs when a media request priority is changed.
	Media Request Timeout Expired	Occurs when a media request has timed out for a backup job. This event occurs when NetVault is unable to find suitable media within the specified timeout interval.
	Media Suspect	Occurs when a piece of media is marked suspect.
	Media Unexpectedly BLANK	Occurs when a piece of media is found to be unexpectedly BLANK when it is selected for use.
	Media Unloaded	Occurs when a piece of media is unloaded.
	Media Unusable	Occurs when a drive rejects a piece of media.
No Suitable Media	Occurs when NetVault is unable to find a suitable media to complete the backup job.	

Table 106. Events types (Continued)

Event class	Event type	Description
Media Database	Request Off Hold	Occurs when a media request is taken off hold.
	Request On Hold	Occurs when a media request is placed on hold.
	Reuse	Occurs when a piece of media is marked for re-use.
	Scan Request	Occurs when a request is received to scan a foreign media.
	Unload	Occurs when a piece of media is unloaded.
	Update Properties	Occurs when media properties are modified.
	Backup Added	Occurs when a backup record is added to the Media Database.
	Backup Deleted	Occurs when a backup record is deleted from the Media Database.
	Backup Modified	Occurs when a new backup record is modified.
	Backup Retired	Occurs when a saveset is retired and its information is deleted from the NetVault Database.
	Index Compressed	Occurs when a backup index is compressed.
	Index Compression Failure	Occurs when NetVault fails to decompress a backup index due to insufficient disk space.
	Index Loaded	Occurs when an offline index is temporarily loaded.
	Index Offlined	Occurs when the backup index for a saveset is deleted from the NetVault Database.
	Index Offlining Failure	Occurs when NetVault fails to delete the backup index for a saveset from the database.
	Index Read Failure	Occurs when NetVault fails to read the index for a saveset.
	Index Uncompressed	Occurs when a backup index is de-compressed.
	Media Database Check Failed	Occurs when the database integrity or reference check fails for the Media Database.
	Media Database Check Passed	Occurs when the Media Database check completes successfully.
Media Database Check Passed with Warnings	Occurs when the Media Database check generates warnings.	
NetVault Time	Modify Backup Expiry	Occurs when the retirement period is modified for a saveset.
	No Time Source	Occurs when a NetVault Time Server is not configured for the Domain.
	Server Time Inconsistency	Occurs when the NetVault Time on two or more controlling servers do not match. (This event can occur when a client is added to two or more NetVault Servers.)
	Server Time Unknown	Occurs when the Time Server is unable to fetch NetVault Time from the Source.
	Time Server Changed	Occurs when the Time Server is changed for a NetVault Domain.
Policy	Time Server Not Responding	Occurs when the Time Server is not responding to a time request.
	Time Server Removed	Occurs when the Time Server is removed from the NetVault Domain.
	Policy Branch Errors Acknowledged	Occurs when policy branch errors or warnings are acknowledged.
	Policy Created	Occurs when a policy is created.
	Policy Deleted	Occurs when a policy is deleted.
	Policy Errors Acknowledged	Occurs when policy errors or warnings are acknowledged.

Table 106. Events types (Continued)

Event class	Event type	Description
	Policy Modified	Occurs when a policy is modified.
	Policy Quiesce	Occurs when a request is received to quiesce a policy.
	Policy Quiesced	Occurs when a policy is placed in a quiesced state.
Scheduler Database	Scheduler Database Check Failed	Occurs when the database integrity or reference check fails for the scheduler database.
	Scheduler Database Check Passed	Occurs when the Scheduler Database check is completed successfully.
	Scheduler Database Check Passed with Warnings	Occurs when the Scheduler Database check generates warnings.
Stats Collection	Cache Too Small	Occurs when cache memory is running low for the Statistics Manager; this would result in the process running slowly.
	Cannot Accept Records	Occurs when the Statistics Manager refuses to accept data from other processes.
	Lost Server	Occurs when the Statistics Manager discards the I/O data collected for a server on polling timeout.
	Stats Manager Ready	Occurs when NetVault starts the Statistics Manager.

Viewing event logs

You can view the event log messages from the **View Events** page.

To view event logs:

- 1 In the Navigation pane, click **View Events**.
- 2 On the **View Events** page, you can view the following information:
 - **Date:** Date and time when the event was raised
 - **Class:** Event class
 - **Event:** Event type
 - **Message:** Detailed log message or description
- 3 By default, the table is sorted by Date (newest to oldest).

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To set the filter options (For more information on table filter options for **View Events** page, see [Table 9](#)), view the page size setting, column sort order, applied filters in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table.

For more information, see [Customizing tables in NetVault WebUI](#).
- 4 You can use the Search option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string. Type the search text into the box and click the tick icon or press Enter.
- 5 The **View Events** page loads a maximum of 10,000 records initially. The total number of records fetched is displayed at the lower-right corner of the table. You can click **Load more** to load the next set of records, if available. Each load action fetches a maximum of 10,000 records. This button is disabled if there are no more records to load.
- 6 To stop or resume live updates, click the **Pause or Resume Live Updates** button.

- 7 To perform a logs-related task, click the corresponding button in the Operations pane. Alternatively, click a link in the Navigation pane to open a different page.

Using global notification methods

This section includes the following topics:

- [About global notification methods](#)
- [Setting up a global notification method](#)

About global notification methods

Global notification methods provide different mechanisms for sending notifications when an event occurs.

The available methods include the following:

- **Sysop Email:** Use this method to send an email notification to the Administrator when an event occurs.
- **Print Report:** Use this method to print a report after it is successfully generated from the **View Reports** page. This method can only be used for Report Job events.
- **Run a Job:** Use this method to run a job when an event occurs.
- **SNMP Trap:** Use this method to send SNMP traps (notification) to a network manager host when an event occurs.

The Simple Network Management Protocol (SNMP) provides a method to monitor and control the network devices on TCP/IP-based networks. An SNMP trap is a notification (message) sent from a managed device to the network management host when a significant event occurs. The event does not have to be an outage, a fault, or a security violation. The SNMP Trap method allows the administrator to monitor the NetVault events from an SNMP-based network management interface as a part of the network management operations.

Setting up a global notification method

To set up a global notification method:

- 1 In the Navigation pane, click **Configure Notifications**.
- 2 On the **Editing Global Notification Profile** page, open the event class, and then open the event type that you want to be notified about.
- 3 Configure the notification method that you want to use. You can configure multiple notification methods for a single event.

Figure 28. Editing Global Notification Profile page

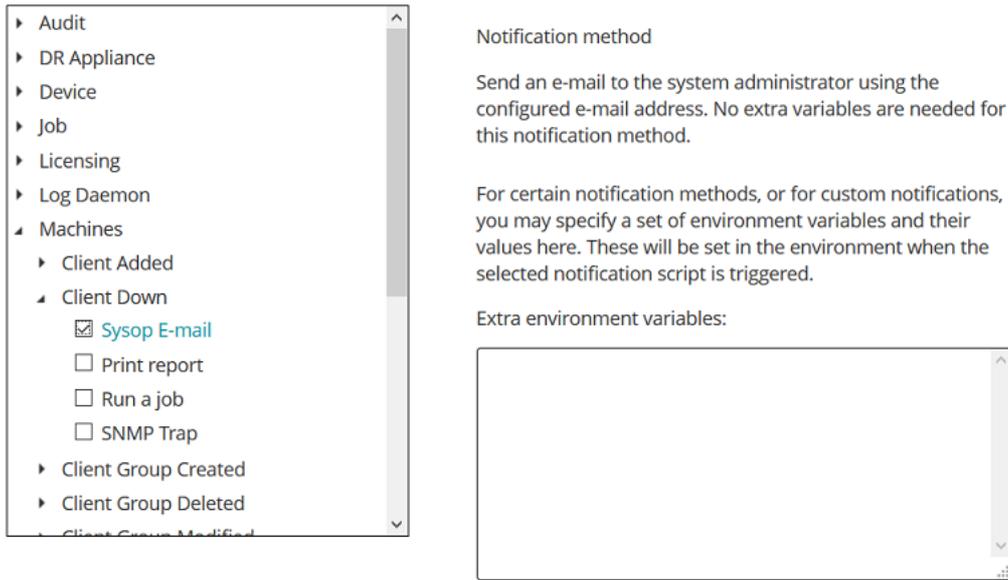


Table 107. Setting up global notification method

Notification method	Description
Sysop Email	<p>Select this method to send an email notification to the Administrator (Sysop) when the event occurs.</p> <p>Before you can use this method for sending notifications, you must complete the following tasks:</p> <ul style="list-style-type: none"> • Configure the Outgoing Email Server (SMTP) settings. For more information, see Configuring email server settings for notifications. • Configure SysOp email address. For more information, see the Configuring SysOp email ID for notifications.
Print Report	<p>Select this method to print a report after it is successfully generated from the View Reports page. This method can only be used for Report Job events.</p> <p>On Linux-based NetVault Servers, no additional configuration is required to use this method. The printer set up as the default printer for the system is used to print the report.</p> <p>On Windows-based machines, you must do one of the following:</p> <ul style="list-style-type: none"> • Configure a default printer in NetVault. For more information, see Configuring a default printer for notifications. The Print Report method uses this printer for all events. • Alternatively, set the NVPRINTER variable in the Extra environment variables box: <pre>NVPRINTER=<Printer URL></pre> <p>The Print Report method uses this printer for an individual event. You can use the NVPRINTER variable to override the default printer setting for an individual event.</p>

Table 107. Setting up global notification method (Continued)

Notification method	Description
Run a Job	<p>Select this method to run a NetVault job when the event occurs.</p> <p>To use this method, you must do the following:</p> <ul style="list-style-type: none">• Create the job that you want to run when the event occurs• Set the JOBID variable in the Extra environment variables box: <code>JOBID=<Job ID number of the job that you want to run></code>
SNMP Trap	<p>Select this method to send SNMP traps (notification) to a network manager host when the event occurs.</p> <p>Before you can use this method for sending notifications, you must complete the following configuration tasks:</p> <ul style="list-style-type: none">• Copy the file nvnotifications.mib to the “.mib” files directory on the host system running the network management software. The “.mib” (Management Information Base) file describes the format of SNMP Traps that NetVault sends. You can find nvnotifications.mib in the etc directory under the NetVault installation directory. For more information about the location of the “.mib” files on the host, see the relevant network management software documentation.• Configure the network manager host settings from the Change Settings page. For more information, see Configuring network manager host settings for notifications.

- 4 Click **Save** to save the global notification profile.

Using custom notification methods

This section includes the following topics:

- [About custom notification methods](#)
- [Creating a custom notification method](#)
- [Setting up a custom notification method](#)

About custom notification methods

Custom Notification Methods are created using external scripts, and typically used to perform some tasks when an event occurs in the NetVault Domain.

The script file contains header and body. The header specifies the script name and provides a brief description about the script. The script body contains the commands. You can use any text editor to create the script file. The script file should reside in the **global** subdirectory under the NetVault **scripts** directory (<NetVault home>\scripts\global on Windows and <NetVault home>/scripts/global on Linux).

Only users with a good knowledge of scripting using NetVault executable files and environmental variables should use this feature.

Creating a custom notification method

To create a custom notification method:

- 1 Open a new script file in a text editor.
- 2 Create the script header:
 - Specify the method name on the first line. Enclose the name within the NVNAMESTART and NVNAMEEND tags.

On Linux and UNIX platforms, specify the interpreter name on the first line of the script (for example, #!\bin\sh).
 - To include a description about the script, enclose the text within the NVCOMMENTSTART and NVCOMMENTEND tags. You can use multiple lines of text. Script description is optional.
 - Begin each header line with a REM tag on Windows and the “#” character on Linux and UNIX to mark the line as a comment.

Example (Windows):

```
REM NVNAMESTART My Custom MethodNVNAMEEND  
  
REM NVCOMMENTSTART Start a new  
REM occurrence of this event.  
REM NVCOMMENTEND
```

- 3 Create the script body and specify the syntax for the commands. You can open the NetVault scripts in the **util** directory using any text editor to view the proper structure and use them as a template for custom scripts.

To use NetVault scripts residing in the **scripts** directory, provide the complete path:

- **Windows:** %NVHOME%\scripts\<<command>
- **Linux and UNIX:** \$NVHOME/scripts/<command>

Verify that the script has a valid exit status.

- 4 Save the script:
 - On Windows, assign “.bat” extension to the file and save it in the <NetVault home>\scripts\global directory.
 - On Linux, assign “.sh” extension to the file and save it in the <NetVault home>/scripts/global directory.

Setting up a custom notification method

To set a custom notification method:

- 1 In the Navigation pane, click **Configure Notifications**.
- 2 On the **Editing Global Notification Profile** page, open the event class, and then open the event type that you want to be notified about.
- 3 Select the custom notification method.
- 4 In the **Notification Method** box, set the environment variables:
 - To specify multiple arguments, you can use a Comma (“,”), Carriage Return (CR), or Line Feed (LF) as a delimiter.

Example: Comma as a delimiter

```
ARG0=NVTEST.NV_Test_SP_v2, ARG1=1, ARG2=NETVAULT, ARG3=NVRelease, ARG4=1,  
ARG5=1, ARG6=p_return_msg, ARG7=p_status
```

Example: CR or LF (that is, a new line) as a delimiter

```
ARG0=NVTEST.NV_Test_SP_v2  
  
ARG1=1  
  
...  
  
ARG7=p_status
```

- To include special characters in the value string, use the Exclamation Mark (“!”) escape character:
 - To specify Comma (“,”) in the value string, enter “!,”
 - To specify Exclamation Mark (“!”) in the value string, enter “!!”
 - To specify Equals (“=”) in the value string, enter “!=”

You cannot use Escape characters in the name string.

- 5 Click **Save** to save the global notification profile.

Removing a user-defined job event

You can delete a user-defined job event if is no longer required. This task can be performed from the **Editing Global Notification Profile** page.

To remove a user-defined job event:

- 1 In the Navigation pane, click **Configure Notifications**.
- 2 On the **Editing Global Notification Profile** page, open the **Jobs — User defined** event class.
- 3 Select the event that you want to delete, and click **Delete**.
- 4 In the confirmation dialog box, click **OK**.

Removing a user-defined report event

You can delete a user-defined report event if is no longer required. This task can be performed from the **Editing Global Notification Profile** page.

To remove a user-defined job event:

- 1 In the Navigation pane, click **Configure Notifications**.
- 2 On the **Editing Global Notification Profile** page, open the **Report Job** event class.
- 3 Select the event that you want to delete, and click **Delete**.
- 4 In the confirmation dialog box, click **OK**.

Reporting in NetVault

- [Reporting system overview](#)
- [Working with reports](#)
- [Available reports](#)

Reporting system overview

The NetVault reporting system provides several canned or predefined reports that gives you quick access to information about your backup system. You can view reports on demand or receive reports by email. The reporting system also lets you export reports to PDF files.

You can use the reporting system in the following ways:

- **Generate and view reports on demand:** You can run and view reports from the **View Reports** page. The reports are generated in HTML format and the information is typically presented in tables.
- **Customize report views:** After a report is generated, you can modify the table views to include or exclude columns, change the column order, and add column totals and other aggregate values. You can save these settings to the job definition so that the custom format is used every time you run the report.
- **Use charts in reports:** You can add charts to your reports to present the information in visual formats. You can choose bar charts or pie charts.
- **Use notification events to send reports by email:** You can configure user-defined events and set up the email notification method or create user notification profiles to distribute reports by email. The reports can be delivered as email attachments in PDF or HTML format. Default format is PDF.
- **Schedule reports:** You can create report schedules to generate the reports once or on a recurring basis. NetVault lets you distribute the generated reports by email or print the reports through the notification system.

i **NOTE:** The following privileges are required to use the reporting system:

- **Run and view reports:** Reports — View and run reports
- **Customize report views:** Reports — Modify/edit report jobs and components

You can also use the **nvreport** utility to run and view reports from the command-line interface. For more information about this utility, see the *Quest NetVault CLI Reference Guide*.

Working with reports

This section includes the following topics:

- [Generating reports](#)
- [Setting a favorite report](#)
- [Setting filters for report](#)
- [Editing report job definitions](#)
- [Using notification methods to send reports by email or print reports](#)

- [Customizing table views for reports](#)
- [Adding charts to reports](#)
- [Exporting reports to PDF files](#)

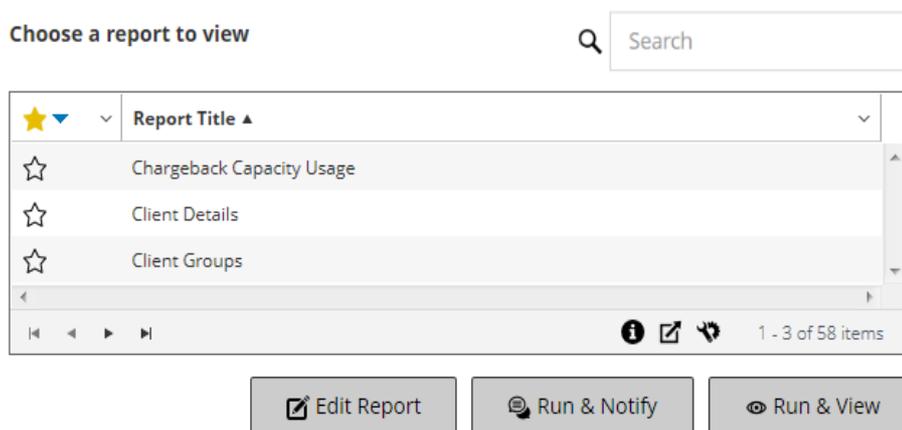
Generating reports

You can access the reports from the **View Reports** page. You can run and view reports on demand or create report schedules to automatically generate reports at specified times. NetVault also allows use of events and notification methods to send reports by email or print reports.

To generate a report:

- 1 In the Navigation pane, click **View Reports**.
- 2 On the **View Reports** page, select the report that you want to generate. For information about available predefined reports, see [Available reports](#).

Figure 29. View Reports page



- 3 To run and view the report immediately, click **Run & View**.

— or —

To raise the associated events and receive a notification, click **Run and Notify**.

i NOTE: For information about creating report schedules and configuring user-defined events for report jobs, see [Editing report job definitions](#). For information about notification methods for report events, see [Using notification methods to send reports by email or print reports](#).

To view the page size setting, column sort order, applied filters in the table, export the records, or edit table settings, click the icons at the lower-right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

You can use the Search option to filter the table data and quickly find entries that contain the specified string in any column values. You can also include wildcards (“?” or “*”) in the search filter string.

- 4 If the report includes any filter conditions (for example, start date, job ID number, client name, and others), the **Set filters for report** dialog box is displayed.

Set the filters that you want to use, and click **OK**. For more information about filter options, see [Setting filters for report](#).

- 5 The report is displayed in a new browser window.

You can customize the table views and add charts to your reports. For more information, see the following topics:

- [Customizing table views for reports](#)
- [Adding charts to reports](#)

You can also export the report to a PDF file. For more information, see [Exporting reports to PDF files](#).

Setting a favorite report

On **View Reports** page, you can select any report (s) to mark it as a favorite report to easily access it from the list of predefined reports.

To set a favorite report:

- 1 In the Navigation pane, click **View Reports**.
- 2 On the **View Reports** page, click the empty star next to the report in the first column that you want to favorite.

The favorite reports are displayed on top of the list in alphabetical order, followed by the remaining reports.

To remove a favorite from the reports list, click a filled-in star next to the report.

Setting filters for report

When you run a canned report that includes any filter conditions, the **Set filters for report** dialog box is displayed. You can set one or more conditions, based on which the reporting system generates the output. For each filter field that you want to use, you must select the filter operator and specify the comparison value. For example, the following **Set filters for report** dialog box is displayed for the report- Data Stored By Client.

Figure 30. Set filters for report dialog box

Inclusion filter values for component 'Data Stored By Client'

Start Date (Date)	>=	TODAY-7DA	<input type="checkbox"/> State
Start Date (Date)	<=	TODAY	<input type="checkbox"/> State <input type="checkbox"/> As above
Client Name (St)	=	*	<input type="checkbox"/> State <input type="checkbox"/> Regex

To set report filters:

- 1 In the operator list corresponding to the filter field, select the comparison operator. You can use the following comparison operators: =, !=, >, <, >= or <=.
- 2 In the associated box, type the comparison value. The value must match the data type of the field.

Fields can be of the following types: date, integer, string, or time.

To specify date values, you can use the following formats:

- YYYY/MM/DD
- YYYYMMDD
- **Relative date:** TODAY-n[time variable]

You can use the following time variables: YE = Year, MO = Month, WE = Week, DA = Date, HO = Hour, MI = Minute, and SE = Second

Example: TODAY-7DA

To specify time values, you can use the following formats:

- HH:MM:SS
- HHMMSS
- **Relative time:** NOW-n[timevariable] or TODAY-n[time variable]

You can use the following time variables: YE = Year, MO = Month, WE = Week, DA = Date, HO = Hour, MI = Minute, and SE = Second.

Example: NOW-12HO

3 Some reports may also include the following filter options:

- **State:** You can select this check box to compare the field state as opposed to the value it contains. The field state comparison option is designed for advanced users who have a good understanding of the NetVault reporting system.

You can set the field status to one of the following values:

- Normal
- N/A
- Unknown
- Never
- Unlimited

For example, you can specify the state **Unknown** to find or exclude records that are not available in one of the tables.

- **Regexp:** You can select this check box to match a regular expression instead of a constant value. The expression can contain text and wildcard characters.
- **As above:** The **As above** check box is shown when the same filter field is applied to multiple report components. You can select this check box to use the same comparison value that is configured for the preceding field.

4 After setting the filters, click **OK** to close the dialog box.

Editing report job definitions

You can edit a report job definition to modify the report schedule or configure user-defined events for the report. NetVault lets you distribute the generated reports by email or print the reports through the notification system.

To edit a report job:

- 1 In the Navigation pane, click **View Reports**.
- 2 On the **View Reports** page, select the report that you want to edit, and click **Edit Report**.
- 3 On the **Edit Report Job Definition** page, configure the following options.

Table 108. Edit report job definition

Option	Description
Schedule	<p>You can use this option to run the report at specified times. The scheduling option is useful for generating reports on a regular basis.</p> <p>To configure a report schedule, select an existing Schedule Set, or click Create New, and specify the schedule type and schedule method. For more information, see Creating Schedule Sets.</p> <p>For scheduled events, make sure that you specify the user-defined event in the Report Completed Event and Report Failure Event field.</p>
Report Completed Event	Specify the event that you want to raise when the job completes successfully.
Report Failure Event	Specify the event that you want to raise when the job fails.

i | **NOTE:** Report events are only raised if you specify a report schedule or use the **Run & Notify** method.

- 4 Click **Save**, and in the **Save Report Job** dialog box, click **OK**.

When you save the job, the configured events are added to the **Report Job** event class. You can use the global notification methods or user notification profiles to receive event notifications when the report is generated. For more information, see [Using notification methods to send reports by email or print reports](#).

Using notification methods to send reports by email or print reports

You can configure user-defined events for a report job and receive notifications when a job completes successfully or fails. NetVault adds these events to the **Report Job** event class. You can use this feature to send reports by email or to send the reports to a designated printer.

Depending on whether you want to distribute the report by email or print the report, do the following:

- To send the report to the Administrator, set up the **Sysop E-mail** notification method. For more information, see [Setting up a global notification method](#). By default, the report is generated in PDF file format and delivered as an email attachment.
- To send the report to other users, set up the notification profiles for the users. For more information, see [Setting up user notification profile](#). By default, the report is generated in PDF file format and delivered as an email attachment.
- To print the report, set up the **Print a report** notification method. For more information, see [Setting up a global notification method](#).

Customizing table views for reports

After a report is generated, you can modify the table views to include or exclude columns, change the column order, and add column totals and other aggregate values. You can add multiple table views for a component. The reporting system lets you save these settings to the job definition so that the custom format is used every time you run the report.

To customize table views for reports:

- 1 Access the report that you want to customize. For instructions on generating a report, see [Generating reports](#).
- 2 In the report window, click the **Settings** link at the upper-right corner of the table. If the report contains multiple components, click the link associated with the table that you want to modify.
- 3 In the **Table Settings** dialog box, the available columns are listed. The columns that have the corresponding **Show** check box selected are included in the table view. For example, the following **Table Settings** dialog box is displayed for the report- Data Stored By Client.

Figure 31. Table Settings dialog box

Start Time	<input checked="" type="checkbox"/> Show		
Start Date	<input checked="" type="checkbox"/> Show		
Run Length	<input checked="" type="checkbox"/> Show		
End Time	<input checked="" type="checkbox"/> Show		
End Date	<input checked="" type="checkbox"/> Show		
Job ID	<input checked="" type="checkbox"/> Show		
Client Name	<input checked="" type="checkbox"/> Show		
Transfer Size	<input checked="" type="checkbox"/> Show	Aggregate	None ▼
Plugin	<input checked="" type="checkbox"/> Show		
Selection Set	<input checked="" type="checkbox"/> Show		
Selection options	<input checked="" type="checkbox"/> Show		
Schedule set	<input checked="" type="checkbox"/> Show		
Backup target set	<input checked="" type="checkbox"/> Show		
Advanced Options Set	<input checked="" type="checkbox"/> Show		

You can make the following modifications:

- **Hide columns:** To hide one or more columns, clear the check marks for those columns.
- **Show columns:** To show hidden columns, click the corresponding **Show** check boxes.
- **Display aggregate values:** To add fields that display the total, minimum, maximum, and other aggregate values for the columns, select the field type in the **Aggregate** list. This list is only displayed for columns that contain numeric values.

The available options include the following:

- **Total:** Displays the sum of all the values in the column.
- **Average:** Displays the average value of all the values in the column.
- **Min:** Displays the maximum value of all values in the column.
- **Max:** Displays the minimum value of all values in the column.
- **Count:** Displays the count of records.
- **None:** Removes the aggregate field.

Click **Apply** to apply the settings to the table view and close the dialog box.

- 4 To change the column order, drag the header of the column to its new location.
- 5 To change the column width, drag the separator until the column is the width that you want.
- 6 To add a table, click the **Add another view** link at the lower-left corner of the component table.

If the component only supports table views, a new table is automatically added. If different view types are supported, then the **Add a view** dialog box is displayed. To add a table, select **Table** in the **Type of view** list.

The new table includes all columns. To customize the table, see [Step 2](#) and [Step 3](#).

- 7 To remove a view, click the **Remove this view** link at the lower-left corner of the component table.
 - 8 To save the custom settings, click the **Save report format** link at the upper-left corner of the report.
- The custom settings are saved to the report job definition and a message is displayed.

Adding charts to reports

You can use the graphical capabilities of the reporting system to view the output in the form of bar charts and pie charts. You can use multiple chart views to compare different sets of data. The reporting system lets you save these settings to the job definition so that the custom format is used every time you run the report.

i | **NOTE:** Charts are only available to report components that contain numeric values.

To add charts to a report:

- 1 Access the report that you want to customize. For instructions on generating a report, see [Generating reports](#).
- 2 In the report window, click the **Add another view** link at the lower-left corner of the table. If the report contains multiple components, click the link associated with the component for which you want to add a chart.

The **Add a view** dialog box is displayed if the component supports graphical views.

- 3 Select the chart type, and configure the applicable options. You can add bar charts and pie charts.

Table 109. Adding charts to reports

Chart type	Options
Bar charts	<p>You can use bar charts to provide a visual comparison of data values.</p> <p>To add a bar chart, do the following:</p> <ul style="list-style-type: none"> • In the Type of view list, select Barchart. • In the Category list, select a column to represent the vertical axis. A bar is generated for each record in the category column. The column values are used as labels. • In the Value list, select a column to represent the horizontal axis. The data values determine the length of the bars. <p>Click OK.</p> <p>Figure 32 displays a bar chart view.</p> <p>NOTE: The chart displays individual bars for the first 20 records. The remaining records are grouped in the chart as Other.</p>
Pie charts	<p>You can use pie charts to show the relative contribution of different categories to the total.</p> <p>To add a pie chart, do the following:</p> <ul style="list-style-type: none"> • In the Type of view list, select Piechart. • In the Category list, select a column to represent the categories or parts of the pie. The number of records in the category column determines the number of slices in the chart. The column values are used as legends. • In the Value list, select the column that contains the data for the pie chart. The data values determine the size of the slices. <p>Click OK.</p> <p>Figure 33 displays a pie chart view.</p> <p>NOTE: The chart displays individual slices for the first 20 records. The remaining records are grouped in the chart as Other.</p>

- 4 To remove a view, click the **Remove this view** link at the lower-left corner of the component table.
- 5 To save the custom settings, click the **Save report format** link at the upper-left corner of the report.
The custom settings are saved to the report job definition and a message is displayed.

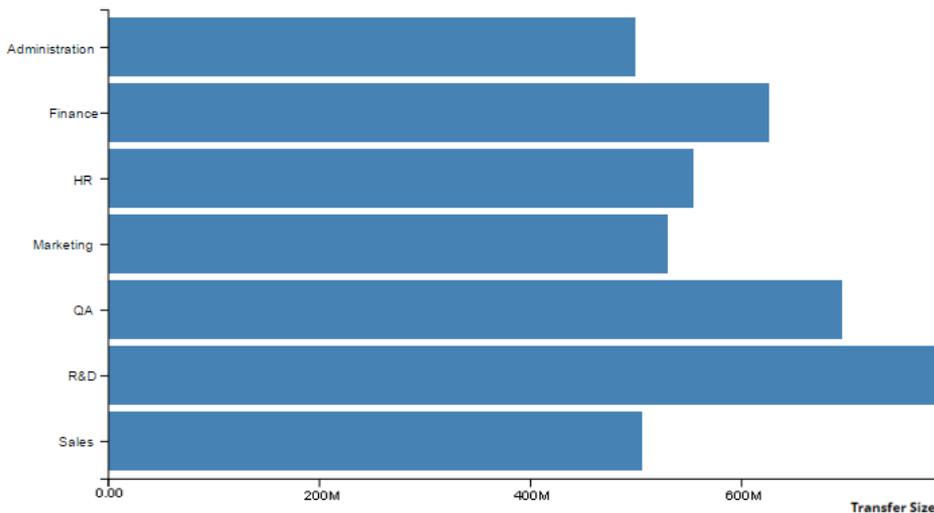
Chart view examples

This section includes some sample charts generated using the reporting system.

Bar chart view

The following figure displays the bar chart view for a sample data set.

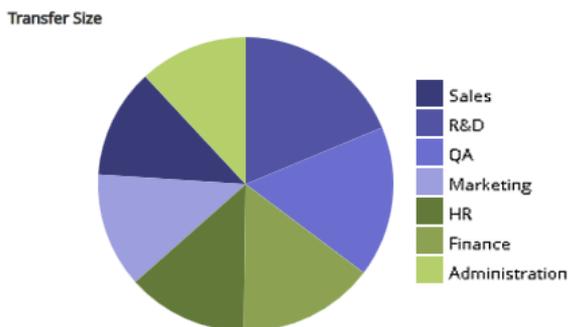
Figure 32. Bar chart view



Pie chart view

The following figure displays the pie chart view for a sample data set.

Figure 33. Pie chart view



Exporting reports to PDF files

After a report is generated, you can choose to export the report to a PDF file. The file is displayed in a browser window. You can download, save, or print the file, as required.

To export a report to a PDF file:

- 1 Access the report that you want to export. For instructions on generating a report, see [Generating reports](#).
- 2 In the report window click the **Export as PDF** link at the upper-left corner of the report.
- 3 The report is converted to a PDF file and displayed in the browser window. Use the applicable browser option to download, save or print the file.

Available reports

NetVault provides the following types of predefined reports listed on **View Reports** page in alphabetical order.

Table 110. Predefined reports

Report title	Description
Chargeback Capacity Usage	<p>This report lets you view the storage usage information for existing client groups. It can be used for chargeback calculations based on storage usage. The report includes the following sections:</p> <ul style="list-style-type: none">• Total Data Stored by Client Group: For each client group, the report shows the following details:<ul style="list-style-type: none">- Client Group: Name of the client group.- Client Group Description: Description of the client group.- Total Data Stored: Total amount of data stored by the member clients. It shows the total size of all backups, including the total size of Duplicate and Data Copy backups stored by the member clients.- Total Source Data Protected: Total amount of source data protected by the member clients. It shows the total size of the last instance of all Full Backups stored by the member clients.• Client Groups: This section provides information about existing client groups. The table is sorted by client group name. You can view the following details: group name, group description, whether the group includes all clients, and list of members (only displayed for groups that do not have the All clients option selected).
Client Details	<p>This report provides information about existing NetVault Clients. The table is sorted by client name.</p> <p>You can view the following details: client name, machine ID number, client type (OS), description, NetVault version, NetVault release, client accessibility status, and client status.</p>
Client Groups	<p>This report provides information about existing client groups. The table is sorted by client group name.</p> <p>You can view the following details: group name, group description, whether the group includes all clients, and list of members (only displayed for groups that do not have the All clients option selected).</p>
Client Statuses	<p>You can use this report to view the status of existing NetVault Clients.</p> <p>The table includes the following information: client name, NetVault version, accessibility status, and client status.</p>

Table 110. Predefined reports

Report title	Description
Configuration	<p>The Configuration report provides a consolidated view of clients, job definitions, sets, users, notifications, and licenses.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> <p>• Client Details: This section provides information about existing NetVault Clients.</p> <p>You can view the following information: NetVault machine name, machine ID number, product version, machine type (server or client), and client status.</p> <p>• Backup Job Definitions: This section shows the definitions of existing backup jobs.</p> <p>You can view the following information: job ID number, job title, client name, and set names (Backup Selection Set, Backup Options Set, Schedule Set, Target Set, and Backup Advanced Options Set).</p> <p>• Scheduled Jobs: This section displays the scheduled jobs.</p> <p>The table shows the job ID number and job title of scheduled jobs.</p> <p>• Selection Sets: This section provides information about the available Backup Selection Sets.</p> <p>You can view the set name, filer name (if applicable), and items that are included in the set.</p> <p>• Selection Option Sets: This section provides information about the available Backup Options Sets.</p> <p>The table includes the following details: set name, selection options, selection options for the Web Service, and plug-in details.</p> <p>• Schedule Sets: This section provides information about the available Schedule Sets.</p> <p>The table includes the following details: set name, schedule type, method, days of week, repeat period, scheduled date, scheduled time, weeks of month, and number of job retries.</p> <p>• Backup Targets: This section provides information about the available Backup Target Sets.</p> <p>The table shows the set name, device selections (any device or the device tree for selected devices), and media options (media group, media labeling and reusability settings, and media request timeout settings).</p> <p>• Advanced Option Sets: This section provides information about the available Backup Advanced Options Sets.</p> <p>You can view the following information: set name, retirement settings, deduplication setting, Secondary Copy details (sets used to define the copy, retirement settings, and other information), and user-defined events.</p> <p>• User Details: This section provides information about the existing NetVault users.</p> <p>The table displays the user name and email ID.</p>

Table 110. Predefined reports

Report title	Description
	<ul style="list-style-type: none"> • User's Notifications (Configuration): This section provides information about user notification profiles. The table includes the following details: notify class, notify event, account name, real name of user, notification method, and environment variables. • Global Notifications (Configuration): This section provides information about global notification profiles. The table includes the following details: notify class, notify event, notification method, and environment variables. • Licenses: This section shows the license capabilities and usage details for various components, including clients, libraries, NDMP devices, tape drives, optical drives, and slots. The table also shows the online and total tape capacity details and the license information for heterogeneous clients.
Data Stored by Client	<p>You can use this report to view the amount of data stored by existing NetVault Clients during a specified period.</p> <p>By default, the report displays the data for the last 7 days; you can set the report filters to change the period. You can also set the Client Name filter to view the details for specific clients.</p> <p>The table includes the following details: Start Time, Start Date, Run Length, End Date, End Time, Job ID, Client Name, Transfer Size, Plugin, and set names (Selection Set, Selection Options, Schedule set, Backup Target Set, and Advanced Options Set).</p>
Disk Storage Devices — General	<p>You can use this report to view information about existing disk-based storage devices.</p> <p>The table includes the following details: device name, host name, media group, device status, licensed capacity, actual protected capacity, staging space, deduplication store space, and total space available.</p>
Expired Offline Media	<p>This report provides information about retired offline media items. The table is sorted by media label.</p> <p>You can view the following details: media label, media group label, barcode, total data written, last read date, offsite location, and marked for reuse.</p>
Failed ULA Requests	<p>This report provides information about user requests that were denied due to insufficient privileges.</p> <p>For failed user requests, you can view the following details: request date and time, user name, source machine, and type of operation</p> <p>By default, the report displays the data for the last 7 days. You can set the report filters to change the period.</p>

Table 110. Predefined reports

Report title	Description
Front Side Data	<p>This report shows the amount of source data protected for client machines and the storage utilization details for disk-based devices.</p> <p>By default, the report includes the backups that were run in the last 28 days. You can set the report filters to change the period.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> <p>Front Side Data: This section displays the Phase 1 Full Backups performed during the specified period. It shows the last instance of a job. The report does not include Incremental Backups, Consolidated Backups, or Phase 1 or Phase 2 Duplicate or Data Copy Backups. You can view the following details: job ID number, instance, job title, protected data size, plug-in name, phase, and backup type (if the type is Incremental).</p> <p>The Protected Data Size column shows backup size before de-duplication. You can add the Total field to this column to display the total size of the backups.</p> <p>RAS Devices — Storage Utilization: This section displays the storage utilization details for disk-based storage devices. You can view the following details: device name, media group, licensed and actual protected capacity, staging space, deduplication store space, and total space available.</p>
Full Online Storage	<p>This report shows which online storage devices and media items are full. This report includes the following sections:</p> <ul style="list-style-type: none"> <p>Full Online Media: This section contains information about online tapes that are full. You can view the following details: media label, media group label, barcode, last write and read date, media expiry date, library name, logical slot position, offsite location, needs import, space left, and marked for reuse.</p> <p>Full Online RAS: This section contains information about online disk-based devices that are full. You can view the device name and media group label.</p>
Global Notifications	<p>This report shows which events are included in the global notification profile. The report shows the event class, event type, notification method, and environment variables.</p>
Historic Jobs — by date	<p>You can use this report to view information about all backup and restore jobs that were performed during the specified period. The records are sorted by start date (newest to oldest).</p> <p>By default, the report includes the jobs that were run in the last 7 days. You can set the report filters to change the period.</p> <p>You can view the following details: start time, start date, run length, job title, job ID number, instance, client name, Backup Selection Set name, data transfer size, and job status.</p>
Historic Jobs — by size	<p>You can use this report to view information about all backup and restore jobs that were performed during the specified period. The records are sorted by data transfer size (largest to smallest).</p> <p>By default, the report includes the jobs that were run in the last 7 days. You can set the report filters to change the period.</p> <p>You can view the following details: start time, start date, run length, job title, job ID number, instance, client name, Backup Selection Set name, data transfer size, and job status.</p>

Table 110. Predefined reports

Report title	Description
Historic Jobs Duration — by date	<p>This report provides information about jobs that exceeded the specified run length. The records are sorted by start date (newest to oldest).</p> <p>By default, the run length limit is set to 8 hours; you can set the report filters to change the value. You can also set the filters to change the period, which is set to 7 days by default.</p> <p>You can view the following details: start time, start date, run length, job title, job ID number, instance, client name, Backup Selection Set name, data transfer size, and job status.</p>
Historic Jobs Encrypted — by date	<p>You can use this report to view information about all encrypted primary and secondary copy backups that were performed during the specified period. The records sorted by start date (newest to oldest). This report does not include jobs which were performed without encryption.</p> <p>By default, the report includes the jobs that were run in the last 7 days. You can set the report filters to change the period.</p> <p>You can view the following details: start time, start date, run length, job title, job ID number, instance, client name, Backup Selection Set name, data transfer size, and job status.</p>
Historic Jobs Encrypted — by size	<p>You can use this report to view information about all encrypted primary and secondary copy backups that were performed during the specified period. The records are sorted by data transfer size (largest to smallest). This report does not include jobs which were performed without encryption.</p> <p>By default, the report includes the jobs that were run in the last 7 days. You can set the report filters to change the period.</p> <p>You can view the following details: start time, start date, run length, job title, job ID number, instance, client name, Backup Selection Set name, data transfer size, and job status.</p>
Historic Jobs Failed — by date	<p>You can use this report to view information about jobs that failed during the specified period. The records, sorted by start date (newest to oldest).</p> <p>By default, the report includes the jobs that were run in the last 7 days. You can set the report filters to change the period.</p> <p>You can view the following details: start time, start date, run length, job title, job ID number, instance, client name, Backup Selection Set name, data transfer size, and job status.</p>
Historic Jobs Successful — by date	<p>You can use this report to view information about jobs that completed successfully during the specified period. The records are sorted by start date (newest to oldest).</p> <p>By default, the report includes the jobs that were run in the last 7 days. You can set the report filters to change the period.</p> <p>You can view the following details: start time, start date, run length, job title, job ID number, instance, client name, Backup Selection Set name, data transfer size, and job status.</p>
Historic Jobs Warnings — by date	<p>You can use this report to view information about jobs that completed with warnings during the specified period. The records are sorted by start date (newest to oldest).</p> <p>By default, the report includes the jobs that were run in the last 7 days. You can set the report filters to change the period.</p> <p>You can view the following details: start time, start date, run length, job title, job ID number, instance, client name, Backup Selection Set name, data transfer size, and job status.</p>

Table 110. Predefined reports

Report title	Description
Index Media	<p>This report displays the index media of a backup job.</p> <p>You can view the following details: Job title, job ID, job instance, job phase, client name, server name, plug-in name, backup time, backup date, media label, off-site location, and Media Online or Offline.</p>
Job Definitions	<p>This report displays the job definitions of backup, restore, and report jobs.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> <p>Backup Job Definitions: This section shows the definitions of backup jobs created by various users. The jobs are sorted by user name.</p> <p>You can view the following details: job ID number, job title, client name, and plug-in name, set names (Backup Selection Set, Backup Options Set, Schedule Set, Target Set, and Backup Advanced Options Set), and creator ID.</p> <p>Restore Job Definitions: This section shows the definitions of restore jobs created by various users. The jobs are sorted by user name.</p> <p>You can view the following details: job ID number, job title, client name, plug-in name, set names (Restore Selection Set, Restore Source Set, Schedule Set, and Restore Advanced Options Set), and creator ID.</p> <p>Report Job Definitions: This section provides definitions of report jobs.</p> <p>You can view the following details: job ID number, job title, client name, schedule set, and creator ID.</p> <p>NOTE: Canned or predefined reports always run in the context of “default” user.</p>
Library Contents	<p>You can use this report to view information about online tape media items. The report also shows the blank tape count for existing libraries.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> <p>Libraries’ Media Contents: This section shows information about online tapes.</p> <p>The table includes the following details: library name, drive name, slot position, media label, barcode, group label, space used, and space left.</p> <p>Blank Media Items In Libraries: This section shows the number of blank tapes available in existing libraries.</p>
Media General	<p>This report provides general information about available tape media items. The table is sorted by media label.</p> <p>You can view the following details: media label, barcode, media type, format, machine name, media expiry date, offsite location, space used, space left, usable or not, and read-only or not.</p>
Media Quotas and Usage	<p>You can use this report to view the media quota and media usage information for existing NetVault users.</p> <p>The table includes the following details: user name, media quota, and media used.</p>

Table 110. Predefined reports

Report title	Description
Monthly Job Summary	<p>You can use this report to view the overall status of backup jobs performed in the last 30 days. The report also shows the jobs that are currently in progress and jobs that failed during the current month.</p> <p>The Monthly Job Summary report includes the following sections:</p> <ul style="list-style-type: none"> • Summary: The Summary section displays the following information: <ul style="list-style-type: none"> - Total amount of data stored - Number of backup jobs that completed successfully - Number of backup jobs that failed - Number of backup jobs that completed with warnings • Currently Active Jobs: This section shows the jobs that are currently in progress. <p>The table includes the following information: job ID number, job title, client name, policy name, and job status.</p> • Failed Month Backup Jobs: This section shows the jobs that failed during the current month. <p>The table includes the following information: job ID number, job title, client name, policy name, and job status.</p>
NDMP Jobs	<p>You can use this report to view all jobs performed using the Plug-in <i>for NDMP</i> or Plug-in <i>for SnapMirror to Tape</i>. The records are sorted by filer name and start date.</p> <p>By default, the report includes the jobs that were run in the last 7 days. You can set the report filters to change the period.</p> <p>The table includes the following information: start time, start date, run length, job title, job ID number, instance, client name, filer name, plug-in name, Backup Selection Set name, data transfer size, transfer rate, and job status.</p>
NetVault Error Logs	<p>You can use this report to view the error log messages generated during the specified period.</p> <p>By default, the report includes the error messages generated in the last 7 days. You can set the report filters to change the period.</p> <p>The table includes the following details: date and time, client name, job ID, instance ID, error message, and warning level.</p>
NetVault Events	<p>You can use this report to view the NetVault events that occurred during the specified period. By default, the report includes the events generated in the last 7 days. You can set the report filters to change the period.</p> <p>The table includes the following details: date and time, event type, event class, event description, and event message.</p>
NetVault Logs	<p>You can use this report to view all log messages generated during the specified period.</p> <p>By default, the report includes the log messages generated in the last 7 days. You can set the report filters to change the period.</p> <p>The table includes the following details: job ID number, class, warning level, time, date, client name, and log message.</p>
NetVault Logs — Most Recent	<p>You can use this report to view the recent log messages.</p> <p>Only the logs generated in the last two days are read to improve the report processing time and reduce the use of system memory for the report generation activity.</p> <p>You can set the report filters to change the period.</p> <p>The table includes the following details: job ID number, class, warning level, time, date, client name, and log message.</p>

Table 110. Predefined reports

Report title	Description
Offline Devices	<p>This report shows which storage devices that are currently offline.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> • Offline Devices: The section contains information about tape-based devices that are offline. The table includes the following details: device name, host name, vendor, library name, and status. • Offline RAS Devices: The section contains information about disk-based devices that are offline. The table displays the device name, device machine (address:port), and device status.
Overnight Duplication Job Summary	<p>You can use this report to view the overall status of duplication jobs performed in the last 16 hours. The report also shows the duplication jobs that are currently in progress and the duplication jobs that failed during the period.</p> <p>The Overnight Duplication Job Summary report includes the following sections:</p> <ul style="list-style-type: none"> • Summary: The Summary section displays the following information: <ul style="list-style-type: none"> - Total amount of data stored by duplication jobs - Number of duplication jobs that completed successfully - Number of duplication jobs that failed - Number of duplication jobs that completed with warnings • Currently Active Duplication Jobs: This section shows the duplication jobs that are in progress. You can view the following information: job ID, job title, client name, policy name, and job status. • Failed Overnight Duplication Jobs: This section shows the duplication jobs that failed during the period. You can view the following information: job ID, Instance, job title, client name, policy name, and job status.
Overnight Job Summary	<p>You can use this report to view the overall status of backup jobs performed in the last 16 hours. The report also shows the jobs that are currently in progress and the jobs that failed during the period.</p> <p>The Overnight Job Summary report includes the following sections:</p> <ul style="list-style-type: none"> • Summary: The Summary section displays the following information: <ul style="list-style-type: none"> - Total amount of data stored - Number of backup jobs that completed successfully - Number of backup jobs that failed - Number of backup jobs that completed with warnings • Currently Active Jobs: This section shows the backup jobs that are in progress. You can view the following information: job ID number, job title, client name, policy name, and job status. • Failed Overnight Backup Jobs: This section shows the backup jobs that failed during the period. You can view the following information: job ID number, job title, client name, policy name, and job status.

Table 110. Predefined reports

Report title	Description
Overnight Job Summary (Extended)	<p>You can use this report to view the overall status of backup jobs performed in the last 24 hours. The report also shows jobs that are currently in progress, and the jobs that completed successfully or failed during the period.</p> <p>By default, the report includes jobs that were performed in the last 24 hours. You can set the Start systime and End systime filters to view jobs that ran during a specified period.</p> <p>The report includes the following components:</p> <ul style="list-style-type: none"> • Summary: The Summary section displays the following information: <ul style="list-style-type: none"> - Total amount of data stored - Number of backup jobs that completed successfully - Number of backup jobs that failed - Number of backup jobs that completed with warnings • Currently Active Jobs: This section shows the jobs that are in progress. <p>You can view the following information: job ID number, job title, client name, policy name, start date and time, and job status. The jobs are sorted by job ID number.</p> • Failed Overnight Backup Jobs: This section shows jobs that failed. <p>You can view the following information: job ID number, job title, client name, start date and time, end date and time, and job status.</p> • Completed Overnight Backup Jobs: This section shows the jobs that completed successfully. <p>You can view the following information: job ID number, job title, client name, policy name, start date and time, end date and time, job status, transfer size, transfer rate, run length, and file count.</p>
Overnight VMware Backup Summary	<p>You can use this report to view the overall status of virtual machine backup jobs that are completed successfully or completed with warning and performed in the last 16 hours.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> • Summary: The Summary section displays the following information: <ul style="list-style-type: none"> - Total amount of data stored by virtual machines - Number of virtual machines backed up - Number of virtual machines that backed up successfully - Number of virtual machines that failed • Overnight Failed VMs Details: This section shows the details of the failed virtual machine whose backup completed with warnings during the period. <p>You can view the following information: VM name, server, job title, job ID number, job instance, start time, start date, end time, end date, file indexing, and CBT.</p>

Table 110. Predefined reports

Report title	Description
Policies Summary	<p>This report provides information about the policy-based backups performed during the specified period.</p> <p>By default, the report includes the policies that were run in the last 7 days. You can set the report filters to change the period.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> • Total Policy Summary details: This section displays the total count of Policies and Policy jobs. The table shows the following details: Total count of policies, count of policies completed successfully, count of policies completed with warnings, count of failed policies, total count of policy jobs, Count of policy jobs completed successfully, count of policy jobs completed with warnings, count of failed policy jobs. • Policy Basics: This section displays all policy-based backups that were performed during the specified period. The table shows the following details: policy name, client count, job count, successful jobs, warning jobs, failed Jobs, status, client list, total transfer size, failure events, and warning events. • Successful Policy Jobs: This section displays the policy backups that completed successfully during a specified period. The table shows the following details: Job ID, last instance number of job, policy name, client name, job title, plug-in name, start time, start date, run length, end time, end date, selection set, transfer size, transfer rate, and job status. • Policy Jobs With Warnings: This section displays the policy backups that completed with Warning during a specified period. The table shows the following details: Job ID, last instance number of job, policy name, client name, job title, plug-in name, start time, start date, run length, end time, end date, selection set, transfer size, transfer rate, and job status. • Failed Policy Jobs: This section displays the policy backups that failed during a specified period. The table shows the following details: Job ID, last instance number of job, policy name, client name, job title, plug-in name, start time, start date, run length, end time, end date, selection set, transfer size, transfer rate, and job status.
Restore Summary	<p>This report provides a summary of restore jobs.</p> <p>The report shows the following information: start time, start date, client name, job title, job ID number, instance ID, run length, and job status.</p>

Table 110. Predefined reports

Report title	Description
Server Daily Summary	<p>You can use this report to view the daily backup summary for the server. The report also includes information about media items residing in the existing libraries.</p> <p>The report includes the following components:</p> <ul style="list-style-type: none"> • Jobs in Progress: This section shows the jobs that are in progress. The table includes the following information: client name, job title, policy name, start date, start time, transfer rate, transfer size, and job status. • Jobs Status Check: This section shows the jobs that failed. The table includes the following information: client name, job title, policy name, start date, start time, end date, end time, run length, transfer rate, transfer size, and job status. • Jobs Completed with Warnings: This section shows the jobs that completed with warnings. The table includes the following information: client name, job title, policy name, start date, start time, end date, end time, run length, transfer rate, transfer size, and job status. • Jobs Completed OK: This section shows the jobs that completed successfully. The table includes the following information: client name, job title, policy name, start date, start time, end date, end time, run length, transfer rate, transfer size, and job status. • Media: The report also provides information about the tapes residing in the specified libraries. The table includes the following information: barcode, media group label, media expiry date, expiry time, media label, slot position, reusable or not, space left, and space used.
Server License — Capabilities and Usage	<p>This report displays the license capabilities and usage details for clients, SmartClients, and various types of devices. The report also includes the license flag information for various components.</p>

Table 110. Predefined reports

Report title	Description
Single Job Definition	<p>This report provides information about ALL job options used by a single job ID. You can view the consolidated data for the single job with detailed description. The report includes the following components:</p> <ul style="list-style-type: none"> <p>• Job Definition: This section shows the job details and sub components included in the report for a job ID.</p> <p>The table includes the following information: Job ID, Job Type, Client Name, Job Title, Selection Set, plug-in Option (for a backup job), Source Set (for a restore job), Target Set, Schedule Set, and Advance Option Set.</p> <p>• Backup Selection Set: This section shows the information related to the backup selection set.</p> <p>The table includes the following information: Selection Set Name, Auto Created, Client Name, plug-in Name, Selection Tree (files selected in the drive or directory), and Operating System.</p> <p>• Restore Selection Set: This section shows the information related to the restore selection set.</p> <p>The table includes the following information: Restore Selection Set Name, Auto Created, Restore Selection set Type details, Restore Selection Tree (files selected in the drive or directory), Restore Options set, and Target Client Tree (client name or client tree included for restore).</p> <p>• Plugin Options: This section shows plug-in details for a backup job.</p> <p>The table includes the following information: Selection Set Name, and Backup Selection Set Option details.</p> <p>• Source Set: This section shows the set option for a restore job for source media selection.</p> <p>The table includes the following information: Source Set Name, Force Local Drives, and Device Tree.</p> <p>• Target Set: This section shows the target set details for a backup job.</p> <p>The table includes the following information: Backup Target Set Name, Auto Created, Any Device, Device Tree, Media Target, Media ID, Media Group, Auto Label Media, Restore Media, Minimum Space Specified, Minimum Space Required, Write Protected Media, and Force Local Drive.</p> <p>• Schedule Set: This section shows the schedule summary for the specified Job ID.</p> <p>The table includes the following information: Set Name, Run Hour, Run Minute, Start Day, Start Month, Start Year, Repeat Method, Repeat Frequency, Repeat Count, Week Number, Day Of Week, Delete always, Last Day of Month, Trigger Name, Priority, Retire on Failure, and Retry Delay.</p> <p>• Backup Life: This section shows the components for advanced backup selection set related to backup life.</p> <p>The table includes the following information: Set Name, Backup type, Keep Time, Offline Index Time, Duplication Life Length, Use Keep of Generation, Use Keep for Time, Number Generations, Time Interval units, Time Interval Count, Force Expiry, and Offline Index.</p> <p>• Additional Options: This section shows the components for advanced backup selection set related to additional options.</p> <p>The table includes the following information: Do Encryption, Do Duplication, Verify Backup, and User Compression.</p>

Table 110. Predefined reports

Report title	Description
	<ul style="list-style-type: none"> <p data-bbox="608 277 1378 331">• Secondary Copy: This section shows the components for advanced backup selection set related to secondary copy.</p> <p data-bbox="644 342 1390 483">The table includes the following information: Create Secondary Copy, Dup Run On Server, Dup Run On Client, Dup Schedule Set, Dup Target Set, Dup Source Set Dup Max Concurrent Sessions, Dup Media Request Time Out,Migrate, Dup Use Minmum Media, Dup Do Replication, and Dup Source Media First.</p> <p data-bbox="608 499 1378 553">• Pre & Post Script: This section shows the components for advanced backup selection set related to pre and post script.</p> <p data-bbox="644 564 1326 618">The table includes the following information: PreScriptName, PreScriptArgument, PostScriptName, and Post Script Argument.</p> <p data-bbox="608 633 1347 687">• Events: This section shows the components for advanced backup selection set related to events.</p> <p data-bbox="644 698 1374 752">The table includes the following information: Success Event,Warning Event, Failure Event, and Delete Always.</p> <p data-bbox="608 768 1337 822">• Restore Type: This section shows the components for advanced backup selection set related restore type.</p> <p data-bbox="644 833 1366 887">The table includes the following information: Restore From Selected Backup, Restore Form Latest Backup.</p> <p data-bbox="608 902 1358 983">• Additional Options (for Restore Job): This section shows the components for advanced restore selection set related to additional options.</p> <p data-bbox="644 994 1257 1048">The table includes the following information: Use Network Compression.</p> <p data-bbox="608 1064 1390 1144">• Pre & Post Script (for Restore Job): This section shows the components for advanced restore selection set related to Pre and Post scripts.</p> <p data-bbox="644 1155 1326 1209">The table includes the following information: PreScriptName, PreScriptArgument, PostScriptName, and Post Script Argument.</p> <p data-bbox="608 1225 1353 1279">• Events (for Restore Job): This section shows the components for advanced restore selection set related to events.</p> <p data-bbox="644 1290 1378 1344">The table includes the following information: Success Event, Warning Event, Failure Event, and Delete Always.</p>

Table 110. Predefined reports

Report title	Description
Single Job Summary	<p>This report provides information about a single job. The report includes job information, drive events, data transfer data, and media usage information. You can set the report filters to view information pertaining to one or multiple instances of the specified job.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> • Single Job Main Summary: This section provides the job details for the specified job. The table includes the following information: job ID number, job title, job type, plug-in name, instance ID, start date, start time, end date, end time, Job status, transfer size, and transfer rate. • Single Job’s Drive Events: This section shows the device events for the specified job. The table includes the following information: date, time, drive name, and event. • Single Job’s Media Transfers: This section shows the data transfer details for the specified job. The table includes the following information: received date, received time, started date, started time, finished date, finished time, transfer type (read or write) and amount of data transferred. • Media used by single backup: This section lists the media items used by the specified job. The table includes the following information: job ID number, instance, barcode, media label, barcode, media type, and offsite location.
Single Policy Summary	<p>This report provides information about a single backup policy. The report provides the policy status summary, client and job details, failed jobs, and data transfer summary.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> • Policy status information: This section provides the status summary. The table includes the following information: policy name, clients count, job count, successful policy job count, warning policy job count, failed policy job count, status of policy, client list, total transfer size, failure events, and warning event. • Policy clients: This section provides the details of policy jobs on each client included in the policy-based backup. The table shows the client name, status, job count, successful jobs, warning jobs, failed jobs, and total transfer size. • Successful jobs within single policy: This section displays the policy backups that completed successfully during a specified period. The table shows the following details: Job ID, last instance number of job, policy name, client name, job title, plug-in name, start time, start date, run length, end time, end date, selection set, transfer size, transfer rate, and job status.

Table 110. Predefined reports

Report title	Description
Single User's Audit Trail	<ul style="list-style-type: none"> • Warning jobs within single policy: This section displays the policy backups that completed with warnings during a specified period. The table shows the following details: Job ID, last instance number of job, policy name, client name, job title, plug-in name, start time, start date, run length, end time, end date, selection set, transfer size, transfer rate, and job status. • Failed jobs within single policy: This section displays the policy backups that failed during a specified period. The table shows the following details: Job ID, last instance number of job, policy name, client name, job title, plug-in name, start time, start date, run length, end time, end date, selection set, transfer size, transfer rate, and job status. <p>This report provides information about all attempted operations by a specified user.</p> <p>You can view the following details: date, time, source machine, account, operation type, and request status (granted or not).</p> <p>By default, the report shows the messages that were generated in the last 7 days. You can set the report filters to change the period.</p>
Storage Contents Query	<p>This report provides information about the contents stored in tape media items.</p> <p>By default, the report includes all existing tape media items. You can view information about specific media items by setting one or more filter options. The available filter options are client name, media label, media group label, job title, and plug-in name.</p> <p>You can view the following details: job title, client name, transfer size, media label, barcode, media group label, backup date, backup time, segment expiry date, segment expiry time, and plug-in name.</p>
Storage Segment Contents Query	<p>This report provides information about the contents stored in individual tape segments.</p> <p>You can view the following details: media label, barcode, job title, client name, plug-in name, backup date, backup time, and segment length.</p>
Storage Utilization	<p>You can use this report to view the utilization details for disk- and tape-based storage devices.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> • Media Utilization: This section shows the used and free space details for tape-based devices. The table includes the following information: media label, barcode, media group label, space used, and space left. • Blank Media Items In Libraries: This section shows the number of blank tapes available in existing libraries. • RAS Devices — Storage Utilization: This section shows the used and free space details for disk-based devices. The table includes the following information: device name, media group, licensed and actual protected capacity, staging space, deduplication store space, and total space available.
Storage Utilization — RAS	<p>You can use this report to view the utilization details for disk-based storage devices.</p> <p>The table includes the following information: device name, media group, licensed and actual protected capacity, staging space, deduplication store space, and total space available.</p>

Table 110. Predefined reports

Report title	Description
Tape Return	<p>You can use this report to get information about retired offline media items. You can filter the report based on Media Group Label and Offsite Location.</p> <p>The table shows the following details: media label, media group label, barcode, last write and read date, offsite location, number of times the tape has been initialized, number of write errors, and marked for reuse.</p>
Tapes to Send Offsite	<p>You can use this report to get information about tapes that are available for offsite storage for disaster recovery purposes. You can filter the report based on the Last Write Date and Media Group Label options.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> <p>Media to Send Offsite: This section displays the media items that are available for offsite storage.</p> <p>The table includes the following information: media label, media group label, barcode, last write date and time, media expiry date, library name, slot position, amount of data stored, and number of segments stored.</p> <p>Media Segments (to Send Offsite): This report provides information about the contents stored in individual tape segments.</p> <p>The table shows the following information: barcode, client name, job title, job ID number, instance, backup date, backup time, segment length, segment position, and segment expiry date.</p>
User Details	<p>This report provides information about existing NetVault users.</p> <p>The table includes the following fields: user name, real name, description, location, email ID, other contact details, last logon date and time, and access type (if local access only).</p>
User Notifications	<p>This report shows which events are included in the notification profiles for users.</p> <p>The table shows the following information: notify class, notify event, account name Real name of the user, notification method, and environment variables.</p>
User Privileges	<p>You can use this report to view the privileges granted to individual user accounts.</p> <p>The table displays the account name and list of privileges granted to the user (All, if the user has all privileges).</p>
User-Defined Event Types	<p>You can use this report to view information about user-defined events.</p> <p>The table shows the event name and event description.</p>
VMware Backup Summary	<p>You can use this report to view the backup summary of the virtual machine backup jobs that are competed successfully or completed with warning. You can use the filter option to specify the Date.</p> <p>The report includes the following sections:</p> <ul style="list-style-type: none"> <p>Summary: The Summary section displays the following information:</p> <ul style="list-style-type: none"> - Total amount of data stored by virtual machines - Number of virtual machines backed up - Number of virtual machines that backed up successfully - Number of virtual machines that failed <p>Failed VMs Details: This section shows the details of the failed virtual machine whose backup completed with warnings during the period.</p> <p>You can view the following information: VM name, server, job title, job ID number, job instance, start time, start date, end time, end date, file indexing, and CBT.</p>

Table 110. Predefined reports

Report title	Description
VMware Job History	<p>This report provides information about virtual machine backup jobs. You can use the filter options to specify the job ID number and instance number.</p> <p>The report shows the following details for each virtual machine included in the job:</p> <ul style="list-style-type: none"> • VM Name: Name of the virtual machine. • Server: Name of the vCenter or ESX/ESXi Server. • Job Title: Name of the job. • Start Time: Time when the first backup was attempted for the virtual machine. • Start Date: Date when the first backup was attempted for the virtual machine. • end Time: Time when the last backup was attempted for the virtual machine. • end Date: Date when the last backup was attempted for the virtual machine. • Job ID: Job ID number. • Job Instance: Job instance number. • Transferred Size: Amount of data transferred. • Backup State: Shows the status of the virtual machine backup (true = backup is successful; false = backup has failed). • File indexing: Shows the status of file-level indexing (true = file-level indexing is successful; false = file-level indexing has failed or file-level indexing is disabled). • CBT: Shows if CBT is enabled for the virtual machine (true = CBT is enabled; false = CBT is disabled).
Weekly Job Summary	<p>You can use this report to view the overall status of backup jobs performed in the last 7 days. The report also shows the jobs that are currently in progress and the jobs that failed during the period.</p> <p>The Weekly Job Summary report includes the following sections:</p> <ul style="list-style-type: none"> • Summary: The Summary section displays the following information: <ul style="list-style-type: none"> - Total amount of data stored - Number of backup jobs that completed successfully - Number of backup jobs that failed - Number of backup jobs that completed with warnings • Currently Active Jobs: This section shows the backup jobs that are currently in progress. <p>You can view the following information: job ID number, job title, client name, policy name, and job status.</p> • Failed Week Backup Jobs: This section provides a list of backup jobs that failed during the period. <p>You can view the following information: job ID number, job title, client name, policy name, and job status.</p>

Working with client clusters

- [About client cluster support](#)
- [Device configuration in cluster environment](#)
- [Installing and upgrading cluster-aware plug-ins](#)
- [Configuring a cluster-aware plug-in](#)
- [Managing virtual clients](#)
- [Backups using cluster-aware plug-ins](#)
- [Restores using cluster-aware plug-ins](#)

About client cluster support

NetVault offers cluster-aware versions of various plug-ins that enable data protection for distributed data. These plug-ins require Cluster Support Licenses.

The cluster nodes are grouped into a Virtual Client on which the cluster-aware plug-in is installed. The backups and restores of cluster nodes are performed through the virtual client.

The following table lists the NetVault plug-ins that can be used in a cluster setup.

Table 111. Cluster-aware plug-ins

Plug-in	Description
Quest NetVault Plug-in <i>for FileSystem</i>	<p>This plug-in is shipped with the NetVault software and can be used to back up the shared file system data on the following platforms:</p> <ul style="list-style-type: none"> • Windows Server Clusters • Linux Clusters • Sun Clusters (Solaris SPARC) <p>For more information about the supported cluster software versions, see the <i>Quest NetVault Compatibility Guide</i>. You can download this guide from https://support.quest.com/technical-documents.</p> <p>A default installation of NetVault does not require licensing of its native Plug-in <i>for FileSystem</i>. However, to use this plug-in a cluster setup, a File System Cluster Support license is required.</p>
Quest NetVault Plug-in <i>for Exchange</i>	<p>This plug-in can be deployed in an Exchange Server Single Copy Cluster (SCC)/Failover Cluster or Cluster Continuous Replication (CCR) setup to back up the distributed Exchange Server data. For more information, see the <i>Quest NetVault Plug-in for Exchange User's Guide</i>.</p>
Quest NetVault Plug-in <i>for Hyper-V</i>	<p>This plug-in can be deployed in a Hyper-V failover cluster setup to back up cluster data through a virtual client. For more information, see the <i>Quest NetVault Plug-in for Hyper-V User's Guide</i>.</p>

Table 111. Cluster-aware plug-ins

Plug-in	Description
Quest NetVault Plug-in <i>for Oracle</i>	This plug-in can be used in Oracle's Real Application Clusters (RAC) setup to back up the distributed Oracle database. For more information, see the <i>Quest NetVault Plug-in for Oracle User's Guide</i> .
Quest NetVault Plug-in <i>for SQL Server</i>	This plug-in can be used in a SQL Server Failover Cluster setup to back up the distributed SQL Server database. For more information, see the <i>Quest NetVault Plug-in for SQL Server User's Guide</i> .
Quest NetVault Plug-in <i>for MySQL</i>	This plug-in can be used in a MySQL Server Failover Cluster setup on Linux platform to back up the distributed MySQL Server database. For more information, see the <i>Quest NetVault Plug-in for MySQL User's Guide</i> .

Virtual clients

A virtual client is created when you install a cluster-aware plug-in. All nodes in a cluster are grouped to form a virtual client.

A virtual client is managed like any other NetVault Client. It can be browsed and added to client groups and policies, granted user access, and included in reports. The NetVault Server administers the creation and configuration of a virtual client. The cluster-aware version of the plug-in runs locally on the cluster nodes and the data is processed locally. A cluster node configured as a SmartClient sends data directly to the locally attached storage device.

Device configuration in cluster environment

In a cluster setup, a backup device can be connected in different ways. This section describes the pros and cons of some of the device configuration methods.

- **Connecting a device to the NetVault Server or Client:** This type of configuration allows the control of a robotic arm. However, during backups and restores, the data is transferred over the network.
- **Sharing drives:** A derivative of the previous method can be used by connecting the physical library to the NetVault Server, thus, giving it the control of the robotic arm, and sharing the drive with the cluster nodes. This configuration allows the control of the robotic arm and at the same time enables local data transfers.

This configuration offers high drive availability. With all cluster nodes sharing the control of drives, the drives are always available. However, the cluster node that currently controls the drive does not need to be the node that currently controls the cluster.

- **Connecting a device to a cluster node:** This configuration offers the fastest method of data transfer as the data is routed directly to a locally attached device.

However, the disadvantage is that the robotic arm cannot be controlled by a machine within the cluster, limiting the device type usage for this type of configuration to standalone drives. Moreover, the drive becomes unavailable when the cluster node is down.

Installing and upgrading cluster-aware plug-ins

This section includes the following topics:

- [Prerequisites](#)
- [Installing a cluster-aware plug-in](#)
- [Upgrading a cluster-aware plug-in](#)

Prerequisites

Before you start the installation procedure for a cluster-aware plug-in, verify that the following requirements are met:

- **Install NetVault Server:** Install the NetVault Server software on the designated machine. The server must be a separate machine outside the cluster setup. For instructions on installing the server software, see the *Quest NetVault Installation Guide*.
- **Install the NetVault Client:** Install the NetVault Client software on the individual cluster nodes. For instructions on installing the client software, see the *Quest NetVault Installation Guide*.
- **Add NetVault Clients:** Add the clients to the NetVault Server. For information about adding a client, see [Adding clients](#).
- **Copy the installation file:** Copy the “.npk” installation file for the cluster-aware plug-in to the NetVault Server.

The path to copy the file is <NetVault home>\packages\standard on Windows and <NetVault home>/packages/standard on Linux. You can also copy the installation files to subdirectories in the **standard** directory.

The installation file for the Plug-in for FileSystem (for the Server operating system) is already available in the **standard** directory.

Installing a cluster-aware plug-in

In a cluster setup, the plug-in is installed by creating a virtual client on the NetVault Server. All nodes in a cluster are grouped to form a virtual client. During the process, the plug-in is installed on the specified cluster nodes.

To install a cluster-aware plug-in:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 On the **Manage Clients** page, and click **Add Virtual Client**.
- 3 On the **Virtual Client** page, provide the following information:
 - In **Virtual Client Name**, type a name for the virtual client. The name must be unique. Spaces are not recognized in a virtual client name and are replaced with an underscore (“_”) character. The virtual client name cannot be changed once it is configured.
 - In **Virtual Client Address**, type the IP address of the cluster application.
 - In the Package list, select the installation file for the plug-in. This list includes all cluster-aware “.npk” files copied to the packages directory and its subdirectories. The following table provides the filenames for the installation files. Here **x-x- x-x** represents the version, build, and platform numbers.

Table 112. Installation files for cluster-aware plug-ins

Option	Description
Plug-in for <i>FileSystem</i> (for Windows)	win-x-x-x-x.npk
Plug-in for <i>FileSystem</i> (for Linux)	nvf-x-x-x-x.npk
Plug-in for <i>FileSystem</i> (for Solaris (SPARC))	nvf-x-x-x-x.npk
Plug-in for <i>Exchange</i>	exs-x-x-x-x.npk
Plug-in for <i>Hyper-V</i>	hv-x-x-x-x.npk
Plug-in for <i>Oracle</i>	ora-x-x-x-x.npk
Plug-in for <i>SQL Server</i>	sql-x-x-x-x.npk
Plug-in for <i>MySQL</i>	mys-x-x-x-x.npk

- To add a cluster node to the virtual client, select the client node in the **Available Clients** table, and click the Add button to the left of the item. The selected client is moved to the **Chosen Clients** table.
- To remove a cluster node from the virtual client, select the client node in the **Chosen Clients** table, and click the Remove button to the left of the item. The selected client is moved to the **Available Clients** table.

4 Click **Create Virtual Client**.

The NetVault Server starts installing the plug-in on the selected cluster nodes. This process overwrites the standard version of the same plug-in installed on the cluster nodes. However, you can use the cluster-aware version to perform backups of the local non-shared data. When the installation completes, the virtual client is added to the NetVault Clients table on the Manage Clients page.

Upgrading a cluster-aware plug-in

To upgrade a cluster-aware plug-in, the existing virtual client is removed and a new virtual client is created using the latest version of the plug-in.

To upgrade a cluster-aware plug-in:

- 1 Remove the virtual client that was created with the previous version of the plug-in. For more information about removing a virtual client, see [Removing a virtual client](#).
- 2 Create a new virtual client using the upgraded version of the cluster-aware plug-in. For more information about installing the plug-in, see [Installing and upgrading cluster-aware plug-ins](#).

You must assign the old virtual client's name to the new virtual client. If you change the name for the new virtual client, you cannot run the jobs that were defined for the old virtual client.

Configuring a cluster-aware plug-in

The configuration procedure for a cluster-aware plug-in includes the steps outlined in the following sections:

- [Configuring preferred network address for cluster nodes](#)
- [Configuring default settings for a cluster-aware plug-in](#)

Configuring preferred network address for cluster nodes

A cluster node has at least two network addresses:

- **Public IP Address:** The address that is used by machines outside the cluster to communicate with cluster nodes.
- **Private IP Address:** The address that is used by a cluster node to communicate with other machines within the cluster.

For each cluster node, you must configure the machine's public IP address as the "Preferred Network Address" for the node.

To configure preferred address for cluster nodes:

- 1 Obtain the IP address for the cluster node.
You can use the **ifconfig** utility on Linux and UNIX and the **ipconfig** utility on Windows to complete this step.
- 2 In the Navigation pane, click **Change Settings**.
- 3 On the **NetVault Server Settings** page, click **Choose machine**.
- 4 In the **Choose machine** drawer, select the target client, and click **Go to settings**.
- 5 On the **Client Settings** page, under **Services**, click **Network Manager**.
- 6 In the **Network Manager** dialog box, specify the preferred network address for the cluster node in the corresponding box.
- 7 Click **Apply** to apply the new settings and close the dialog box.
- 8 Repeat [Step 1](#) through [Step 7](#) for each cluster node.

Configuring default settings for a cluster-aware plug-in

To configure default settings for a cluster-aware plug-in:

- 1 In the Navigation pane, click **Create Backup Job**, and then click the **Create New** button next to the Selections list.
- 2 On the **NetVault Selections** page, double-click the virtual client node. In the list of plug-ins, select the cluster-aware plug-in, and in the **Actions** list, click **Configure**.

i **NOTE:** The default settings for a cluster-aware plug-in can only be set from the **NetVault Selections** page. For cluster-aware plug-ins, configuring these options from the **Change Settings** page is not supported.
- 3 In the **Configure** dialog box, set the required options. The configuration options for the cluster-aware version are the same as the standard version of the plug-in. For more information about these options, see the relevant plug-in user's guide.
- 4 Click **OK** to save the settings.

These settings are stored on the NetVault Server in configuration files specific to the virtual client, and applied during backups and restores of shared data performed through the virtual client.

Managing virtual clients

This section includes the following topics:

- [Modifying a virtual client](#)
- [Checking access to a virtual client](#)
- [Determining the current real client](#)
- [Removing a virtual client](#)

Modifying a virtual client

Once a virtual client is created, you can add or remove the cluster nodes or change the IP address for the cluster application.

To modify a virtual client:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the list of NetVault Clients, select the target virtual client, and click **Manage**.
- 3 On the **Virtual Client** page, modify the required settings:
 - To modify the IP address for the cluster application, type the new address in **Virtual Client Address**.
 - To add a cluster node to the virtual client, select it in the **Available Clients** table, and click the Add button to the left of the item. The selected client is moved to the **Chosen Clients** table.
 - To remove a cluster node from the virtual client, select it in the **Chosen Clients** table, and click the Remove button to the left of the item. The selected client is moved to the **Available Clients** table.
- 4 To save the modified settings, click **Save Virtual Client**.

Checking access to a virtual client

For a backup or restore job to complete successfully for a virtual client, at least one member client must be online and active. You can use the **Check Access** option to find out the accessibility status of the member clients.

To check the status of a virtual client:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the list of NetVault Clients, select the target virtual client, and click **Manage**.
- 3 On the **Virtual Client** page, click **Check Access**.

NetVault tries to connect to each member of the virtual client, and returns a message indicating the current accessibility status of the member clients.

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

Determining the current real client

You can use the **Current Real Client** option to find out which machine is in control of the cluster application.

To determine the machine that is currently in control of the cluster application:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the list of NetVault Clients, select the target virtual client, and click **Manage**.
- 3 On the **Virtual Client** page, click **Current Real Client**.

In the dialog that appears, the NetVault name of the controlling node is displayed.

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

Removing a virtual client

You can remove a virtual client if you no longer want to use the associated plug-in in a cluster setup. When you remove a virtual client, the cluster nodes added as NetVault Clients are not removed from the server. However, it removes the ability of the plug-in to back up the cluster data.

To remove a virtual client:

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the list of NetVault Clients, select the target virtual client, and click **Manage**.
- 3 On the **Virtual Client** page, click **Remove**.
- 4 In the **Confirm** dialog box, click **Remove**.

Backups using cluster-aware plug-ins

The procedure for performing backups using the Plug-in *for FileSystem* is similar for both standard and virtual clients. The cluster-aware version of the plug-in can be used in the same manner as the standard version to select the data items. However, when you open the virtual client node on the **NetVault Selections** page, only the cluster-aware plug-in is listed under the node. The plug-in displays both shared and local drives and mount points in the selection tree. You must make a note of the drive letter or mount point for the shared resource and select the data accordingly. The backup options that can be set for the standard version are also available with the cluster-aware version.

For information about backing up data in the Exchange Server Clustered environment, Oracle RAC setup, SQL Server Failover Cluster, and MySQL Server Failover Cluster, see the relevant plug-in user's guide.

Note the following:

- When you back up a virtual client, the data is backed up from a single client and it is accessed from the controlling node.
- When backing up cluster data using a virtual client, only the LUNs owned by the active node are backed up; the LUNs owned by passive nodes are not backed up.
- In an Active/Active cluster setup, the NetVault Server may start a backup on a secondary node even if the "primary only" option is selected for the backup. In such cases, the backup is redirected to the primary node, but the server only communicates with the secondary node to which it sent the message to start the backup. If the secondary node becomes unavailable while the backup is in progress, the server loses communication with the parent process running on this client. This process is not restarted even if the secondary node becomes available immediately. Consequently, the backup job becomes unresponsive.

i | **NOTE:** For cluster backups and restores, the virtual client name is displayed on the **Job Status** page and the actual client name is displayed on the **View Logs** page.

Cluster failover during backups

If a failover occurs during a backup, the job is aborted and the status “Job Failed” is returned. You can use the **Job Retry** scheduling option to run the job again after the failover completes.

With the Plug-in *for FileSystem*, when a failover occurs on Windows, the reason for the failover has a direct bearing upon the status of the failed job. Depending on the reason, the job returns the following exit status:

- If the machine in control of the cluster resources goes offline during a backup, the job is aborted and the status “Job Failed” is returned. You can use the job retry feature to run the backup again.
- If the machine in control of the cluster remains online but the actual cluster resource that is being backed up fails, the job is aborted and the status “Backup Completed with Warnings” is returned. The scheduled job retries do not work for such jobs. In this scenario, examine the job logs to find the missing data or run the job again to back up the data.

Restores using cluster-aware plug-ins

The procedure for performing restores using the Plug-in *for FileSystem* is similar for both standard and virtual clients. The backups are restored from the virtual client node and not the actual client node. When you submit a restore job, the plug-in communicates with the cluster service to determine the controlling node and targets this machine for the restore.

For information about restoring data in the Exchange Server Clustered environment, Oracle RAC setup, SQL Server Failover Cluster, and MySQL Server Failover Cluster, see the relevant plug-in user’s guide.

- i** | **NOTE:** For cluster backups and restores, the virtual client name is displayed on the **Job Status** page and the actual client name is displayed on the **View Logs** page.

Configuring default settings for NetVault

- [About configuring default settings](#)
- [Configuring encryption settings](#)
- [Configuring plug-in options](#)
- [Configuring default settings for post-scripts](#)
- [Configuring default settings for Verify Plug-in](#)
- [Configuring Deployment Manager Settings](#)
- [Configuring Job Manager settings](#)
- [Configuring Logging Daemon settings](#)
- [Configuring Media Manager settings](#)
- [Configuring Network Manager settings](#)
- [Configuring Process Manager settings](#)
- [Configuring RAS device settings](#)
- [Configuring Schedule Manager settings](#)
- [Configuring Web Service settings](#)
- [Configuring Auditor Daemon settings](#)
- [Configuring firewall settings](#)
- [Configuring general settings](#)
- [Configuring security settings](#)
- [Synchronizing NetVault Time](#)
- [Configuring default settings for global notification methods](#)
- [Configuring the reporting utility](#)
- [Configuring NetVault WebUI default settings](#)
- [Configuring NetVault to use a specific VSS provider](#)
- [Configuring default settings using Txtconfig](#)

About configuring default settings

NetVault runs with some default settings that can be customized to suit your environment. You can view and modify these settings from the **Change Settings** link in the Navigation pane. The default settings are available for the following services and components.

Figure 34. Server Settings page

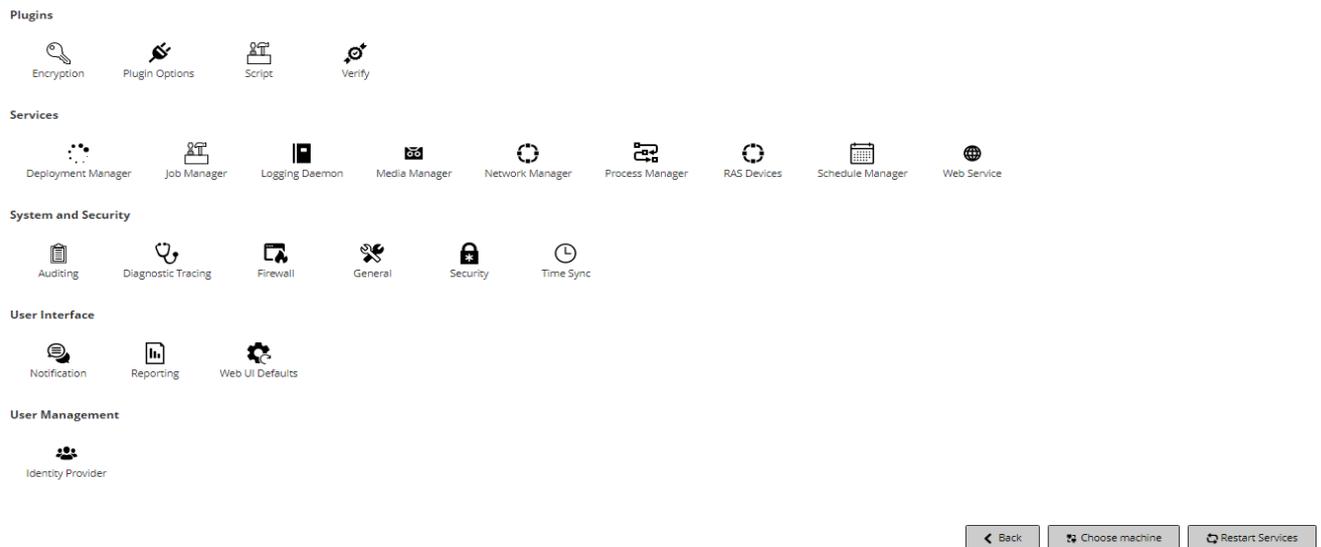


Table 113. Server settings

Group	Setting
Plugins	<ul style="list-style-type: none"> • Encryption • Plugin Options • Script • Verify
Services	<ul style="list-style-type: none"> • Deployment Manager • Job Manager • Logging Daemon • Media Manager • Network Manager • Process Manager • RAS Devices • Schedule Manager • Web Service
System and Security	<ul style="list-style-type: none"> • Auditing • Diagnostic Tracing • Firewall • General • Security • Time Sync
User Interface	<ul style="list-style-type: none"> • Notification • Reporting • WebUI Defaults
User Management	<ul style="list-style-type: none"> • Identity Provider

You can view and modify the NetVault Client settings by clicking **Choose Machine** on the Server Settings page or from the **Manage Clients** link in the **Navigation** pane.

Other configuration utilities

You can also use the Txtconfig utility to customize default settings for the server and client machines.

Txtconfig

The Txtconfig utility provides a text-based interface to customize various settings for the server and client machines. This utility is available on all supported operating systems. The Txtconfig utility resides in the **bin** directory under the NetVault installation directory. You must be logged-in with Administrator privileges on Windows and root user privileges on Linux and UNIX to use the Txtconfig utility. For more information, see [Configuring default settings using Txtconfig](#).

Configuring encryption settings

Before you can start using the Plug-in for Standard Encryption or Plug-in for Advanced Encryption, you must configure the encryption settings, which specify the encryption algorithm and encryption key that you want to use. You can configure these settings from the **Change Settings** link.

The encryption plug-ins provide support for CAST-128, AES-256, and CAST-256 algorithms to meet regulatory backup security requirements. You can install these plug-ins on the NetVault Server or Clients to perform encrypted backups for that machine. For more information about these plug-ins, see the *Quest NetVault Built-in Plug-ins User's Guide*.

To configure default settings for the encryption plug-ins:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, go to Step 3.
 - To configure default settings for a NetVault Client, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**.
- 3 Under **Plugins**, click **Encryption**, and configure the following settings.

Table 114. Default settings for encryption plug-ins

Setting	Description
Encrypt ALL Backups on this Client	<p>Once the Plug-in for Standard Encryption or Plug-in for Advanced Encryption is installed on a client, you can do either of the following:</p> <ul style="list-style-type: none">• Encrypt all backups performed for that client• Encrypt specific backups performed for that client <p>To enable encryption for all backups, select this check box. When you enable encryption for all backups, you cannot change the setting on a per-job basis.</p> <p>For more information about performing job-level encryption, see Specifying additional options.</p> <p>NOTE: To perform job-level encryption for backups originating from a NetVault Server or Client, the plug-in should not be configured for encrypting all backups.</p>
Encryption Key String	<p>Type the string that serves as the encryption key for the NetVault machine.</p> <p>Different platforms allow varying characters and password lengths. We recommend that you use passwords of 32 characters or less. You can use characters from the following set: "A-Z", "a-z", "0-9", and "_". Key strings that do not conform to these specifications may work on one platform, but may be invalid in another environment.</p>
Available Encryption Algorithms	<p>Select the encryption algorithm that you want to use for backups and restores. Depending on the products that you have installed, the list includes the following options: CAST-128, CAST-256, and AES-256.</p>

- 4 Click **Apply** to apply the new settings and close the dialog box.

i | **NOTE:** An encrypted backup can be restored to the original client or an alternate client. In both cases, the plug-in must be installed on the target client and it must be configured as it was when the backup was performed — using the same **Encryption Key String** and **Encryption Algorithm**.

Configuring plug-in options

You can configure default settings for various built-in and licensed plug-ins from the **Change Settings** or **Create Backup Job** page. The plug-in configuration procedures are covered in the respective user guides.

This section includes the following topics:

- [Configuring default settings for Disk Devices Plug-in](#)
- [Configuring default settings for nvjobstart](#)

Configuring default settings for Disk Devices Plug-in

The Disk Devices Plug-in is used to create Virtual Tape Libraries (VTLs). You can configure default settings for this plug-in from the **Change Settings** page.

To configure default settings for the Disk Devices Plug-in:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, go to Step 3.
 - To configure default settings for a NetVault Client, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**.
- 3 Under **Plugins**, click **Plugin Options**.
- 4 In the **Plugin Options** dialog box, under **Disk Devices Plugin**, configure the following options.

Table 115. Default settings for Disk Devices Plug-in

Setting	Description
Allow disk libraries to have entry/exit ports	Virtual tape libraries do not have an entry/exit port by default. To create Virtual Tape libraries that have entry/exit ports (EEPorts), select this check box.
Check available disk space before creating disk libraries	This check box is selected by default. Before creating a VTL, NetVault performs a disk space check to ensure that the target disk has sufficient space to accommodate the new VTL. On normal file systems you can use this feature to avoid errors during VTL creation. When creating a VTL on a third-party deduplication appliance or compressed file system, you should disable this feature. To disable disk space checks for VTLs, clear this check box.

Table 115. Default settings for Disk Devices Plug-in

Setting	Description
Free disk space margin to be used when calculating available disk space	<p>During disk space checks, the free space required on the disk is calculated as follows:</p> $\text{Number of Slots} * \text{Media Capacity} + \langle x \rangle$ <p>$\langle x \rangle$ is the additional disk space considered for the following requirements:</p> <ul style="list-style-type: none"> • Disk space required to create the directory structure for the VTL. This requirement varies for different file systems. • Disk space required by other applications running on the system. <p>The default value is 20MB. To change this requirement, type or select the new value.</p>
Unit used to express the free disk space margin in	Type or select the unit used to specify disk space margin. The unit can be MB or GB. The default unit is MB.

- 5 Click **Apply** to apply the new settings and close the dialog box.

Configuring default settings for nvjobstart

By default, the **nvjobstart** command returns either 0 or 1 (0 to indicate success, and 1 to indicate failure). You can configure the utility to return enhanced job completion status codes and messages that indicate the way in which a job failed.

To configure default settings for the nvjobstart utility:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, go to Step 3.
 - To configure default settings for a NetVault Client, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**.
- 3 Under **Plugins**, click **Plugin Options**.
- 4 In the **Plugin Options** dialog box, under **CLI**, configure the following setting:
 - **nvjobstart Enhanced Job Completion Status:** By default, the **nvjobstart** utility returns the following messages depending on the job exit codes.

Table 116. Default job exit codes and status messages

Exit code	Status message
0	Job completed successfully
1	<p>Job failed with error: 'Job Failed'</p> <p>This message is returned for the following job completion states:</p> <ul style="list-style-type: none"> • Job failed • Job completed with warnings • Job aborted • Job stopped • Job died

If you select the **nvjobstart Enhanced Job Completion Status** check box, the utility returns the following codes and messages.

Table 117. Enhanced job exit codes and status messages

Exit code	Status message
0	Job completed successfully
1	Job Failed
2	Job Completed with Warnings
3	Job Aborted
4	Job Stopped
5	Job Died
-1	Job Failed with Undefined Error

NOTE: You must configure this option on the NetVault machine on which you run the **nvjobstart** utility.

- 5 Click **Apply** to apply the new settings and close the dialog box.

Configuring default settings for post-scripts

By default, the execution of a post-script is terminated when a backup or restore job is aborted. You can modify the configuration setting for post-scripts to change this behavior.

To change the default setting for post-scripts:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, go to Step 3.
 - To configure default settings for a NetVault Client, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**.
- 3 Under **Plugins**, click **Script**.
- 4 In the **Script** dialog box, configure the following setting:
 - **Terminate running script on Job Abort:** This check box is selected by default. If you want to continue script execution even when the associated job is aborted, clear this check box.

This setting applies to all post-scripts defined for the backup, restore, and report jobs performed on the given NetVault Client.
- 5 Click **Apply** to apply the new settings and close the dialog box.

Configuring default settings for Verify Plug-in

With a backup job, you can choose to run the verification phase to check the correctness and completeness of a backup at the end of data transfer. NetVault uses the built-in Verify plug-in to perform the verification phase. The Verify plug-in verifies the stream length written to the media and ensures that no blocks were dropped during backup. While the actual backup runs as Phase 1, backup verification runs as Phase 2 of the backup job.

By default, the backup verification job runs on the NetVault Server. You can modify the configuration settings for the plug-in to run the verification on an alternate client or locally on specified clients.

To configure default settings for the Verify Plug-in:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Plugins**, click **Verify**. In the **Verify** dialog box, configure the following settings.

Table 118. Default settings for Verify plug-in

Setting	Description
Always run Verify locally	By default, the verification phase runs on the NetVault Server. To avoid data transfers over the network, you can choose to run verification locally on the clients to which the device used for backup is attached. This option is globally applied to all clients. It does not work for clients that do not have a locally attached backup device.
Comma separated list of clients that verify locally	If your backup jobs are distributed across multiple client-attached devices, use this setting to specify a comma-separated list of clients that can run the verification phase locally.
Default client to run Verify	To configure an alternate client to run all verification jobs, specify the client name. This setting is useful if you use a client-attached device for your backups. It allows you to select the same client to run backup verification jobs.

NOTE: You must configure these settings on the NetVault Server.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring Deployment Manager Settings

This section includes the following topics:

- [About Deployment Manager](#)
- [Configuring default settings for Deployment Manager](#)

About Deployment Manager

The Deployment Manager process (**nvdeploymentmgr**) runs on the NetVault Server. This process manages the push installation tasks.

You can use the push installation method to perform the following tasks:

- Install and upgrade the NetVault Client software simultaneously on multiple machines
- Install and upgrade NetVault plug-ins simultaneously on multiple machines
- Add new clients to the NetVault Server

For more information about the push installation method, see [Installing software from WebUI \(push installation\)](#).

You can modify the Deployment Manager settings from the **Change Settings** page.

Configuring default settings for Deployment Manager

To modify default settings for Deployment Manager:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Deployment Manager**. In the **Deployment Manager** dialog box, configure the following settings.

Table 119. Deployment Manager settings

Setting	Description
Maximum simultaneously active deployment processes on server	By default, 50 concurrent deployment processes can run on the server. To change the number of deployments that can be run in parallel at any time, type or select the new value.
Maximum time a deployment task should wait when processing a client	This setting determines the amount of time the Deployment Manager waits for a response from a remote client during the client addition phase. The default value is 300 seconds. The timeout interval can be between 30 and 900 seconds.
Skip Remote Machine Cleanup	This setting lets users change the behavior of client-side log cleanup. By default, it skips cleanup on the remove target machine.

- 3 Click **Apply** to apply the new setting and close the dialog box.

Configuring Job Manager settings

This section includes the following topics:

- [About Job Manager](#)
- [Configuring default settings for Job Manager](#)

About Job Manager

The Job Manager process (**nvjobmgr**) runs on the NetVault Server and manages the execution of a job.

The Schedule Manager starts the Job Manager. A single instance of Job Manager runs for each job until the completion of the job. The Job Manager reports on the job run status and exit status. This process coordinates with the Data Plug-in and fetches the required information from the NetVault Server. It is also responsible for sending drive and media requests to the Media Manager process.

You can modify the Job Manager settings from the **Change Settings** page.

Configuring default settings for Job Manager

To configure default settings for Job Manager:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Job Manager**. In the **Job Manager** dialog box, configure the following settings.

Table 120. Job Manager settings

Setting	Description
Job Keep Alive rate	Keep-alive messages are used to verify network connection between two NetVault machines and to keep this connection intact. The keep-alive rate setting controls how often keep-alive messages are sent between the Job Manager running on the server and the Data Plug-in running on a client. The default value is 5 minutes. To change the interval, type or select a new value. The keep-alive rate is specified in number of minutes.
CLI utility nvjobcreate will return exit status	By default, the nvjobcreate utility returns the Job ID on success and "0" on failure. When you select this check box, the nvjobcreate utility returns the exit code "0" on success and "1" on failure. To determine the Job ID number when the nvjobcreate utility is configured to return the exit status, you can use the -jobidfile option. The Job ID number is written to the specified file. For more information about this option, see the <i>Quest NetVault CLI Reference Guide</i> .
Allow VMware plugin jobs to be distributed on other VMware backup proxy	To distribute the backup jobs running on a Plug-in for VMware to the other VMware backup proxies, select this check box. By default, this check box is clear.
Jobs threshold for VMware Backup Proxy	If you want to change the default threshold from 2, enter the applicable number for the jobs threshold for VMware Backup Proxy setting. This value indicates the number of VMware proxy jobs that can run on the same VMware proxy before the next job is distributed to the same proxy or a different one, depending on the load balancing.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring Logging Daemon settings

This section includes the following topics:

- [About Logging Daemon](#)
- [Configuring disk space alert thresholds](#)
- [Modifying the purge policy for log messages](#)
- [Configuring additional settings for Logging Daemon](#)

About Logging Daemon

The Logging Daemon (**nvlogdaemon**) runs on the NetVault Server. This process starts along with the NetVault Service.

The Logging Daemon manages the log messages generated by various NetVault processes and writes these messages to the NetVault Database. Log messages contain information that can be used for diagnosing and troubleshooting problems. The Logging Daemon also performs periodic disk space checks, and issues alert messages when the space usage level reaches the Warning or Critical Threshold. These checks apply to the NetVault Home, Database, Logs, and Reports directories.

You can modify the Logging Daemon settings from the **Change Settings** page.

Configuring disk space alert thresholds

The Logging Daemon performs periodic disk space checks, and issues alert messages when the space usage level reaches the Warning or Critical Threshold. These checks are performed for the following directories: NetVault Home, Database, Logs, and Reports directories. The Warning and Critical thresholds are set to 85 and 95 percent of the total disk space, respectively. The default interval between two disk space check events is one hour.

You can use the configuration settings for Logging Daemon to change the alert thresholds and default interval for disk space checks.

i | **NOTE:** You must restart the NetVault Service to apply any changes to the Warning and Critical Threshold values.

To change the default alert threshold settings:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Logging Daemon**. In the **Logging Daemon** dialog box, configure the following settings.

Table 121. Disk space alert threshold settings

Setting	Description
Disk Space Warning Threshold	Type or select the Warning Threshold level. The default value is 85 percent of the total disk space. When the disk space usage reaches or exceeds the Warning Threshold, the following events occur: <ul style="list-style-type: none">• The Logging Daemon reports an error.• The Auditor Daemon logs a message.• NetVault raises the Home Drive Becoming Full event. You can set up a global notification method to receive notification when this event is raised. For more information, see Using global notification methods.
Disk Space Critical Threshold	Type or select the Critical Threshold level. The default value is 95 percent of the total disk space. When the disk space usage reaches or exceeds the Critical Threshold, the following events occur: <ul style="list-style-type: none">• NetVault sends an error message to the Event Viewer on Window and syslog on Linux and UNIX.• The Logging Daemon reports an error. NOTE: If the disk space is critically low when the NetVault Service starts, the errors are only logged through the Event Viewer or syslog. <ul style="list-style-type: none">• The NetVault Service shuts down automatically and the service status is set to “Stopped Disk Full.” You can restart the service only when the disk usage percentage drops below the configured Critical Threshold level.
Time interval between disk space full checks	Type or select the interval between two disk space check events. It is specified in number of hours. The default value is one hour. To disable disk space checks, set the value to zero (0). NOTE: Regardless of this setting, the Logging Daemon performs disk space check when the NetVault Service is restarted.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Modifying the purge policy for log messages

Log messages that are older than 30 days are automatically purged from the database. You can use the configuration settings for Logging Daemon to change the maximum age for log messages.

i | **NOTE:** You can also manually delete the log messages by using the **Purge Logs** method available on the **View Logs** page. For more information, see [Manually purging the log messages](#).

To modify the purge policy for logs:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Logging Daemon**. In the **Logging Daemon** dialog box, configure the following settings.

Table 122. Auto-purge settings for logs

Setting	Description
Auto-purge method	By default, the Auto-purge method is set to “ Exceeds log age .” Use this setting to delete logs that are older than a specified number of days. To disable automatic purging of log files, set the Auto-purge method to “ None .” NOTE: Log messages can consume a considerable amount of disk space. Therefore, periodic purging of the log messages is necessary. If automatic purging is disabled, use the Purge method available on the View Logs page to manually delete the log messages at regular intervals.
Auto-purge entries that are older than	When the Auto-purge method is set to “ Exceeds log age ,” use this setting to specify the maximum age for logs. The log age is specified in number of days. The default value is 30 days.
Select the time interval to auto-purge	The default interval between two auto-purge events is three hours. To change the interval, type or select a new value. The purge interval is specified in number of hours.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring additional settings for Logging Daemon

To configure additional settings for the Logging Daemon:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Logging Daemon**. In the **Logging Daemon** dialog box, configure the following settings.

Table 123. Additional settings for Logging Daemon

Setting	Description
Message Response Interval	<p>This setting determines the amount of time the NetVault WebUI waits for the Logging Daemon to send all log messages before displaying a progress dialog box. The dialog box displays the number of log messages loaded against the total number of log messages. If the response is delayed, you can cancel the operation by clicking Cancel on the progress dialog box.</p> <p>The default response time for the Logging Daemon is 256 milliseconds. To change default setting, type or select a new value. The minimum value can be 100 milliseconds.</p> <p>If the NetVault WebUI receives all log messages within the specified interval, the progress dialog box is not displayed (for example, if there are only a few log messages, the dialog box is not displayed).</p>
Outgoing message bundle size	<p>To improve performance, the logging daemon sends messages in bundles. Each bundle contains 64 messages by default. To change the number of messages per bundle, type or select the bundle size in Outgoing Message Bundle Size.</p>
Minimum time between progress updates	<p>By default, the progress dialog box is refreshed every 256 milliseconds. To change the refresh rate, type or select the time interval.</p>
Minimum warning level to send to system log	<p>By default, NetVaults sends log messages with warning code 64 and above to the OS. You can view these messages from the Event Viewer (on Windows) or syslog (on Linux and UNIX).</p> <p>To change the severity level of messages that NetVault sends to the OS, specify the warning level code in the Minimum warning level to send to system log box. The following list provides the warning codes and the corresponding severity level of the messages:</p> <ul style="list-style-type: none"> • 112: Severe errors • 96: Severe errors and errors • 80: Severe errors, errors and warnings • 64: Severe errors, errors, warnings, and startup messages • 48: Severe errors, errors, warnings, startup messages, and job messages • 32: Severe errors, errors, warnings, startup messages, job messages, and informational messages • 16: Severe errors, errors, warnings, startup messages, job messages, informational messages, and background messages • 0: All messages

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring Media Manager settings

This section includes the following topics:

- [About Media Manager](#)
- [Configuring general settings for Media Manager](#)
- [Configuring retirement rules for dependent backups](#)
- [Configuring Media Manager settings for tape devices](#)
- [Configuring Media Manager settings for backup indexes](#)
- [Configuring Media Manager settings for RAS devices](#)

- Configuring transfer update settings for foreign RAS devices
- Configuring data transfer stall timeout period for NetVault SmartDisk
- Configuring media request weightings
- Configuring default interval for backup retirement scans
- Configuring an alternate index read block size for Quest DR Series systems

About Media Manager

The Media Manager process (**nvmedmgr**) runs on the NetVault Server.

This process manages the Media Database, which contains information about the media contents and online backup savesets. The Media Manager also stores and manages the device configuration details. It manages the backup devices through the Device Manager processes. The Media Manager issues high-level instructions for loading and unloading media; the Device Manager processes carry out these instructions. The Media Manager controls the selection of device and media for a job according to the media requests submitted by the Job Manager.

You can modify the Media Manager settings from the **Change Settings** page.

NOTE: You must restart the NetVault Service to apply any changes to the Media Manager settings.

Configuring general settings for Media Manager

To configure general settings for Media Manager:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Media Manager**. In the **Media Manager** dialog box, under **General Settings**, configure the following settings.

Table 124. General settings for Media Manager

Setting	Description
Minimum life for imported backups	<p>This setting specifies the minimum life for backups that are imported to the NetVault Database by scanning the backup media. This setting only applies to the backups that are not available in the NetVault Database.</p> <p>When you import a backup, its data is temporarily stored in the NetVault Database for seven days. To change the default setting, type or select a new value. The minimum life setting is specified in number of days.</p> <p>You can override the global setting by configuring the minimum life for imported backups at the time of scanning. For more information, see Scanning a disk-based storage device and Scanning tape media.</p> <p>This setting does not affect the backups that have not reached their retirement age or generation count. For such backups, the backup life setting determines the retirement time.</p>
Use barcodes as labels	<p>You can configure NetVault to automatically label blank media during backups by setting the Automatically Label Blank Media option in the Target Set.</p> <p>By default, NetVault assigns a system-generated string as the media label to blank media items. The string consists of the NetVault Server Name, the current date, and a seed number. To assign media barcodes as the default label for blank media during backups, select this check box.</p>

Table 124. General settings for Media Manager

Setting	Description
Continuation Restore Request Priority Boost	<p>A continuation restore request occurs when an active restore job requires additional media for completion.</p> <p>By default, these requests are assigned a priority level of 5 so that they are not preempted by other media requests and the active job is completed without any interruption. To change the default setting, type or select a value from 1 through 100. 1 denotes highest priority, while 100 denotes lowest priority. A job with a priority level of zero (0) runs as a background task.</p>
Continuation Backup Request Priority Boost	<p>A continuation backup request occurs when an active backup job requires additional media for completion.</p> <p>By default, these requests are assigned a priority level of 5 so that they are not preempted by other media requests and the active job is completed without any interruption. To change the default setting, type or select a value from 1 through 100. 1 denotes highest priority, while 100 denotes lowest priority. A job with a priority level of zero (0) runs as a background task.</p>
Label Request Priority	<p>NetVault assigns a priority level of 10 to bulk media labeling requests. To change the default priority setting, type or select a value from 1 through 100. 1 denotes highest priority, while 100 denotes lowest priority. A job with a priority level of zero runs as a background task.</p>
Blank Request Priority	<p>NetVault assigns a priority level of 10 to bulk media blanking requests. To change the default setting, type or select a value from 1 through 100. 1 denotes highest priority, while 100 denotes lowest priority. A job with a priority level of zero (0) runs as a background task.</p>
Bulk Label Callback Timeout	<p>The callback timeout for bulk media labeling controls how long NetVault waits for user inputs before ending these requests.</p> <p>By default, NetVault waits for 120 seconds for user confirmation. If confirmation is not received within the specified interval, the request is not carried out. To change the default setting, type or select a new value. The timeout value is specified in number of seconds.</p>
Bulk Blank Callback Timeout	<p>The callback timeout for bulk media blanking controls how long NetVault waits for user inputs before ending these requests.</p> <p>By default, NetVault waits for 120 seconds for user confirmation. If confirmation is not received within the specified interval, the request is not carried out. To change the default setting, type or select a new value. The timeout value is specified in number of seconds.</p>
Minimum interval between reporting online capacity	<p>This setting determines how frequently the online media status is updated on the NetVault WebUI. The default value is 10 minutes.</p> <p>To change the default setting, type or select a new value. The update interval is specified in number of minutes.</p>
Mark expired Read Only media as writable	<p>This option prevents automatic reuse of retired read-only media.</p> <p>When you select this check box, the read-only tag for a media item is automatically removed when the last saveset stored on it is expired, and the media item becomes available for re-use.</p> <p>Existing read-only and expired media (all saveset are expired) prior to selecting this check box, will be writable when the next saveset is expired in the NetVault after selecting this check box.</p>
Maximum characters in a media request diagnosis	<p>By default, NetVault supports a maximum of 64,000 characters in the Diagnose Media Requests for Job dialog box. To change this setting, type or select a new value. The value is specified in thousands of characters.</p> <p>Lesser number of characters may result in quicker output. However, the output may be truncated and you may not get the complete information. Therefore, you should not change the default value for this option.</p>

Table 124. General settings for Media Manager

Setting	Description
Use Target media request timeout for Source media request	<p>This option can be used to set timeout interval for source media requests. It applies to Plug-in <i>for Consolidation</i>, Plug-in <i>for Data Copy</i>, and Secondary Copy jobs.</p> <p>When you select this check box, the timeout setting for the source media is automatically obtained from the Target Set defined for the job. If the required piece of media is not available within the specified interval, the job is aborted.</p>
Force Process Media Request Queue Throttling	<p>This setting allows NetVault to interrupt the continuous processing of media request queues in very busy systems to allow other processing to occur. It determines the minimum response time between two media requests. The default value is 5 seconds. The value must be specified in number of seconds.</p> <p>NOTE: This setting should only be changed under the direction of Quest Technical Support.</p>
Life Of Segments Table Scan Results For Duplication	<p>When performing Phase 2 Data Copy and Duplication jobs, the Media Manager stores the results of the Segments table scan so that it does not have to scan the table each time a media request is serviced. This setting determines the amount of time the cached results are retained by NetVault.</p> <p>The default value is 10 seconds. You can change it to any value between 10 and 1800 seconds.</p>
Maximum number of sessions to retire at a time	<p>By default, the Media Manager removes a maximum of 250 sessions per iteration. To change the default setting, type or select a new value.</p>
Delete Unused Media Groups	<p>This check box is selected by default.</p> <p>When configuring the target media options for a backup job, the media group list only shows those groups that contain any media items. The Media Manager automatically removes unused media groups.</p> <p>If you do not want to remove unused groups from the media group list, clear this check box.</p>
Media Online Threshold (secs)	<p>This setting allows you to select the time interval (in seconds) at which the NetVault media manager updates the device status in the database.</p> <p>Type or select the value. The default value is 0 second. The range for this setting is 0-600 seconds.</p> <p>If the value of this setting is set to low, the device status is updated frequently resulting in overhead on media manager. If the number of media devices attached are few, then a lower value is preferred.</p> <p>However, if more number of media devices are added then, enter a high value for this parameter.</p> <p>NOTE: Increasing this value leads to delay in updating the device online status and does not affects the performance.</p>

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring retirement rules for dependent backups

By default, the retirement of backups from a backup series is deferred until all dependent backups are ready for retirement. The backups are retired only when the last dependent backup is retired. You can use the **Retirement Timing Control** setting to modify this behavior and force removal of backups according to their retirement schedule.

To configure the retirement rules dependent backups:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Media Manager**. In the **Media Manager** dialog box, under **General Settings**, configure the following setting:
 - **Retirement Timing Control:** You can use the following rules to retire dependent backups:
 - **Force Always:** When this rule is applied, a backup is removed when it reaches its retirement time. As a result, all backups from an Incremental Backup Series or a Differential Backup Series are retired when the first backup is retired. Forcing this behavior can cause early retirement of dependent Incremental and Differential Backups.
 - **Keep Duplicates with Dependents:** This rule specifies the retirement behavior of Duplicate Backups. When this rule is applied, Duplicates with dependents are retired only after all dependent backups have reached their retirement date. The Duplicate Backups are marked for removal according to the retirement schedule, but the removal is deferred until all dependent backups are ready for retirement.
 - **Allow Duplicates with Dependents to Retire:** This rule specifies the retirement behavior of Duplicate Backups. When this rule is applied, Duplicates with dependents are removed when their retirement date is reached only if the original backups or other copies are available for these backups. In the absence of the original backup or other copies, the Duplicate Backups are marked for deletion, and removed when the last dependent backup is retired.
- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring Media Manager settings for tape devices

To configure Media Manager settings for tape devices:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Media Manager**. In the **Media Manager** dialog box, under **General Settings**, configure the following settings.

Table 125. Device settings controlled by Media Manager

Setting	Description
Do not describe device activity in device window	Select this check box to stop status updates on the Device Activity page. It may help reduce the load on the NetVault Server when several backup devices are added to the server.
Do not display drives that are available but not active	Select this check box to hide the offline devices on the Device Activity page.
Only check available drives and media when processing a media request	Select this check box to only check for available drives and media when a media request is received. It can be useful in large environment to reduce the network traffic generated by automatic checks. However, it may cause a delay in job startup when several jobs are started at the same time.
Only check available drives when processing a media request	Select this check box to only check for available drives when a media request is received.
Only check available media when processing a media request	Select this check box to only check for available media when a media request is received.

Table 125. Device settings controlled by Media Manager

Setting	Description
Do not issue Prevent/Allow Media Remove commands to drives	During backup and restore operations, the Media Manager issues PREVENT/ ALLOW MEDIUM REMOVAL SCSI commands to move a tape to or from a drive. If your library software can handle tape removal or ejection by itself and does not require any explicit commands to be sent to the drive, you can select this check box. When this check box is selected, the Media Manager skips the PREVENT/ALLOW MEDIUM REMOVAL SCSI commands during device operations.
Do not open a device if the serial number has changed	Select this check box to stop issuing commands to a device whose serial number has changed since its last configuration for NetVault.
Unavailable device retry interval	Type or select the interval at which NetVault tries to locate unavailable devices. The retry interval is specified in minutes. The default value is 30 minutes.
Allow library modification when jobs are running	Select this check box to allow a user to modify the library settings while it is in use.
Do Not Scan Unknown Media	When you add tapes to a library, and close the door (or restart the library), NetVault takes inventory by reading the media barcodes. If NetVault cannot find the barcode for a piece of media, it marks that tape as "UNKNOWN." NetVault loads each unknown tape into one of the drives to scan for backups and on-tape indices. On large systems, this process can increase the burden on the drive resources. If you do not want to scan unknown tapes and leave them as "UNKNOWN," select this check box. This setting applies to all new libraries that are added to the NetVault Server. It does not change the setting for existing libraries. To apply this setting to an existing library, you must remove the library and re-add it.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring Media Manager settings for backup indexes

To configure Media Manager settings for backup indexes:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Media Manager**. In the **Media Manager** dialog box, under **General Settings**, configure the following settings.

Table 126. Index-related settings

Setting	Description
Days of inactivity before an index is compressed	The online indexes are automatically compressed after 30 days of inactivity. To change the default inactivity period for index compression, type or select the new value. To disable automatic compression of online indexes, set this option to zero (0).

Table 126. Index-related settings

Setting	Description
Hours between scans for indexes to be compressed	<p>This setting determines the interval at which the Media Manager scans the Media Database to identify backup indexes that can be compressed. The default value is 12 hours. To change the default interval, type or select a new value.</p> <p>To disable Media Manager scans for index compression, set this option to zero (0).</p>
Maximum number of indexes to be compressed/uncompressed/offlined at a time	<p>This setting determines the maximum number of backup indexes that can be simultaneously compressed, decompressed, or taken offline. The default value is 200 indexes. To change the default setting, type or select a new value.</p>
Minimum Space For Index Decompression	<p>The compressed backup indexes are automatically decompressed when you try to browse or restore the corresponding saveset. The minimum amount of space for index decompression is set to 500MB by default. This space is reserved on the drive or partition on which the NetVault Database resides. To change the minimum space for index decompression, type or select the new value. The space requirement is specified in MB.</p> <p>A decompression request fails in the following circumstances:</p> <ul style="list-style-type: none"> • The available space on the target drive or partition is less than the minimum required space. • The specified amount of free disk space is not available on the target drive or partition after index decompression. <p>When NetVault Database receives a decompression request, it reads the index header to determine its decompressed file size. The request is not completed if after decompression the free disk space on the target drive or partition would be less than the specified amount.</p> <p>If the decompression request fails, an error message is displayed. If the request was issued manually, the message is displayed on the Create Restore Job page. When index decompression is performed automatically, the message is displayed on the View Logs page.</p>
Days of inactivity before an index is offlined	<p>To automatically delete online indexes after a specified period of inactivity, type or select the value. The inactivity period is specified in number of days. The default value is zero (0), which disables automatic deletion of online indexes.</p>
Hours between scans for indexes to be offlined	<p>This setting determines the interval at which the Media Manager scans the Media Database to identify backup indexes that can be taken offline. The default interval is 24 hours. To change the default interval, type or select a new value.</p> <p>To disable Media Manager scans for online index deletion, set this option to zero (0)</p>
Life Of Index When Loaded From Offline	<p>The nvrestore CLI utility automatically imports an offline index if it is needed for a restore job.</p> <p>This setting determines how long the indexes imported by the nvrestore utility are retained in the NetVault Database. The default value is one (1) day. To change the default setting, type or select the number of days you want to retain the index. This value must be specified in number of days.</p>

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring Media Manager settings for RAS devices

To configure Media Manager settings for RAS device:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Media Manager**. In the **Media Manager** dialog box, under **General Settings**, configure the following settings.

Table 127. Media Manager settings for RAS devices

Setting	Description
Offline RAS device after timeout time	The timeout value for determining that a RAS device is not responding. The default value is 7200 seconds. To change the default setting, type or select a new value. After the timeout expires, NetVault sets the device status to offline.
Number of attempts (retries) for onlining a RAS device	The number of times NetVault tries to contact an offline RAS device in an attempt to bring it back online. The default value is 3. If you set this value to zero (0), no attempts are made to bring the device back online.
Cancel online RAS device batch after timeout time	The amount of time NetVault waits for a response from the RAS device before canceling the Media Manager batch that is trying to bring the device back online. The default value is 300 seconds.
RAS device start transfer batch timeout time	The amount of time NetVault waits for a RAS device to begin data transfers. The default value is 30 seconds. You can change it to any value between 10 and 300 seconds.
RAS segment deletion batch timeout time	<p>This setting determines the timeout interval for a batch request for deleting backup segments and indexes from disk-based storage devices. The default value is 300 seconds. If the task completion response is not received within this interval, the Media Manager cancels the current request, and sends a new request.</p> <p>When deleting large savesets, you can increase the timeout interval to allow the request to be completed before the timeout period elapses. The timeout interval can be between 300 and 3600 seconds.</p>
Max number of RAS segments can be deleted in one batch	By default, the Media Manager removes a maximum of 500 RAS segments per iteration. The value for this setting can be between 100 and 100000.
Max number of RAS Indexes can be deleted in one batch	By default, the Media Manager removes a maximum of 500 RAS segments per iteration. The value for this setting can be between 100 and 100000.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring transfer update settings for foreign RAS devices

NetVault writes the data transfer statistics for Data Domain Systems to the **stats.stnz** file. The **nvstatsmgr** process uses this file and requires that it is regularly updated. However, frequent updates can have a significant performance impact on the system.

By default, NetVault updates the **stats.stnz** file after every 5 seconds or 10 blocks of data transfer. You can change this default setting from the **Change Settings** page.

To configure transfer update settings for Data Domain Systems:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Media Manager**. In the **Media Manager** dialog box, under **Foreign RAS Device Configuration**, configure the following settings:
 - **Transfer Update Frequency (Blocks):** By default, NetVault updates the stats.stnz file after every 10 blocks of data transfer. To change the default setting, type or select a new value.
 - **Transfer Update Frequency (Seconds):** By default, NetVault updates the stats.stnz file after every 5 seconds. To change the default interval, type or select a new value. The time interval is specified in number of seconds.
- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring data transfer stall timeout period for NetVault SmartDisk

To configure data transfer stall timeout period for NetVault SmartDisk:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Media Manager**. In the **Media Manager** dialog box, under **DAV RAS Device Configuration**, configure the following settings:
 - **Data transfer stall timeout:** This setting determines the amount of time NetVault waits for a “stream end” response from a NetVault SmartDisk before reporting a stall. The default value is 1800 seconds. To change the default setting, type or select the new value. The stall timeout interval is specified in number of seconds.
- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring media request weightings

The Media Manager process uses media request weightings while assigning a media request for a backup or restore job.

i | IMPORTANT: These settings should only be changed under the direction of Quest Technical Support.

To change relative priority for backup devices:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Media Manager**. In the **Media Manager** dialog box, under **Media Request Weightings**, configure the following settings.

Table 128. Media request weightings

Setting	Description
Local Device Weighting	The media request weighting for local devices is set to 32 by default. To change this setting, type or select the new value. You can assign any value between 2 and 32.
RAS Device Weighting	The media request weighting for RAS devices is set to 16 by default. To change this setting, type or select the new value. You can assign any value between 2 and 32.
Non NDMP Device Weighting	The media request weighting for non-NDMP devices is set to 8 by default. To change the default setting, type or select the new value. You can assign any value between 2 and 32.
Media Loaded Weighting	The media request weighting for media-loaded devices is set to 4 by default. To change the default setting, type or select the new value. You can assign any value between 2 and 32.
Device Empty Loaded Weighting	The media request weighting for empty loaded devices is set to 2 by default. To change the default setting, type or select the new value. You can assign any value between 2 and 32.

By default, the Media Manager gives preference to a local device. To use any other suitable device for a job, you can set a higher media request weight for that device. For example, to ensure that the Verification phase for a job uses the same NDMP device that was used for the backup, you can set a higher value for the **Media Loaded Weighting** option. If you use the default media request weights, the Media Manager gives preference to a local device even if the required piece of media is loaded into the NDMP device.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring default interval for backup retirement scans

In time-based retirement, the time component (HH:MM) does not represent the actual retirement time. This component only represents the time due for backup retirement. The actual time of retirement is determined by the interval at which Media Manager scans the Media Database to identify the backups that it needs to retire. The default interval between two scans is 60 minutes. Thus, if the retirement time is set to 10:20, the backup is actually retired at 11:00.

You can change the default interval for backup retirement scans in the **mediamgr.cfg** file.

To configure the default interval for backup retirement scans:

- 1 Open the **mediamgr.cfg** file in a text editor. You can find this file in the **config** directory under the NetVault installation directory.
- 2 Add the following lines:

```
[Defaults:Retirement Check Granularity in Mins]
Value = <Minutes>
```

By default, the **mediamgr.cfg** file does not include the `[Defaults:Retirement Check Granularity in Mins]` section. To change the default interval, you must add this section and specify the interval. If you do not add the section, the default interval (60 minutes) is used.

For example, to set the scan interval to 30 minutes, add the following lines:

```
[Defaults:Retirement Check Granularity in Mins]
Value = 30
```

i | **NOTE:** If you set this value to 0 (zero), the savesets are retired according to the time specified in the Advanced Options set or the **Change Expiry** dialog box.

3 Save the file.

Configuring an alternate index read block size for Quest DR Series systems

If a backup stored on a Quest DR Series system was performed using a non-standard block size, the scan process is unable to read the index for that backup from the device. To import indexes for such backups, you can configure an alternate index read block size in the **mediamgr.cfg** file.

The alternate block size is used only when index scanning fails using the current block size. When the alternate block size is used, the following message is added to the NetVault logs: "Scanned index for job '<job name>' found using alternate block size <xx>, after a failed scan using original request block size <yy>."

To configure an alternate index read block size for Quest DR Series systems:

1 Open the file **mediamgr.cfg** in a text editor. You can find this file in the **config** directory under the NetVault installation directory.

2 Add the following lines to this file:

```
[Defaults:Alternate Index Read Block Size]
Type = Range
Range = 500,2147483647
Value = <Original non-standard block size>
```

3 Save the file.

Configuring Network Manager settings

This section includes the following topics:

- [About Network Manager](#)
- [Configuring timeout settings for Network Manager](#)
- [Enabling or disabling availability broadcasts](#)
- [Configuring network addresses for multihomed machines](#)
- [Reducing startup delay](#)
- [Configuring default port for Network Manager](#)
- [Configuring default port for Communications Manager](#)

About Network Manager

The Network Manager (**nvnmgr**) and Communications Manager (**nvcmgr**) support the inter-process messaging system. Both run as processes on Linux and UNIX systems and as threads within the **nvpmgr** process on Windows.

These processes perform the following functions:

- The Network Manager and Communications Manager work in tandem to transmit inter-process messages to remote clients. While the Communications Manager handles communication between the NetVault processes on local machines, Network Manager transmits the inter-process messages to remote clients.
- The Network Manager broadcasts availability messages, which help determine the status of the NetVault Clients.

You can modify the Network Manager settings from the **Change Settings** page.

i | **NOTE:** You must restart the NetVault Service to apply any changes to the Network Manager settings.

Configuring timeout settings for Network Manager

To configure timeout settings for Network Manager:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, go to Step 3.
 - To configure default settings for a NetVault Client, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**.
- 3 Under **Services**, click **Network Manager**.
- 4 In the **Network Manager** dialog box, under **Timeouts**, configure the following settings.

Table 129. Network Manager timeout settings

Setting	Description
Time to complete a remote connection	This setting controls how long the NetVault Server tries to establish a connection with a remote client. The default setting is 300 seconds. To change the timeout interval, type or select a new value. The timeout interval is specified in number of seconds.
Time to wait before dropping inactive connection(s)	This setting controls how long the NetVault Server waits before ending an inactive connection. It helps to reduce the amount of resources consumed by idle connections. The default setting is 600 seconds. To change the timeout interval, type or select a new value. The timeout interval is specified in number of seconds.
Keep Alive rate	Keep-alive messages are used to verify that a connection between two NetVault machines is still intact. The keep-alive rate controls how often the server sends keep-alive messages. The default setting is 7 seconds. To change the interval, type or select a new value. The keep-alive rate is specified in number of seconds.

Table 129. Network Manager timeout settings

Setting	Description
Time between availability broadcasts	<p>The NetVault Clients broadcast availability messages at regular intervals that provide their status and location on the network. These broadcasts are used to discover new clients and update the client status on the NetVault WebUI.</p> <p>The default interval for availability broadcasts is 600 seconds. To change the interval, type or select a new value. The broadcast interval is specified in number of seconds.</p> <p>NOTE: Setting a very small interval can increase the network traffic, while setting a large interval can cause delays in discovering the client and updating client status on the NetVault WebUI.</p>
Time between security broadcasts	<p>Security broadcasts discover the password-protection status of clients, and notify whether its password has been enabled or disabled on a client. The Client Status icons on the NetVault WebUI are updated based on these broadcasts.</p> <p>The default interval for security broadcasts is 600 seconds. To change the interval, type or select a new value. The broadcast interval is specified in number of seconds.</p>
Time between availability checks	<p>The NetVault Server performs availability checks at regular intervals to scan for changes in the network settings. If a change is detected, the server sends an interim broadcast to propagate the new settings. The broadcast system is then reset to regular pulse, reducing network traffic.</p> <p>The default interval for availability checks is 10 seconds. To change the interval, type or select a new value.</p> <p>You can reduce the interval between two checks to quickly detect and transmit changes. If you do not want to apply the changes immediately, use the default value or set it equal to the interval for Availability Broadcasts.</p>

- 5 Click **Apply** to apply the new settings and close the dialog box.

Enabling or disabling availability broadcasts

Availability broadcasts are used to discover NetVault Clients and update the client status on the NetVault WebUI. You can enable or disable availability broadcasts from the **Change Settings** page.

i | **NOTE:** Although you can disable availability broadcasts, it is not recommended.

To enable or disable availability broadcasts:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Choose machine**. In the **Choose machine** drawer, select the client, and click **Go to settings**.
- 3 Under **Services**, click **Network Manager**.
- 4 In the **Network Manager** dialog box, under **Connections**, configure the following setting:
 - **Broadcast details to machines on local network:** This check box is selected by default. Although you can disable this option, it is not recommended.

Availability broadcasts are used to discover NetVault Clients and update the client status on the NetVault WebUI. Without these broadcasts, you can only add clients by using the client FQDN or IP address.

- 5 Click **Apply** to apply the new settings and close the dialog box.

Configuring network addresses for multihomed machines

For multihomed machines, you can specify the preferred network address, fallback addresses, and barred addresses from the **Change Settings** page.

When a connection is initiated, the addresses are attempted in the following order:

- Preferred network address
- Fallback addresses
- Any other addresses in the order defined by the binding order of the host machine

To configure the network addresses for multihomed machines:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Choose machine**. In the **Choose machine** drawer, select the client, and click **Go to settings**.
- 3 Under **Services**, click **Network Manager**.
- 4 In the **Network Manager** dialog box, under **Connections**, configure the following settings.

Table 130. Network settings for multihomed machines

Setting	Description
Preferred network address	Specify the primary IP address for network connection. You can only specify a single IP address.
Comma separated list of barred address(es)	Specify the barred or blocked addresses that you do not want to use for NetVault connections. To specify multiple addresses, use a comma as delimiter. NOTE: You must restart the NetVault Service to apply any changes to the barred addresses setting.
Comma separated list of fallback address(es)	Specify the fallback addresses to use when the preferred network address is not available. To specify multiple addresses, use a comma as delimiter.

i **IMPORTANT:** You should use these settings only if machine is connected to multiple networks, either through multiple Network Interface Cards or through virtual IP addressing.
When configuring the network addresses, make sure that the preferred, barred, and fallback addresses do not conflict with each other

- 5 Click **Apply** to apply the new settings and close the dialog box.

Important notes

If the preferred address is unavailable and the NetVault Server starts using a fallback address, it does not automatically revert to the preferred address when the IP becomes available.

To force the NetVault Server to use the preferred address, do one of the following:

- Remove the client and add it again.

When you remove the client, the message “Machine <NetVault Machine Name> Has Gone Down” (Warning Level: Background) is displayed on the **View Logs** page. Wait until the timeout for remote

connection expires. The timeout interval is determined by the **Time to complete remote connection** setting (set to 60 seconds by default).

When the client is up and detected on the network, the message “Machine <NetVault Machine Name> Has Come Up” (Warning Level: Background) is displayed on the **View Logs** page. To add the client, use the **Find Machine** command, and specify its preferred network address in the box.

- Alternatively, disable or disconnect the network interface currently in use until the timeout for remote connection expires. The value configured for the **Time to complete remote connection** setting determines the duration for which the network interface needs to be disabled or disconnected. Restart the NetVault Service when the client is reported as unavailable. The NetVault Server uses the preferred address when the next connection attempt is made after you restart the service on the client.

Reducing startup delay

When the NetVault Service starts, it attempts to resolve all client IP addresses listed in the **machines.dat** file, which resides in the **etc** directory under the NetVault installation directory. Resolving all client addresses can cause a significant delay in the service startup and the machine can appear unresponsive during this time. You can reduce or avoid the startup delay by specifying the addresses that are not to be resolved during startup.

To specify the addresses that are not to be resolved:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Choose machine**. In the **Choose machine** drawer, select the client, and click **Go to settings**.
- 3 Under **Services**, click **Network Manager**.
- 4 In the **Network Manager** dialog box, under **Connections**, configure the following setting:
 - **Comma separated list of networks and addresses not to resolve:** To reduce or avoid the startup delay, specify the addresses that are not to be resolved during startup.

You can specify a single address or a range, for example, 192.168.1.2 or 192.168.x.x. If you want to configure multiple addresses or networks, use a comma as delimiter.

To get the list of networks from the **machines.dat** file, issue the following command on Windows:

```
findstr Network <NetVault home>\etc\machines.dat
```

The following is an example output:

```
C:\NetVault\etc>findstr Network machines.dat
Networks=192.168.203.1,192.168.65.1,172.16.245.1
Networks=10.1.40.81,172.16.211.1,172.16.62.1
Networks=10.1.2.37,172.16.22.1,172.16.128.1
Networks=10.1.240.222,172.16.4.1
Networks=192.168.122.1,10.1.240.52
Networks=10.1.80.83,10.1.2.68,172.16.116.1
Networks=192.168.172.1,10.1.40.98,192.168.147.1
Networks=192.168.174.1,10.1.8.71,192.168.120.1
Networks=192.168.122.1,10.1.8.79
Networks=10.1.8.132,192.168.91.1,192.168.106.1
Networks=10.1.8.163,192.168.233.1,192.168.207.1
```

```
Networks=10.1.8.16,200.0.0.1
```

...

Based on this output, you can configure the following values in the **Comma separated list of networks and addresses not to resolve** box:

```
10.0.0.0, 172.0.0.0, 192.0.0.0, 200.0.0.0
```

Some networks may have names instead of IP addresses. To find IP addresses for such clients, you can use the **nslookup** tool.

Note the following:

- To prevent the NetVault Service from resolving all networks, specify the first octet of a network address (that is, 192.0.0.0, 10.0.0.0).
- Using 0.0.0.0 does not prevent the service from resolving the networks.
- The clients that are offline and listed on the **Manage Clients** page can also cause the startup delay. To speed up the process, you can remove the clients that are offline or no longer in use.
- In a domain managed by a Windows-based NetVault Server, you may experience a long delay as the service attempts to resolve all client IP addresses using the NBNS (NetBIOS Name Service) protocol. In this environment, you can use the **Comma separated list of networks and addresses not to resolve** setting to reduce the startup delay.

You can also use this setting in a NetVault Client Cluster setup to prevent the service from resolving the private cluster IP addresses.

- 5 Click **Apply** to apply the new settings and close the dialog box.

Configuring default port for Network Manager

The Network Manager is configured to use port 20031 to open TCP and UDP sockets on a client. If this port is in use by any other application, the NetVault Service fails immediately after startup. When this error occurs, you must change the default port for Network Manager.

Use the following procedure to change the default port on the NetVault Server and all clients.

To change the default port for Network Manager:

- 1 Open the file **nvnmgr.cfg** in a text editor. You can find this file in the **config** directory under the NetVault installation directory.
- 2 Add the following lines to this file:

```
[network]
```

```
UdpPort=<port number>
```

```
TcpPort=<port number>
```

Verify that the ports are not in use by any other application. In a firewall-protected environment, make sure that the ports are open and specified in the firewall settings for the client.

- 3 Save the file.
- 4 Restart the NetVault Service to apply the new settings.

Configuring default port for Communications Manager

The Communications Manager is configured to use port 20032 to open TCP sockets on a client. If this port is in use by any other application, the NetVault Service fails immediately after startup. When this error occurs, you must change the default port for Communications Manager.

Use the following procedure to change the default port on the NetVault Server and all clients.

To change the default port for Communications Manager:

- 1 Open the file **configure.cfg** in a text editor. You can find this file in the **config** directory under the NetVault installation directory.
- 2 In the `[machine]` section, append the following line:

```
[machine]  
Comms TcpPort=<port number>
```

Verify that the port is not in use by any other application. In a firewall-protected environment, make sure that the port is open and specified in the firewall settings for the client.
- 3 Save the file.
- 4 Restart the NetVault Service to apply the new settings.

Configuring Process Manager settings

This section includes the following topics:

- [About Process Manager](#)
- [Configuring shared memory settings](#)

About Process Manager

The Process Manager (**nvpmgr**) runs on all NetVault machines.

This process manages all other NetVault processes. It creates and destroys the transient processes. The Process Manager also manages the allocation of the shared memory area for the process table and trace buffers.

You can modify the Process Manager settings from the **Change Settings** page.

i | **NOTE:** You must restart the NetVault Service to apply any changes to the Process Manager settings.

Configuring shared memory settings

To configure shared memory settings for Process Manager:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Choose machine**. In the **Choose machine** drawer, select the client, and click **Go to settings**.

- Under **Services**, click **Process Manager**.
- 3 In the **Process Manager** dialog box, configure the following settings.

Table 131. Shared memory settings for Process Manager

Setting	Description
Shared Memory Allocated to the Trace Buffer of Each Process	<p>This setting controls the amount of shared memory allocated to the individual trace buffers of each NetVault process.</p> <ul style="list-style-type: none"> • The default value is 31KB on all platforms. • The minimum value is 1KB on all platforms. • The maximum value is 32KB on all platforms. <p>To increase the shared memory for individual trace buffers, type or select a new value. The shared memory is allocated in KB.</p> <p>NOTE: Increasing the value of this setting could affect the performance and scalability of NetVault. For example, it could reduce the maximum number of concurrent data transfers.</p>
Shared Memory Used for Process Table	<p>This setting controls the amount of shared memory allocated to the Process Table that maintains the details of all current NetVault processes.</p> <ul style="list-style-type: none"> • The default value for server and client is 16384KB and 3864KB respectively on all platforms. • The minimum value is 2000KB on Windows and 500KB on Linux and UNIX. • The maximum value is 128000KB on all platforms. <p>On NetVault machines running a number for concurrent processes, you can increase the amount of shared memory for the Process Table to get optimum performance. To increase the shared memory, type or select a new value. The shared memory is allocated in KB.</p> <p>On Windows, the shared memory is allocated dynamically.</p> <p>On Linux and UNIX, the shared memory for process table is allocated from the shared memory pool. To increase the shared memory for process table on these platforms, you must first increase the overall shared memory pool. You can increase the shared memory by configuring the shmmax variable in the system configuration file. For more information about increasing the shared memory pool, consult the relevant O/S documentation.</p>

- 4 Click **Apply** to apply the new settings and close the dialog box.

Configuring RAS device settings

This section includes the following topic:

- [Configuring connection settings for NetVault SmartDisk](#)

Configuring connection settings for NetVault SmartDisk

The connection settings specify the period for which NetVault makes retry attempts to establish a TCP connection to a NetVault SmartDisk and the time between retry attempts.

By default, the connection retry mechanism is disabled. Under most circumstances, you do not need to configure the connection settings for NetVault SmartDisk. If your backups report “failure to bind port” errors, then you can configure these settings to enable the connection retry mechanism.

To configure connection settings for NetVault SmartDisk devices:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **RAS Devices**. In the **RAS Devices** dialog box, configure the following settings:
 - **Time limit for retrying a SmartDisk connection:** This setting specifies how long the NetVault Server tries to establish a connection to a NetVault SmartDisk device after the initial attempt fails.

The default value for this setting is zero (0). When set to zero (0), NetVault does not make any retry attempts. If NetVault fails to establish a connection due to insufficient ports, you can configure this setting to retry the connections for a specified period that allows the operating system to recycle the sockets for use. You can set the retry time to any value between 0 and 300 seconds.
 - **Time (in seconds) between retries for a SmartDisk connection:** This setting specifies the time between retry attempts. You can set the retry interval to any value between 0 and 60 seconds.
- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring Schedule Manager settings

This section includes the following topics:

- [About Schedule Manager](#)
- [Configuring default settings for Schedule Manager](#)
- [Configuring default job priority settings](#)

About Schedule Manager

The Schedule Manager (**nvsched**) runs on the NetVault Server. This process manages the job schedules and queues. It also manages the Scheduler Database.

The Schedule Manager initiates the Job Manager to start a job instance, and schedules the next instance for recurring jobs; the Job Manager runs the job. This process also updates the **Job Status** page and provides job scheduling data to the Reporting utility.

You can modify the Schedule Manager settings from the **Change Settings** page.

Configuring default settings for Schedule Manager

To configure default settings for Schedule Manager:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Schedule Manager**. In the **Schedule Manager** dialog box, configure the following settings.

Table 132. Schedule Manager settings

Setting	Description
Number of days to keep job status	<p>This setting determines display period for job status records.</p> <p>The default setting is seven days. To change the display period, type or select a new value. The display period is specified in number of days.</p> <p>NOTE: This setting just hides the display of job status records on the NetVault WebUI. It does not delete the records from the Scheduler Database.</p>
Number of days to keep report job histories in the database	<p>This setting determines the retention period for report job history.</p> <p>The default setting is 90 days. To change the retention period, type or select a new value. The retention period is specified in number of days. After the retention period expires, these records are removed from the Scheduler Database.</p>
Number of days to keep other job histories in the database	<p>This setting determines the retention period for backup and restore job history.</p> <p>The default setting is 90 days. To change the retention period, type or select a new value. The retention period is specified in number of days. After the retention period expires, these records are removed from the Scheduler Database.</p>
Number of days to keep non-scheduled jobs in the database	<p>This setting determines the retention period for non-repeating jobs.</p> <p>The default value for this setting is 0 (zero). With the default setting, the job definitions of non-repeating jobs are retained indefinitely. To delete these jobs after a specified period, type or select a new value. The retention period is specified in number of days.</p> <p>NOTE: Because the jobs using the Triggered schedule type have no physical scheduled time, they are also deleted after the specified period elapses. If you are using Triggered jobs in your NetVault environment, do not use this setting.</p> <p>NOTE: This setting is also applied to report jobs that use any non-repeating schedule type. Therefore, we recommend that you do not change the default value for this setting. If the default value is changed, canned reports that have run once using any non-repeating schedule are removed from the View Reports page after the specified interval.</p>
Maximum simultaneously active jobs	<p>This setting determines the maximum number of simultaneous jobs that you can run in NetVault.</p> <p>By default, the Schedule Manager supports a maximum of 200 simultaneous jobs, including backup, restore and report jobs. To change the setting, type or select a new value.</p> <p>NOTE: Each active job requires some amount of shared memory. An increase in the number of active jobs might have an impact on the overall performance of NetVault.</p>
Exclude reports jobs from jobs management views	<p>To display report jobs on the Job Status page, clear this check box.</p> <p>These jobs are excluded by default.</p>
Exclude reports jobs from policy management views	<p>To display report jobs on the Manage Policy page, clear this check box.</p> <p>These jobs are excluded by default.</p>
Exclude restore jobs from policy management views	<p>To display restore jobs on the Manage Policy page, clear this check box.</p> <p>These jobs are excluded by default.</p>

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring default job priority settings

The Schedule Manager assigns a default priority level to each job type (backup, restore, and report). These default settings are applied globally to all backup, restore, and report jobs. The priority settings are used to prioritize resource allocation when two or more jobs are scheduled to run at the same time. You can change the default job priority settings from the **Change Settings** page.

By default, the Schedule Manager assigns the following priority levels:

- Backup job: 30
- Restore job: 20
- Report job: 50

To change the priority settings globally for all jobs, use the following procedure. You can override the priority setting for an individual job by configuring the **Job Priority** option in the Schedule Set. For more information, see [Creating Schedule Sets, Table 35](#).

To configure default job priority settings:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Schedule Manager**. In the **Schedule Manager** dialog box, configure the following settings:
 - **Default backup job priority**
 - **Default restore job priority**
 - **Default report job priority**

Type or select a value from 1 through 100. 1 denotes highest priority, while 100 denotes lowest priority. A job with a priority level of zero (0) runs as a background task.

i **NOTE:** If you configured a Secondary Copy, the Secondary Copy job has two requests: one for read and another for write. When you set a priority to the Secondary Copy job, the first request (source request) receives the priority that the user assigns to it; the subsequent request (destination request) served as first priority in the background with priority 0 or 1. This implementation is to avoid waiting for the destination media request to be served if the source media is available.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring Web Service settings

This section includes the following topics:

- [Configuring Web Service to disable ciphers](#)
- [Configuring Web Service to disable protocols](#)
- [Enabling Web Service auditing](#)
- [Configuring timeout period for client operations](#)
- [Configuring timeout period for saveset removal request](#)

Configuring Web Service to disable ciphers

You can modify the Web Service configuration to disable ciphers for incoming web service connections settings to prevent from allowing one or more ciphers. You can configure this setting from the **Change Settings** page.

To configure Web Service to disable ciphers:

- 1 In the Navigation pane, click **Change Settings**.
 - 2 On the NetVault Server Settings page, under **Services**, click **Web Service**. In the **Web Service** dialog box, configure the following setting:
 - **Ciphers to disable for incoming Web Service connections:** To disable a cipher for incoming web service connections, enter its value given in the second column of the table as mentioned in the link <https://www.openssl.org/docs/man1.0.2/apps/ciphers.html>. However, if the entered cipher does not match with ciphers of incoming web service connection, no cipher is disabled. To disable multiple ciphers, enter the values separated by comma. By default, this field is blank and all the ciphers are allowed.
- Click **Apply** to apply the new settings and close the dialog box.

Configuring Web Service to disable protocols

You can modify the Web Service configuration to disable protocols for incoming web service connections settings to prevent from allowing one or more protocols. You can configure this setting from the **Change Settings** page.

To configure Web Service to disable protocols:

- 1 In the Navigation pane, click **Change Settings**.
 - 2 On the NetVault Server Settings page, under **Services**, click **Web Service**. In the **Web Service** dialog box, configure the following setting:
 - **Protocols to disable for incoming Web Service connections:** Enter the protocol to disable for incoming web service connections. However, if the entered protocol does not match with protocols of incoming web service connection, no protocol is disabled. To disable multiple protocols, enter the values separated by comma. By default, this field is blank and all the protocols are allowed.
- Click **Apply** to apply the new settings and close the dialog box.

Enabling Web Service auditing

Web Service auditing is disabled by default. You can enable Web Service auditing from the **Change Settings** page.

To enable Web Service auditing:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Web Service**. In the **Web Service** dialog box, configure the following setting:
 - **Enable Web Service auditing:** To track Web Service requests, select this check box. When Web Service auditing is enabled, NetVault creates an audit log file “**webservice-audit.txt**” in the NetVault installation directory to record the Web Service requests.
- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring timeout period for client operations

You can modify the Web Service configuration settings to resolve any timeout errors that occur during client operations like browsing or opening the selection tree nodes. These settings can be configured from the NetVault Server Settings page.

To configure timeout period for client operations:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Web Service**. In the **Web Service** dialog box, configure the following setting.

Table 133. Web Service timeout settings for client operations

Setting	Description
Physical Client Short Timeout	The amount of time the Web Service waits for short duration operations, like requesting the plug-in list for a client. The default timeout period is 30 seconds. You can change it to any value between 10 and 120 seconds.
Physical Client Medium Timeout	The amount of time the Web Service waits for medium duration operations, like removing a server (for example, Database Server). The default timeout period is 60 seconds. You can change it to any value between 60 and 120 seconds.
Physical Client Long Timeout	The amount of time the Web Service waits for long duration operations, like populating the selection tree. The default timeout period is 3,000 seconds. You can change it to any value between 120 and 60,000 seconds.

ⓘ | **NOTE:** You must restart the NetVault Service to apply any changes to these settings.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring timeout period for saveset removal request

By default, the timeout period for saveset removal request is set to 600 seconds. When the Web Service Worker Process (**nvswsworker**) sends the delete request to Media Manager, it waits for 600 seconds. If the Media Manager is unable to complete the request within this period, the Web Service Worker Process times out and reports an error (“Failed to receive reply from Media Manager”). If you receive this error, change the **SavesetRemoveTimeout** setting in the **webservice.cfg** file.

To configure the timeout period for saveset removal request:

- 1 Open the file **webservice.cfg** in a text editor. You can find this file in the **config** directory under the NetVault installation directory.
- 2 In the [WebService] section, configure the value for the **SavesetRemoveTimeout** setting:

```
[WebService]
```

```
SavesetRemoveTimeout=<Value>
```

The default value for this setting is 600 seconds.

- 3 Save the file.

Configuring Auditor Daemon settings

This section includes the following topics:

- [About Auditor Daemon](#)

- [Configuring Auditor to log only failed requests](#)
- [Modifying the purge policy for audit logs](#)

About Auditor Daemon

The Auditor Daemon (**nvavp**) tracks and controls all user activities in NetVault. This process runs on the NetVault Server. The Auditor Daemon validates each user request, and depending on the assigned privileges, allows or denies a request. Audit log messages are stored in the NetVault Database.

By default, the Auditor Daemon logs every user activity regardless of whether the request is granted or denied. It automatically deletes the log messages that are older than 31 days. You can configure the Auditor Daemon to log only failed user requests. You can also modify the default settings to change the maximum age for log messages.

You can modify the Auditor Daemon settings from the **Change Settings** page.

Configuring Auditor to log only failed requests

By default, the Auditor Daemon logs every user request, regardless of whether the request is granted or denied. You can configure the Auditor Daemon to log only failed user requests.

To configure the Auditor Daemon to log only failed requests:

- 1 In the Navigation pane, click **Change Settings**.
 - 2 On the NetVault Server Settings page, under **System and Security**, click **Auditing**. In the **Auditing** dialog box, configure the following setting:
 - **Only record failed requests in the audit trail:** To log only failed user requests, select this check box.
- i** | **NOTE:** You must restart the NetVault Service to apply any changes to this setting.
- 3 Click **Apply** to apply the new settings and close the dialog box.

Modifying the purge policy for audit logs

Audit log messages that are older than 31 days are automatically purged from the database. You can use the configuration settings for Logging Daemon to change the maximum age for log messages.

To modify the purge policy for audit logs:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **System and Security**, click **Auditing**. In the **Auditing** dialog box, configure the following settings.

Table 134. Auto-purge settings for audit logs

Setting	Description
Purge entries by age	To purge log entries that are older than the maximum age set for the log messages, leave this check box selected. The maximum log age is specified in the Purge entries that are older than box. Automatic purging of log messages is disabled if you clear this check box. NOTE: The audit logs can consume a considerable amount of disk space. Therefore, you must periodically purge the log files. Since you cannot purge the audit logs manually, we recommend that you do not disable this feature.
Purge entries that are older than (days)	Type or select the maximum age for the log messages. The log age is specified in number of days. The default value is 31 days.
Time interval between purges (hours)	Type or select the interval between two purge events for audit logs. The time interval is specified in number of hours. The default interval is 24 hours.
Purge events that are older than (days)	Type or select the maximum age for the event. The event age is specified in number of days. The default value is 31 days.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring firewall settings

Firewall settings are required to communicate with NetVault Clients that reside outside the firewall. You can use these settings to specify the TCP/IP ports for establishing data transfer channels, message channels, and broadcast channels through the firewall.

You can configure the firewall settings when adding a client, or update these settings from the **Change Settings** page. You can also use the NetVault WebUI or Txtconfig utilities to configure the firewall settings for new or existing clients.

To configure firewall settings:

- 1 Access the firewall settings by using any of the following methods:
 - **NetVault WebUI:**
 - a In the Navigation pane, click **Change Settings**.
 - b On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, go to Step C.
 - To configure default settings for a NetVault Client, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**.
 - c Under **System and Security**, click **Firewall**.
 - **Txtconfig:**
 - a In a terminal or command prompt window, type `txtconfig`, and press **Enter** or **Return**.
 - b On the **Machine** page, press **p** to view the main menu, and then press the option number for the **Firewall** page.
- 2 Configure the following settings:
 - **Listen ports for devices**
 - **Connect port for devices**
 - **Listen ports for NetVault message channels**
 - **Connect ports for NetVault message channels**

- **Connect ports for NDMP control channels**
- **Listen ports for NDMP data channels**
- **Connect ports for inter-machine setup**

For more information, see [About firewall settings](#).

- 3 Save the settings or click **Next** to complete the configuration procedure.

Configuring general settings

This section includes the following topics:

- [Relocating default directories](#)
- [Relocating the NetVault Database directory](#)
- [Configuring TCP/IP buffer sizes](#)
- [Changing language and locale settings](#)
- [Disabling pre-installation package compatibility check](#)
- [Configuring license expiration warning period](#)

Relocating default directories

You can relocate database, trace, logs, reports, stats, and temporary directories to a different drive or volume to alleviate low disk space issues.

i | **NOTE:** You must restart the NetVault Service to apply any changes to these settings.

To relocate the default directories:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**.
- 3 Under **System and Security**, click **General**. In the **General** dialog box, configure the following settings.

Table 135. Default directory paths

Setting	Description
Database Directory	<p>This directory stores the module list, licensefile or files, and NetVault Database.</p> <p>The default path for the database directory is <NetVault Home>\db on Windows and <NetVault Home>/db on Linux.</p> <p>For instructions on relocating the database directory, see Relocating the NetVault Database directory.</p>
Trace Directory	<p>This directory stores the NetVault trace logs.</p> <p>The default path for the trace directory is <NetVault Home>\trace on Windows and <NetVault Home>/trace on Linux. To relocate the directory, type the full path. The specified path must exist on the target drive or volume. If the NetVault Server is unable to find the path, it fails to relocate the directory.</p> <p>NOTE: The trace logs contain large volume of data. Therefore, this directory should not be relocated to a network share.</p>
Log Directory	<p>This directory stores the NetVault log files.</p> <p>The default path for the logs directory is <NetVault Home>\logs on Windows and <NetVault Home>/logs on Linux. To relocate the directory, type the full path. The specified path must exist on the target drive or volume. If the NetVault Server is unable to find the path, it fails to relocate the directory.</p>
Temporary Directory	<p>This directory stores the temporary files generated during various NetVault operations.</p> <p>The default path for the temporary directory is <NetVault Home>\tmp on Windows and <NetVault Home>/tmp on Linux. To relocate the directory, type the full path. The specified path must exist on the target drive or volume. If the NetVault Server is unable to find the path, it fails to relocate the directory.</p> <p>NOTE: The temporary directory holds the Backup Saveset Index. The free disk space on this directory should be at least three times the size of the index file to perform backups and restores properly. For example, if the backup index file is 3GiB in size, the temporary directory should have at least 9GiB of free disk space. If the temporary directory does not have sufficient space, warnings and logs are generated.</p>
Reports Directory	<p>This directory stores the report templates.</p> <p>The default path for the reports directory is <NetVault Home>\reports on Windows and <NetVault Home>/reports on Linux. To relocate the directory, type the full path.</p> <p>Before configuring a new path, you must copy the directory contents to the new path. If a new path is configured without relocating the directory contents, an error message appears (“Provider 'NVBUPhysicalClient' failed”), and the NetVault Service fails to restart.</p>
Statistics Directory	<p>This directory stores data collected by the nvstatsmgr process.</p> <p>The default path for the statistics directory is <NetVault Home>\stats on Windows and <NetVault Home>/stats on Linux. To relocate the directory, type the full path. The specified path must exist on the target drive or volume. If the NetVault Server is unable to find the path, it fails to relocate the directory.</p>

- 4 Click **Apply** to apply the new settings and close the dialog box.

Relocating the NetVault Database directory

To relocate the NetVault Database directory, use the procedures described in the following sections:

- [Changing the NetVault Database directory on a Linux-based machine](#)
- [Changing the NetVault Database directory on a Windows-based machine](#)

Changing the NetVault Database directory on a Linux-based machine

To change the Database Directory on a Linux-based NetVault Server:

- 1 Stop the NetVault Service.
 - If the **systemd** is supported on the system:
Run the `systemctl stop netvault` command.
 - If the **systemd** is not supported on the system:
 - a In a terminal window, type `txtconfig`, and press **Enter** or **Return**.
 - b On the **Machine** page, press **p** to view the main menu, and then press the option number for the **Services** page.
 - c If the service is running, press the option number to stop the service.
- 2 Stop the **netvault-pgsql** service.
 - If the **systemd** is supported on the system:
Run the `systemctl stop netvault-pgsql` command.
 - If the **systemd** is not supported on the system:
Run the `/etc/init.d/netvault-pgsql stop` command
- 3 Manually move the **db** directory to the new location.
- 4 From the **Txtconfig** utility, change the Database Directory path.
 - a On the **Machine** page, press **p** to view the main menu, and then press the option number for the **General** page.
 - b Press the option number for the **Database Directory** setting, and type the full path.
The specified path must exist on the target drive or volume. If the NetVault Server is unable to find the path, it fails to relocate the directory.
 - c Press **s** to save the setting, and then press **q** to quit Txtconfig.
- 5 In the NetVault installation directory, open the **config** directory, and then open the **pgdb.cfg** configuration file in the text editor.
In the **[InstallInfo]** section, change the **datadir** to specify new path for **pgsql**.
datadir=<new db path>/pgsql
- 6 Change the database path from the service script:
 - `/etc/init.d/netvault-pgsql` file.
 - `/usr/lib/systemd/system/netvault-pgsql.service` file.
- 7 Change the file permissions of the owner and group for **pgsql** directory from the new path:
`chmod 700 <new DB path>/pgsql`
`chown netvault-pgsql:netvault-pgsql <new DB path>/pgsql`
- 8 Start the **netvault-pgsql** service.
 - If the **systemd** is supported on the system:
Run the `systemctl start netvault-pgsql` command.

- If the **systemd** is not supported on the system:
Run the `/etc/init.d/netvault-pgsql start` command

9 Start the NetVault Service.

Changing the NetVault Database directory on a Windows-based machine

To change the Database Directory on a Windows-based NetVault Server:

- 1 From the Windows Services Management Console, stop the **netvault-pgsql** service and **NetVault Process Manager** service.
- 2 Manually move the **db** directory to the new location. Make sure that the user under which the **netvault-pgsql** service runs has the ownership of the new Database Directory.
- 3 In the NetVault installation directory, open the config directory, and then open the file `configure.cfg` in a text editor.

In the [Directories:Database] section, change the Value to specify the full path.

Value=<full path>
- 4 In the NetVault installation directory, open the **config** directory, and then open the **pgdb.cfg** file in the text editor.

In the [InstallInfo] section, change the **datadir** to specify new path for **pgsql**.

datadir=<new db path>\pgsql
- 5 Open the Registry Editor.
- 6 Expand the key `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\netvault-pgsql`, and change the value for the **ImagePath** property to reflect the new path for the **Database Directory**.
- 7 On the Windows Services Management Console, right-click the **netvault-pgsql** service, and select **Properties**. On the **General** tab, verify that the Path to executable parameter displays the new database path.
- 8 Start the **netvault-pgsql** and **NetVault Process Manager** service.

Configuring TCP/IP buffer sizes

In most cases, the default buffer sizes for the TCP/IP sockets are adequate, and should not be changed. For guidance on optimal buffer sizes and TCP/IP tuning, see your OS documentation.

To change buffer sizes for the TCP/IP sockets:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**.
- 3 Under **System and Security**, click **General**. In the **General** dialog box, configure the following settings.

Table 136. TCP/IP send and receive buffer sizes

Setting	Description
Minimum network send buffer size	This setting determines the minimum send buffer size for a TCP/IP socket. The default value is 1KB. To adjust the buffer size, type or select the new value. The buffer size must be set in KB.
Maximum network send buffer size	This setting determines the maximum send buffer size for a TCP/IP socket. The default value is 16384KB. To adjust the buffer size, type or select the new value. The buffer size must be set in KB.
Minimum network receive buffer size	This setting determines the minimum receive buffer size for a TCP/IP socket. The default value is 1KB. To adjust the buffer size, type or select the new value. The buffer size must be set in KB.
Maximum network receive buffer size	This setting determines the maximum receive buffer size for a TCP/IP socket. The default value for this option is 16384KB. To adjust the buffer size, type or select the new value. The buffer size must be set in KB.

- 4 Click **Apply** to apply the new settings and close the dialog box.

i **NOTE:** On Windows, the registry settings must be modified to increase the TCP/IP buffer size. For more information about these settings, see the KB article <https://support.microsoft.com/en-us/kb/823764> from Microsoft.

Changing language and locale settings

To change the language and locale for NetVault:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**.
- 3 Under **System and Security**, click **General**. In the **General** dialog box, configure the **Language Selection** setting.

Select the preferred language for NetVault. The available options are:

- Chinese (Simplified)
 - English
 - French
 - German
 - Japanese
 - Korean
- 4 Click **Apply** to apply the new settings and close the dialog box.
 - 5 Close the NetVault WebUI, and open it in a new browser tab or window.

Changing language settings for NetVault WebUI

To change the display and input language for NetVault WebUI:

- On the Header pane, click **Language**, and from the list of available languages, select the display and input language for NetVault WebUI. The available options are:
 - Chinese (zh)
 - English (en)
 - French (fr)
 - German (de)
 - Japanese (ja)
 - Korean (ko)

Disabling pre-installation package compatibility check

Before installing a plug-in, NetVault ensures that the installation package is compatible with the client OS type and bitness. You can disable the pre-installation compatibility checks, if necessary.

To disable pre-installation compatibility checks for packages:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**.
- 3 Under **System and Security**, click **General**. In the **General** dialog box, configure the following setting:
 - **Package Install System Check:** This check box is selected by default. Before installing any plug-in, NetVault ensures that installation package is compatible with the client OS and bit-type. The installer reports an error if the package is incompatible.

To disable pre-installation package compatibility checks, clear the check box.

i | **NOTE:** When you disable package check, you risk installing an incompatible package.

- 4 Click **Apply** to apply the new settings and close the dialog box.

Configuring license expiration warning period

By default, the NetVault WebUI shows the license expiration message seven days before the product expiration date. The dialog box is displayed when you log on to the NetVault WebUI.

To change the license expiration warning period:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Choose machine**. In the **Choose machine** drawer, select the client, and click **Go to settings**.

- 3 Under **System and Security**, click **General**. In the **General** dialog box, configure the following setting:
 - **License expiry warning threshold:** Specify how many days in advance NetVault notifies you about the license expiration. The message is displayed every time you log on to the WebUI. The default period is seven days.
- 4 Click **Apply** to apply the new settings and close the dialog box.

Configuring security settings

This section includes the following topics:

- [Disabling password protection for a client](#)
- [Changing NetVault password](#)

Disabling password protection for a client

To disable password-protection for a NetVault Client:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Click **Clients**.
- 3 In the **NetVault Clients** table, select the client, and click **Next**.
- 4 Under **System and Security**, click **Security**. In the **Security** dialog box, configure the following setting:
 - **Disable Security:** To add or access a client without using its NetVault password, select this check box.
- 5 Click **Apply** to apply the new settings and close the dialog box.

Changing NetVault password

To change the NetVault Server or Client password:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Choose machine**. In the **Choose machine** drawer, select the client, and click **Go to settings**.
- 3 Under **System and Security**, click **Security**. In the **Security** dialog box, configure the following settings.

Table 137. Change NetVault password

Setting	Description
Master Password	Specify a password for the NetVault machine. The password can contain a maximum of 100 characters. It cannot contain the following characters: \ and spaces. The NetVault password is used to add and access the clients. For more information about NetVault passwords, see the <i>Quest NetVault Installation Guide</i> .
Confirm Password	Type the NetVault password again for confirmation.

- 4 Click **Apply** to apply the new settings and close the dialog box.

Synchronizing NetVault Time

This section includes the following topics:

- [About NetVault Time](#)
- [Configuring an alternate NetVault Time Server](#)

About NetVault Time

NetVault designates a Time Server, called NetVault Time Server, to synchronize time on the server and all client machines in the NetVault Domain. Typically, the NetVault Server acts as the NetVault Time Server, and the system time on this machine is the NetVault Time. However, you can designate any other NetVault machine as the NetVault Time Server. NetVault ignores the local time on the clients and uses NetVault Time for all time-specific operations, such as job scheduling, reporting, and tracing.

Configuring an alternate NetVault Time Server

To establish an alternate NetVault machine as the NetVault Time Server:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Choose machine**. In the **Choose machine** drawer, select the client, and click **Go to settings**.
- 3 Under **System and Security**, click **Time Sync**. In the **Time Sync** dialog box, configure the following settings.

Table 138. Time synchronization settings

Setting	Description
This machine is the NetVault time server	This check box is selected by default on the NetVault Server. To configure an alternate time server, clear this check box on the NetVault Server.
Synchronize NetVault time with system	Type the NetVault name of the alternate time server.
Number of ping-pongs to determine time difference	Type or select the number of time packets to be exchanged while polling. The default value is 5.
Number of hours between time sync updates	Specify how often the NetVault Server checks to see if it is still synchronized with the time server. The default polling interval is 24 hours.
Number of milliseconds of time difference allowed between 2 servers	Type or select the allowed time variance. By default, NetVault allows 1000-millisecond variance between the NetVault Server and the NetVault Time Server.

- 4 Click **Apply** to apply the new settings and close the dialog box.

Configuring default settings for global notification methods

This section includes the following topics:

- [Configuring email server settings for notifications](#)
- [Configuring SysOp email ID for notifications](#)
- [Configuring a default printer for notifications](#)
- [Configuring network manager host settings for notifications](#)

Configuring email server settings for notifications

To send notifications by email, you must configure the Outgoing Email Server (SMTP) settings. You can configure the SMTP settings from the **Change Settings** page.

To configure email server settings for notifications:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Notification**. In the **Notification** dialog box, under **Mail Server**, configure the following options.

Table 139. Mail server settings for notifications

Setting	Description
Outgoing E-mail Server (SMTP)	Type the machine name or IP Address of the Mail Server.
Port number	The default SMTP listener port is port number 25. If the Mail Server is configured to listen on a different port, specify the port number.
NetVault's E-mail Address	Type the sender's email address for notifications.
NetVault's Real Name	Type the sender's name for email notifications.
Perform Authentication	To perform authentication, select this check box. You can use this setting only if your Mail Server supports LOGIN or PLAIN protocols.
Authentication Account	Specify a user account that can be used for SMTP authentication. This option is only required if the Perform Authentication check box is selected. If no user account is specified, the user name from NetVault's E-mail Address is used for authentication. For example, if you specify "User-A@mycompany.com" in the NetVault's E-mail Address box, NetVault uses "User-A" to perform SMTP authentication.
Authentication Password	Type the password for the SMTP authentication account.
Confirmation Authentication Password	Re-type the password for confirmation.
Use hostname instead of nvsendmail	Select this check box to use the Fully Qualified Domain Name (FQDN) instead of nvsendmail with the EHLO and HELO commands. NetVault uses the host name in mail server messages, so the system must be set up with an FQDN to display the host name in FQDN format.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring SysOp email ID for notifications

To use the Sysop Email notification method, you must configure the SysOp email address. You can configure the email address from the **Change Settings** page.

To configure the SysOp email ID for notifications:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Notification**. In the **Notification** dialog box, under **Global Notification**, configure the following settings.

Table 140. SysOp mail ID for notifications

Setting	Description
Email Address of System Administrator	Type the email address of the Sysop (Administrator). Use a comma (,) to separate multiple email addresses.
Real Name of System Administrator	Type the real name of the Sysop or Administrator.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring a default printer for notifications

To use the **Print Report** method on Windows, you can either configure a default printer or set the variable NVPRINTER while setting up a notification method. You can configure the default printer from the **Change Settings** page.

To configure a default printer for notifications:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Notification**. In the **Notification** dialog box, under **Default Printer**, configure the following setting:
 - **Default Printer Name:** To configure a default printer for the **Print Report** notification method, type the printer name. Specify the exact string that is used to identify the printer in the Windows OS.
This setting is only required on Windows-based NetVault Servers.
- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring network manager host settings for notifications

Before you can use SNMP Trap method, you must configure the network manager host settings. You can configure these settings from the **Change Settings** page.

To configure Network Manager host settings for notifications:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Notification**. In the **Notification** dialog box, under **SNMP Traps**, configure the following settings.

Table 141. Default settings for SNMP Trap notifications

Setting	Description
Trap Target Host	Type the network name or FQDN (Fully Qualified Domain Name) of the Network Manager host to which the SNMP traps are to be sent.
Port Number	The default listener port for SNMP traps is port number 162. If the host is configured to listen on a different port, specify the port number.
Community String	An SNMP community string is a password that is used to authenticate messages that are sent between the Network Manager Host and the agent. The community string is included in every packet that is transmitted between the SNMP manager and the SNMP agent. This Community String is set to "public," which is the default read-only community string for most network devices. We recommend that you modify this default string and set a new password for SNMP traps.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring the reporting utility

This section includes the following sections:

- [About reporting utility](#)
- [Customizing CSV report templates](#)
- [Customizing HTML report templates](#)
- [Customizing plain text report templates](#)
- [Configuring timeout setting for report generation](#)
- [Configuring default mail format type for report](#)
- [Configuring default settings for Statistics Manager](#)
- [Creating a global purge policy for the Reports Database](#)
- [Creating table-specific purge policy](#)

About reporting utility

The NetVault reporting utility provides a selection of canned reports that can be generated and viewed in HTML, text, and Comma Separated Value (CSV) formats. For more information about reports, see [Reporting in NetVault](#).

NetVault uses the Statistics Manager (**nvstatsmgr**) and Reports Database Manager (**nvrepdbmgr**) processes to gather and transmit data for the canned reports:

- **Statistics Manager:** This process runs on the NetVault Server and Client machines. The Statistics Manager collects drive statistics, event history, media requests, server capacity, and transfer information.
- **Reports Database Manager:** This process runs only on the NetVault Server. The Reports Database Manager polls the Statistics Manager at regular intervals to retrieve the collected data, and writes the data to the Reports Database. It provides the information in the reports database to the reporting utility and performs periodic purging of the reports database.

You can customize the report templates and change the default settings for the Statistics Manager and Reports Database Manager from the NetVault WebUI.

Customizing CSV report templates

NetVault uses built-in templates to format the output for CSV reports. You can customize these templates to define custom delimiters or add line breaks, tabs, or separators. The formatting styles are applied globally to all CSV report templates. Only users familiar with the use of control characters and escape sequences in CSV output should configure these settings. Improper configuration can cause NetVault to produce incorrect output.

To customize the output format for CSV reports:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Reporting**. In the **Reporting** dialog box, under **CSV Output**, configure the following settings:
 - **Default CSV column header pre-text**
 - **Default CSV column header post-text**
 - **Default CSV header field pre-text**
 - **Default CSV header field post-text**
 - **Default text to output for a CSV report with no records**
 - **Default CSV format pre-text**
 - **Default CSV format post-text**
 - **Default CSV format field pre-text**
 - **Default CSV format field post-text**

Use the pre-text fields to specify formatting styles (for example, line breaks, or separators) or text for headers, rows showing the total or average values, or body cells.

Use the post-text fields to specify formatting styles (for example, line breaks or separators) and custom delimiters.

Use the **Default text to output for a CSV report with no records** field to change the default text “Nothing to display” with any custom text.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Customizing HTML report templates

NetVault uses built-in templates to format the output for HTML reports. You can customize these templates to apply various formatting styles. The formatting styles are applied globally to all HTML report templates. Only users having a good knowledge of HTML should configure these settings. Improper configuration can cause NetVault to produce incorrect output.

To customize the output format for HTML reports:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Reporting**. In the **Reporting** dialog box, under **HTML Text Output**, configure the following settings:
 - **Pre-text for HTML output of plain text**
 - **Post-text for HTML output of plain text**
 - **Default HTML column header pre-text**
 - **Default HTML column header post-text**
 - **Default HTML header field pre-text**
 - **Default HTML header field post-text**

- **Default text to output for an HTML report with no records**
- **Default HTML total row pre-text**
- **Default HTML total row post-text**
- **Default HTML total field pre-text**
- **Default HTML total field post-text**
- **Default HTML average row pre-text**
- **Default HTML average row post-text**
- **Default HTML average field pre-text**
- **Default HTML average field post-text**
- **Default HTML format pre-text**
- **Default HTML format post-text**
- **Default HTML format field pre-text**
- **Default HTML format field post-text**
- **Default job header/footer location**

Use the pre-text fields to specify opening HTML tags for formatting styles (for example, font type or font size) or text for headers, rows showing the total or average values, or body cells.

Use the post-text fields to specify closing HTML tags for custom formatting styles.

Use the **Default text to output for an HTML report with no records** field to change the default text “Nothing to display” with any custom text.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Customizing miscellaneous report templates

NetVault uses built-in templates to format miscellaneous reports. You can customize these templates to add line breaks, tabs, or separators. The formatting styles are applied globally to all plain text report templates. Only users familiar with the use of control characters and escape sequences in text output should configure these settings. Improper configuration can cause NetVault to produce incorrect output.

To customize the output format for miscellaneous reports:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Reporting**.
- 3 In the **Reporting** dialog box, under **Miscellaneous**, configure the following settings:
 - **Default output type**
 - **Default Mail Format**
 - **Default filter view**
 - **Duration to keep produced reports (days)**
 - **Report Generation Timeout (in seconds)**
 - **Default keep filtered out records setting**

By default, the output type of report is HTML Text. You can change it to Plain Text or Comma Separated Values.

i | **NOTE:** If you change the output type, you must also modify the template file to match the selected output type.

- 4 To apply the new settings, click **Apply**, and then close the dialog box.

Customizing plain text report templates

NetVault uses built-in templates to format the plain text reports. You can customize these templates to add line breaks, tabs, or separators. The formatting styles are applied globally to all plain text report templates. Only users familiar with the use of control characters and escape sequences in text output should configure these settings. Improper configuration can cause NetVault to produce incorrect output.

To customize the output format for plain text reports:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Reporting**. In the **Reporting** dialog box, under **Plain Text Output**, configure the following settings:
 - **Default plain text column header pre-text**
 - **Default plain text column header post-text**
 - **Default text to output for a plain text report with no records**
 - **Default plain-text total row pre-text**
 - **Default plain-text total row post-text**
 - **Default plain-text total field pre-text**
 - **Default plain-text total field post-text**
 - **Default plain-text average row pre-text**
 - **Default plain-text average row post-text**
 - **Default plain-text average field pre-text**
 - **Default plain-text average field post-text**
 - **Default plain-text format pre-text**
 - **Default plain-text format post-text**
 - **Default plain-text format field pre-text**
 - **Default plain-text format field post-text**

Use the pre-text fields to specify formatting styles (for example, line breaks or separators) or text for headers, rows showing the total or average values, or body cells.

Use the post-text fields to specify formatting styles (for example, line breaks or separators).

Use the **Default text to output for a plain text report with no records** field to change the default text “Nothing to display” with any custom text.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring timeout setting for report generation

To configure the timeout setting for report generation:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Reporting**. In the **Reporting** dialog box, under **Misc**, configure **Report Generation Timeout (in seconds)**.

By default, the timeout period is set to 120 seconds. If the report generation does not complete within this period, the job fails. To change the timeout period, type or select a new value. The valid timeout range is 60 to 600 seconds.

- 3 Click **Apply** to apply the new settings and close the dialog box.

Configuring default mail format type for report

To configure the default mail format for report:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Reporting**.
- 3 In the **Reporting** dialog box, under **Misc**, configure the following setting:
 - **Default Mail Format:** By default, the mail format for report is set to PDF. To change the default format select HTML.
- 4 Click **Apply** to apply the new settings and close the dialog box.

Configuring default settings for Statistics Manager

To configure default settings for Statistics Manager:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, complete one of the following actions:
 - To configure default settings for the NetVault Server, continue to Step 3.
 - To configure default settings for a NetVault Client, click **Choose machine**. In the **Choose machine** drawer, select the client, and click **Go to settings**.
- 3 Under **User Interface**, click **Reporting**. In the **Reporting** dialog box, under **Stats Collection**, configure the following settings.

Table 142. Stats Collection settings

Setting	Description
Statistics gathering window start	By default, the Statistics Manager collects statistics for 24 hours, starting at 00:00:00 and continuing until 23:59:59 hours. To limit statistics collection to certain times of the day, specify the start and end time in these boxes. Specify the time in HH:MM:SS format.
Statistics gathering window end	The maximum duration for a session is 24 hours. It can last a single calendar day or continue to the next day. For example, if you set 10:00:00 as the start time and 7:00:00 as the end time, the session will begin at 10:00 A.M. on the current day and lasts until 7:00 A.M. the next day.
Frequency	The default polling period for Statistics Manager is 10 seconds. To change the setting, type or select a new value. The polling period is specified in number of seconds.
Records per Storage File	Type or select the maximum number of records per file. The Statistics Manager will close the current file and open a new file when this limit is reached. The default value for this setting is 1000 records.

- 4 In the **Reporting** dialog box, under **Stats Provision**, configure the following settings.

Table 143. Stats Provision settings

Setting	Description
Enable stats collection on this machine	<p>Statistics collection is enabled by default on all NetVault machines.</p> <p>To disable this activity on a client, clear this check box.</p> <p>If you disable statistics collection, the reports related to drive performance, event history, media requests, and other data collected by the Statistics Manager may show inaccurate information.</p>
Absent server threshold	<p>Type or select how long the Statistics Manager holds the collected data for the Reports Database Manager. The threshold is specified in number of hours. The default value is 72 hours. If the Reports Database Manager does not poll for data within this time, a warning is logged and the following events occur:</p> <ul style="list-style-type: none"> • The Statistics Manager discards all files that are not transferred to the Reports Database. • The Statistics Manager does not store any files until the Reports Database Manager establishes contact with it.
Transmit Block Size	Type or select the block size for transmitting data collected by the Statistics Manager. The block size is specified in KB. The default value is 10KB.
Minimum Stats Manager Cache	Type or select the minimum cache for Statistics Manager. It is specified in number of records. The default value is 30 records.

- 5 Click **Apply** to apply the new settings and close the dialog box.

Creating a global purge policy for the Reports Database

By default, the Reports Database Manager deletes all records that are older than 31 days. You can override this behavior with a custom age-based or size-based purge policy for the Reports Database.

To create a global purge policy for the Reports Database:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Reporting**.
- 3 In the **Reporting** dialog box, under **Tables**, configure the following settings.

Table 144. Global purge settings for the Reports Database

Setting	Description
Tables default to being part of the global purge group	<p>By default, the global purge policy is applied to all tables in the Reports Database. You can override this behavior for individual tables with table-specific purge policies. The table-specific policies are only applied when you disable the global policy at the database level.</p> <p>To apply the global purge policy to all report tables, ensure that this check box is selected.</p>
Purge check frequency	<p>By default, the Reports Database Manager performs purge check every 24 hours, and deletes all records that match the purge condition.</p> <p>To change the frequency of purge checks, type or select the time interval between two purge events. The purge frequency is specified in number of hours.</p>

Table 144. Global purge settings for the Reports Database

Setting	Description
Global purge method	Select one of the following options: <ul style="list-style-type: none"> • By date: Select this option to delete records that are older than the specified time limit. • By space used: Select this option to delete records when the space usage exceeds the specified size limit.
Global purge space limit	Type or select the maximum file size for the report tables. The file size is specified in MB. The default value is 50MB.
Global purge time limit	Type or select the maximum age for records stored in the Reports Database. The record age is specified in number of days. The default value is 31 days.

- 4 Click **Apply** to apply the new settings and close the dialog box.

Creating table-specific purge policy

By default, the Reports Database Manager deletes all records that are older than 31 days. You can override this behavior with a custom age-based or size-based purge policy for the individual report tables.

NetVault supports table-specific purge policies for the following report tables: driveevents, events, mediacapacities, mediarequests, mediatransfers, driveperformance, and jobfileallies. These tables are used to store statistical data produced by backups, media usage and drive activities.

To create a purge policy for individual report tables:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Reporting**.
- 3 In the **Reporting** dialog box, under **Tables**, configure the following settings.

Table 145. Table-specific purge policy

Setting	Description
Tables default to being part of the global purge group	By default, the global purge policy is applied to all tables in the Reports Database. The table-specific policies are only applied when you disable the global policy at the database level. To disable the global purge policy at the database level, clear this check box.
Table <table name> purge method	Select one of the following options: <ul style="list-style-type: none"> • By space used: Select this option to delete records when the space usage exceeds the specified size limit. • By date: Select this option to delete records that are older than the specified time limit. • Use global policy: Select this option to apply the global purge policy to the table.
Table <table name> purge space limit	Type or select the maximum file size for the report tables. The file size is specified in MB. The default value is 10MB.
Table <table name> purge time limit	Type or select the maximum age for records stored in the Reports Database. The record age is specified in number of days. The default value is 31 days.

- 4 Click **Apply** to apply the new settings and close the dialog box.

Configuring NetVault WebUI default settings

You can configure the default settings for Default Table Pagesize (global default of auto) navigation pane (global default of auto-hide), page or table actions (global default of buttons), Time formatting (global 24-hour), and UI theme (global default of light) in the NetVault WebUI.

To configure default settings for navigation pane in the NetVault WebUI:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **WebUI Defaults**.
- 3 In the **WebUI Defaults** dialog box, under **Navigation Bar**, configure the **Auto- hide** setting.
By default, the navigation pane is set to auto-hide. To remove the auto-hide setting of navigation pane in the NetVault WebUI, clear this check box.
- 4 Click **Apply** to apply the new settings and close the dialog box.

To configure default settings for page or table actions in the NetVault WebUI:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **WebUI Defaults**.
- 3 In the **WebUI Defaults** dialog box, under **Page/Table Actions**, configure the **Show Actions** setting.
By default, the action options are displayed as buttons on bottom of the NetVault WebUI pages. To view the action options as a page-level and table-level context menus, select **Context menus**.
- 4 Click **Apply** to apply the new settings and close the dialog box.

To configure default settings for UI theme in the NetVault WebUI:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **WebUI Defaults**.
- 3 In the **WebUI Defaults** dialog box, under **UI Theme**, configure the **Default Theme** setting.
By default, the color theme for NetVault WebUI is set to Light. To change the default color theme, select **Dark**.
- 4 Click **Apply** to apply the new settings and close the dialog box.

To configure default settings for default time in the NetVault WebUI:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **WebUI Defaults**.
- 3 In the **WebUI Defaults** dialog box, under **Default Time**, configure the **Time Formatting** setting.
By default, the 24-hour format is set for NetVault WebUI. To change the default time, select **12-hour (AM/PM)**.
- 4 Click **Apply** to apply the new settings and close the dialog box.

To configure default settings of table pagesize in the NetVault WebUI:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **WebUI Defaults**.
- 3 In the **WebUI Defaults** dialog box, under **Default Table Pagesize**, configure the **Table Pagesize** settings.

By default, table page size is set to **Auto**. You can manually enter page size number or change the application settings page size to one of the following options:

- **25**: Sets the page size to display 25 records in a table.
- **50**: Sets the page size to display 50 records in a table.
- **100**: Sets the page size to display 100 records in a table.
- **500**: Sets the page size to display 500 records in a table.
- **1000**: Sets the page size to display 1000 records in a table.

4 Click **Apply** to apply the new settings and close the dialog box.

Configuring NetVault to use a specific VSS provider

When performing VSS-based backups, NetVault uses the default VSS provider as determined by VSS. To use a specific VSS provider, you can configure the provider ID in the **vss.cfg** file.

To configure NetVault to use a specific VSS provider:

1 On the NetVault Client on which you want to run the VSS-based backup, open the **vss.cfg** file. This file resides in the **config** subdirectory under the NetVault installation directory.

If the **vss.cfg** file does not exist on the server or client machine, manually create the file in the **config** directory.

2 Add the following section:

```
[Provider]
name = <VSS provider name>
type = <Provider type>
ID = {<Provider ID/GUID>}
Version = <Provider version>
```

Example:

```
[Provider]
name = Microsoft Software Shadow Copy provider 1.0
type = System
ID = {b5946137-7b9f-4925-af80-51abd60b20d5}
Version = 1.0.0.7
```

3 Save and close the file.

i **IMPORTANT:**

- When you specify a provider ID (GUID), NetVault does not perform any checks to determine if the provider is registered, installed, or capable of performing snapshot for a volume.
- If NetVault is unable to properly read or parse the provider ID, the VSS backup continues with the default provider as normally determined by VSS.

Configuring default settings using Txtconfig

The Txtconfig utility provides a text-based interface to customize various settings for the server and client machines. This utility is available on all supported operating systems. The Txtconfig utility resides in the **bin** directory under the NetVault installation directory.

i | **NOTE:** You must be logged-in with Administrator privileges on Windows and root user privileges on Linux and UNIX to use this utility.

Before you start Txtconfig on a Windows-based machine running the French or German version of NetVault, you must set the font type and code page in the command prompt window:

- **Set the font type to Lucida Console:** Open a command prompt window. Click the Control Menu icon in upper-left corner or right-click the title bar, and select **Properties**. In the **Properties** dialog box, click the **Font** tab, and select Lucida Console in the **Font** list.
- **Set the code page to 1252:** In the command prompt window where you want to run Txtconfig, type the following command:

```
chcp 1252
```

These settings are necessary to display the text correctly in the command prompt window.

To configure default settings using Txtconfig:

- 1 Log in with Administrator privileges on Windows and root user privileges on Linux and UNIX.
- 2 In a terminal or command prompt window, type:

```
txtconfig
```

Press **Enter** or **Return**
- 3 The **Machine** page is displayed when you start the Txtconfig utility. You can view the following details on this page: Machine name, machine ID, network names, IP addresses, NetVault version, OS version, and OS platform.
- 4 To view the main menu, press **p**. On the main menu page, the following menu items are available:
 - Machine
 - License
 - Services
 - Plugins
 - Auditing
 - Plugin Options
 - General
 - RAS Device Options
 - Encryption
 - Firewall
 - GUI
 - Job Manager
 - Language
 - Logging Daemon
 - Notification
 - Media Manager
 - Device Scan
 - Network Manager
 - Process Manager
 - Reporting
 - Schedule Manager
 - Script
 - Security
 - Time Sync
 - Diagnostic Tracing
 - Verify
 - Web Service
- 5 To open a configuration page, press the option number for that page.
- 6 To configure the settings, do the following:
 - To change the value to ON or OFF, press the option number for that setting.
 - To select a value from a list, press the option number for the list item, and then press Enter or Return.
 - To provide a numeric, string, or date value, type the value, and then press Enter or Return.
- 7 To save the changes, press **s**.

8 To quit Txtconfig, press q.

Diagnostic tracing

- [About diagnostic tracing](#)
- [Managing trace filters](#)
- [Enabling tracing](#)
- [Downloading trace files](#)
- [Changing the trace directory location](#)
- [Enabling tracing using Txtconfig](#)
- [Disabling tracing](#)
- [Deleting trace session directories](#)

About diagnostic tracing

Diagnostic tracing is used to capture detailed information about error conditions. Quest Technical Support personnel use this information for debugging and troubleshooting purposes.

Tracing is disabled by default. To capture diagnostic information, you must enable tracing on the relevant NetVault machines. You can configure NetVault to start tracing immediately or when the service restarts. You can also specify a time window to automatically start and stop tracing at the specified time.

i | **IMPORTANT:** Diagnostic trace settings should only be configured under the direction of Quest Technical Support.

Diagnostic information is written to trace files. You can generate trace files for all processes or a subset of processes. These files are generated on the machine where the selected processes are running. NetVault uses a Trace Container directory to store all trace files generated during various sessions. By default, the trace container directory is created under the NetVault installation directory. You can modify the default settings to specify a new location for the container directory. When tracing is enabled, a session directory is created under the container directory; all trace files generated during the session reside in this directory. You can use the WebUI to download trace files from various machines to the local machine where the WebUI is running.

Managing trace filters

NetVault 11.2 or later allows you to manage trace filters to NetVault processes (core and plug-in) from the NetVault WebUI. You can change/ assign trace settings for these processes using the option **Manage Trace Filters**.

To manage trace filters:

- 1 In the Navigation pane, click **Change Settings**.
- 2 To manage trace filters:
 - For the NetVault Server: On the **NetVault Server Settings** page, under **System and Security**, click **Diagnostic Tracing**.

- For the NetVault Client: On the **NetVault Server Settings** page, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**. Under **System and Security**, click **Diagnostic Tracing**.
- 3 In the **Diagnostic Trace Settings** page, click **Manage Trace Filters**.
 - 4 In the list of NetVault (core and plug-in) processes, select the process to edit. Click **Edit**.
 - 5 In the **Edit Trace Filter Mapping** dialog box, configure the following settings:

Table 146. Edit Trace Filter Mapping

Setting	Description
Selected Process	Selected NetVault process (core and plug-in) is displayed.
Filter	<p>Click the down arrow and select any of the following filter options:</p> <ul style="list-style-type: none"> • KillAll: When this filter is assigned to a process, the process will not generate any traces. • MessageOnly: If this filter is assigned to a process that process will only produce trace messages from MESSAGE module. All other trace messages will be filtered out. <p>Specify the following trace filter mapping options for the selected process:</p> <ul style="list-style-type: none"> • Enable Filter: Select this option to enable the selected filter option for the selected process. • Force Disk Tracing: Select this option to allow the selected process to trace the disk (even when disk tracing is disabled). • Circular: Select this option to enable circular tracing and type or select the values in the following fields: <ul style="list-style-type: none"> ▪ Lines per trace file: Number of lines per trace file when circular is enabled. When lines in a trace file exceed this threshold, traces are written to next trace file. ▪ Number of trace files to keep: Number of trace files to keep when circular is enabled. When number of trace files exceed this threshold, oldest trace file will be overwritten.

- 6 Click **Update** to apply the new settings and close the dialog box.

Enabling tracing

Diagnostic tracing is disabled by default. When investigating a problem, Quest Technical Support might ask you to enable tracing on the relevant NetVault machines to capture diagnostic information about the error condition. This information helps in determining the source of error condition. Once the traces are enabled for the server, warning is displayed as **Diagnostics** in the **Header Pane** of NetVault.

i | **IMPORTANT:** Tracing can affect performance of the application and generate large amounts of data on the server and client machines. You should only enable tracing when requested by Quest Technical Support, and disable this option after the issue is resolved.

To enable diagnostic tracing:

- 1 In the Navigation pane, click **Change Settings**.
—or—
In the Navigation pane, click **Manage Clients**.
- 2 To enable diagnostic tracing:

- For the NetVault Server from **Change Settings**: On the **NetVault Server Settings** page, under **System and Security**, click **Diagnostic Tracing**. On the **Diagnostic Trace Settings** page, click **Enable Tracing**.
 - For the NetVault Client from **Change Settings**: On the **NetVault Server Settings** page, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**. Under **System and Security**, click **Diagnostic Tracing**. On the **Diagnostic Trace Settings** page, click **Enable Tracing**.
 - For the NetVault Client from **Manage Clients**:
 - On the **Manage Clients** page, select the **Trace Status** check box or click **Enable Tracing** —or—
 - On the **Manage Clients** page, select the client, and click **Enable Tracing**.
- 3 In the **Trace Enable Options** drawer, configure the following settings.

Figure 35. Trace Enable Options drawer

Trace Enable Options ⓘ

Enable Tracing

Immediately
 At service startup
 At a certain time

Starting at 10/5/2018 2:00 PM
 Stopping at 10/6/2018 2:00 PM

Create Trace for

All current and future processes
 Specific processes

1	Process Manager
2	Communication Manager
3	Network Manager
9	Stats Manager
15	Console Service
18	WebUI Process Viewer

comma separated list of process IDs

i **IMPORTANT:** When you select multiple clients, the **Trace Enable Options** drawer does not show the list of process names. You can either use the **All current and future processes** option, or provide the process ID list in the associated box.

Table 147. Trace Enable Options

Setting	Description
Enable Tracing	<p>Select one of the following options:</p> <ul style="list-style-type: none"> • Immediately: This option is selected by default. To start tracing immediately, use this option. • At service startup: To enable tracing at service startup, select this option. <p>NOTE: You can also enable tracing at service startup by clicking Trace Configuration on the Diagnostic Trace Settings page, and selecting the Enable trace on service startup check box in the Trace Configuration Options dialog box.</p> <p>When you select this option, tracing is enabled on all processes.</p> <ul style="list-style-type: none"> • At a certain time: To automatically start and stop tracing at specified times, select this option, and specify the time window: <ul style="list-style-type: none"> - Starting at: Select this check box, and type or select the date and time when tracing should be started on the machine. - Stopping at: Select this check box, and type or select the date and time when tracing should be stopped on the machine.
Create Trace for	<p>Specify the processes for which trace files are required:</p> <ul style="list-style-type: none"> • All current and future processes: This option is selected by default. To generate trace files for all current and new processes, use this option. • Specific processes: To generate trace files for specific processes, select this option. <p>In the list of process names, select the processes, and click the Add button (+). To select consecutive items, hold down the Shift key while clicking with the mouse button; to select non-consecutive items, hold down the Ctrl key while clicking with the mouse button. To remove a process that was previously included, select the process name in the list, and click the Remove button (-). Alternatively, edit the Process ID list in the associated box to add or remove any process ID numbers.</p>

i | **IMPORTANT:** Trace settings do not persist across restarts of NetVault Service unless you have selected the **At service startup** or **Enable trace on service startup** options. In such cases, tracing is enabled on all processes.

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

- 4 Click **Apply** to apply the settings and close the drawer.

Downloading trace files

Trace files are generated on the machines where the selected processes are running. You can use the WebUI to download trace files from various machines to the local machine where the WebUI is running.

To download trace files:

- 1 In the Navigation pane, click **Change Settings**.

- 2 To download trace files:
 - From the NetVault Server: On the **NetVault Server Settings** page, under **System and Security**, click **Diagnostic Tracing**.
 - From the NetVault Client: On the **NetVault Server Settings** page, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**. Under **System and Security**, click **Diagnostic Tracing**.
- 3 In the list of trace sessions, select the session that you want to download.
- 4 Click **Download**, and in the confirmation dialog box, click **OK**.

In the File Download dialog box, select the Save option, and specify the location, if necessary. Make sure that you do not navigate away from the page until the browser starts downloading the files.

i | **NOTE:** When you send trace files to Quest Technical Support, make sure that you include files from all relevant NetVault machines.

Changing the trace directory location

By default, the trace container directory is created under the NetVault installation directory. You can modify the default settings to specify a new location for the container directory. When tracing is enabled, a session directory is created under the container directory; all trace files generated during the session reside in this directory.

i | **NOTE:** The trace container directory cannot be changed when tracing is enabled.

To change the trace container directory:

- 1 In the Navigation pane, click **Change Settings**.
- 2 To change the trace directory:
 - For the NetVault Server: On the **NetVault Server Settings** page, under **System and Security**, click **Diagnostic Tracing**.
 - For the NetVault Client: On the **NetVault Server Settings** page, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**. Under **System and Security**, click **Diagnostic Tracing**.
- 3 On the **Diagnostic Trace Settings** page, click **Trace Configuration**. In the **Trace Configuration Options** dialog box, configure the following setting:
 - **Trace Container Directory:** Specify the new location for the trace container directory. You must provide the full filepath.
- 4 Click **Apply** to apply the new settings and close the dialog box.

Enabling tracing using Txtconfig

You can use the Txtconfig utilities to enable tracing at service startup. At service startup, tracing is enabled on all processes. The Txtconfig utilities do not provide the options to start tracing immediately, start and stop tracing at the specified time, or enable tracing for specific processes.

To enable diagnostic tracing using Txtconfig:

- 1 In a terminal or command prompt window, type `txtconfig`, and press **Enter** or **Return**.
- 2 On the **Machine** page, press **p** to view the main menu, and then press the option number for the **Diagnostic Tracing** page.
- 3 Press the option number for the **Enable trace on service startup** setting to change the value to ON.

- 4 To save the settings and quit Txtconfig, press **s**, and then press **q**.

i | **NOTE:** Tracing is not enabled until the next service startup. For more information about restarting the NetVault Service, see [Starting or stopping the NetVault Service](#).

Disabling tracing

Tracing can affect performance and generate large amounts of data on the server and client machines. After your issue is resolved, make sure that tracing is disabled on all relevant NetVault machines.

To disable diagnostic tracing:

- 1 In the Navigation pane, click **Change Settings**.
—or—
In the Navigation pane, click **Manage Clients**.
—or—
In the **Header pane** of NetVault WebUI, click **Diagnostics** warning, the **Manage Clients** page is displayed.
- 2 To disable diagnostic tracing from **Change Settings**:
 - For the NetVault Server: On the **NetVault Server Settings** page, under **System and Security**, click **Diagnostic Tracing**.
 - For the NetVault Client from **Change Settings**: On the **NetVault Server Settings** page, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**. Under **System and Security**, click **Diagnostic Tracing**.
- 3 On the **Diagnostic Trace Settings** page, click **Disable Tracing**.
- 4 To disable diagnostic tracing for the NetVault Client from **Manage Clients** page:
 - Clear the **Trace Status** check box for the selected clients. In the **Confirm Disable Trace** dialog box, click **Ok**.
—or—
▪ Select the client, and click **Disable Tracing**. In the **Confirm Disable Trace** dialog box, click **Ok**.

Deleting trace session directories

If trace files are no longer required, you can delete the session directories to remove all trace files generated during the selected sessions.

To delete the trace session directories:

- 1 In the Navigation pane, click **Change Settings**.
- 2 To delete trace session directories:
 - a From the NetVault Server: On the **NetVault Server Settings** page, under **System and Security**, click **Diagnostic Tracing**.
 - a From a NetVault Client: On the **NetVault Server Settings** page, click **Clients**. In the **NetVault Clients** table, select the client, and click **Next**. Under **System and Security**, click **Diagnostic Tracing**.
- 3 In the list of trace sessions, select the sessions that you want to delete. You can select multiple trace session directories.
- 4 Click **Delete**, and in the confirmation dialog box, click **OK**.

i | **NOTE:** You cannot delete an active trace session directory.

Managing diagnostic data

- [About support diagnostics](#)
- [Downloading diagnostic data](#)
- [Uploading diagnostic data](#)

About support diagnostics

Support Diagnostics provides a functionality that allows a user to collect detailed information about selected machine(s) of NetVault domain. Quest Technical Support personnel use this information for debugging and troubleshooting purpose. This information is collected in a compressed file and can be downloaded on the local machine or uploaded to cloud location against corresponding Service Request (SR).

The following information is collected for the selected machine(s):

- Operating System version of the selected machine(s)
- NetVault version
- NetVault machine ID
- Licensing details
- List of all npk installed
- Contents of the netvault/config directory
- Contents of the netvault/dump directory
- Contents of the netvault/etc directory
- Contents of the netvault/bin directory
- Disk space usage for the machine

Downloading diagnostic data

You can use the NetVault WebUI to download the diagnostic data from various machines to the local machine where the WebUI is running.

To download diagnostic data:

- 1 In the Navigation pane, click **Support Diagnostics**.
- 2 On the **Support Diagnostics** page, type a valid Service Request number in the field **SR Number**. The SR Number must be a seven-digit number. This field is optional.
- 3 In the **Client List** table, you can view all the NetVault clients that are added to the server. The table also displays Virtual Clients.

The table shows the following information about the machines:

- **Status:** Displays the status icons, which indicate the type of client and whether the client is online or offline.
- **Client:** Displays the NetVault name assigned to the client.
- **Version:** Displays the NetVault version installed on the machine.
- **Description:** Displays the client description.

The following table provides a brief description of the client icons.

Table 148. Client icons

Icon	Description
	Client is up and running.
	Client is online. It is in the process of being added, or the NetVault password for the client has changed since it was added.
	Client is unavailable. The system is offline or the NetVault Service is not running.
	Represents a Virtual Client that consists of a cluster of Clients. For more information about Virtual Clients, see Working with client clusters .

- 4 By default, the table is sorted by client name.

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower-right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

- 5 You can use the Search option to filter the table data and quickly find entries that contain the specified string in any column values.
- 6 In the **Client List** table, select the client (s) that you want to collect the diagnostic data, and click **Download**. You can select up to five NetVault clients.

i | **NOTE:** The **Download** button is enabled only when the **SR Number** field is blank or has a valid seven-digit number.

- 7 A confirmation dialog is displayed showing the information regarding non accessible client (s). Click **OK**. Ensure that you do not navigate away from the page until the browser starts downloading the files.

Uploading diagnostic data

You can use the NetVault WebUI to upload the diagnostic data corresponding to a SR number and provide it directly to Quest Technical Support for further analysis.

To upload diagnostic data:

- 1 In the Navigation pane, click **Support Diagnostics**.
- 2 On the **Support Diagnostics** page, type a valid Service Request number in the field **SR Number**. The SR Number must be a seven-digit number. This field is mandatory.
- 3 In the **Client List** table, you can view all the NetVault clients that are added to the server. The table also displays Virtual Clients.

The table shows the following information about the machines:

- **Status:** Displays the status icons, which indicate the type of client and whether the client is online or offline.
- **Client:** Displays the NetVault name assigned to the client.
- **Version:** Displays the NetVault version installed on the machine.
- **Description:** Displays the client description.

The following table provides a brief description of the client icons.

Table 149. Client icons

Icon	Description
	Client is up and running.
	Client is online. It is in the process of being added, or the NetVault password for the client has changed since it was added.
	Client is unavailable. The system is offline or the NetVault Service is not running.
	Represents a Virtual Client that consists of a cluster of Clients. For more information about Virtual Clients, see Working with client clusters .

- 4 By default, the table is sorted by client name.

You can sort the table by one or more columns, as required. For more information, see [Sorting records in NetVault WebUI](#).

To view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower-right corner of the table. For more information, see [Customizing tables in NetVault WebUI](#).

- 5 You can use the Search option to filter the table data and quickly find entries that contain the specified string in any column values.
- 6 In the **Client List** table, select the client (s) that you want to collect the diagnostic data, and click **Upload**. You can select up to five NetVault clients.

i | **NOTE:** The **Upload** button is enabled only when the **SR Number** field has a valid seven-digit number.

- 7 A confirmation dialog is displayed showing the information regarding non accessible client (s). Click **OK**. Ensure that you do not navigate away from the page until the upload process is complete.

Using the deviceconfig utility

- [About deviceconfig](#)
- [Configuring default settings for tape libraries](#)
- [Configuring default settings for tape drives](#)

About deviceconfig

The **deviceconfig** utility is a console application that can be used to configure advanced device settings for all types of tape libraries and drives. This utility is automatically installed on the NetVault Server and Client machines.

The **deviceconfig** utility is located in the “**bin**” directory under the NetVault installation directory.

To use this utility, you must be logged-in with Administrator privileges on Windows machine and root user privileges on Linux and UNIX.

Configuring default settings for tape libraries

To configure default settings for a tape library:

- 1 Start a terminal session or command window, and navigate to the “**bin**” directory under the NetVault installation directory.
- 2 Type the following command:

```
deviceconfig [-servername <FQDN> (-httpport <HTTP port> |
-httpsport <HTTPS port>)] -username <NetVault user>
-password <user password>
-librarymachine <library machine name> -libraryname <library name>
```

These options are described in the following table.

Table 150. Options for modifying tape library settings using deviceconfig

Option	Description
-servername	Specifies the FQDN, or IP Address of the NetVault Server. With this option, you must specify either the HTTP or HTTPS port to connect to the remote Web Service. If you omit this option, the local server is used. The deviceconfig utility connects to the local Web Service and automatically detects the local Web Service configuration.
-httpport	Specifies the HTTP port for Web Service. This option is only required if the server name is supplied.

Table 150. Options for modifying tape library settings using deviceconfig

Option	Description
-httpsport	Specifies the HTTPS port for Web Service. This option is only required if the server name is supplied.
-username	Specifies a valid NetVault user name.
-password	Specifies the password for the user account.
-librarymachine	Specifies the name of the NetVault machine to which the library is attached.
-libraryname	Specifies the name of the library.

Press **Enter**.

- The **deviceconfig** utility displays the **Changer Configuration** screen. You can access the configuration settings by pressing the corresponding option number. The available setting types include the following:
 - Configuration:** To configure general settings for a tape library, select this option. For more information about these settings, see [General settings for tape libraries](#).
 - Cleaning:** To configure drive cleaning settings for a tape library, select this option. For more information about these settings, see [Drive cleaning settings](#).
 - Mixed Media:** To configure mixed media settings, select this option. For more information about these settings, see [Mixed media settings](#).
- To save the changes and quit, press **s**.
(To quit without saving the changes, press **q**.)

Examples

- `deviceconfig -username default -librarymachine Server-A -libraryname MyCustomName`
- `deviceconfig -username admin -password mypassword -librarymachine Server-A -libraryname "Server-A: 2-0.4.0 (SONY LIB-304)"`
- `deviceconfig -servername 10.11.25.125 -httpsport 8443 -username admin -password mypassword -librarymachine Client-A -libraryname MyCustomName-2`

General settings for tape libraries

On the **Configuration** screen, the following settings are available.

Table 151. General settings for tape libraries

Setting	Description
Need command to open entry/exit port	The default setting is OFF . To change it, press the option number.
Do not issue unload commands to drives	The default setting is OFF . To change it, press the option number.
Must unload drive(s) to open door	The default setting is OFF . To change it, press the option number.
Do not overlap commands to arm and drive	The default setting is OFF . To change it, press the option number.

Table 151. General settings for tape libraries

Setting	Description
Return inactive media to slot delay	Type the timeout period for media inactivity. The timeout period is specified in number of seconds. NetVault returns the media to the slot if no activity occurs within the specified period. The default value for this option is 30 seconds. To retain the media indefinitely in the drive, set this option to zero (0). The default value of device config in ACSLS libraries is 120 seconds.
Mark Unknown Media Blank	The default setting is OFF . To change it, press the option number. When it is set ON , NetVault marks unknown tapes as BLANK without reading the tape headers. You still need to run the Blank command to actually delete the data and use the media for backups.
Do Not Scan Unknown Media	When you add tapes to a library, and close the door (or restart the library), NetVault takes an inventory by reading the media barcodes. If NetVault cannot find the barcode for a piece of media in the Media Database, it marks that tape as UNKNOWN. You can use the Mark Unknown Media Blank option to mark such tapes as BLANK. If you do not mark the unknown tapes as BLANK, NetVault starts loading each unknown tape into one of the drives to scan for backups and on-tape indices. On large systems, this process can increase the burden on the drive resources. If you do not want to scan the unknown tapes automatically, change the setting to ON . When it is set ON , the unknown tapes are marked as UNKNOWN.
Check if Offline Media is Available	Enter or select the amount of time. NetVault waits for an offline media to become online. The request times out if the media does not become available within the specified period. The timeout for the availability of offline media is specified in minutes. The default value for this parameter is 10 minutes.
Entry / Exit Port Locks Arm	The default setting is OFF . To change it, press the option number.

Drive cleaning settings

On the **Cleaning** screen, the following settings are available.

Table 152. Drive cleaning settings

Setting	Description
Automatic Cleaning Supported	The default setting is ON for libraries that support automatic cleaning.
Barcode Prefix(es) of cleaning media	Specify the barcode prefixes for cleaning media. To configure multiple cleaning tapes, separate the barcodes using a comma. If you do not use the barcode labels of tape media, you can omit this option. Note: The barcode prefixes (if applicable) and slot numbers are used to identify and place the cleaning media in the reserved slots. To set the “cleaning lives” option, the cleaning media should be placed in the designated slots.
Cleaning Slots	Specify the slot numbers that can hold the cleaning media. To configure multiple slots, use a comma.

Mixed media settings

On the **Mixed Media** screen, the following settings are available.

Table 153. Mixed Media settings

Setting	Description
Slot Types	<p>This option specifies the slot ranges for each media type. The media organization in the library during initial setup determines how you configure this setting. To create this list, assign a unique Slot Type Identifier for each media type and specify the corresponding slot range for it. The format for creating is:</p> <pre><SlotTypeIdentifier>=<SlotRange></pre> <p>To specify more than one slot range for a media type, create a separate list item. Comma-separated values are not supported. Use the same Slot Type Identifier for configuring the additional slots or slot ranges. When assigning a Slot Type Identifier, verify that it allows you to easily identify the media type contained in the slot. No spaces are allowed in the values.</p> <p>Example:</p> <pre>SDLTMedia=1-10 STKRMedia=11-30 LTO1Media=31-60 LTO1Media=101</pre>
Drive Types	<p>This option specifies the types of drives that are available on the library. To create this list, assign a unique Drive Type Identifier for each media type and specify the corresponding drive number for it. The format for creating this list is:</p> <pre><DriveTypeIdentifier>=<DriveNumber></pre> <p>For multiple drives of the same type, configure each individual drive as a separate list item. Comma-separated values are not supported. Use the same Drive Type Identifier for configuring multiple drives of the same type. When assigning a Drive Type Identifier, make sure that it allows you to easily identify the drive type. No spaces are allowed in the values.</p> <p>Example:</p> <pre>SDLT=1 STKR=2 STKR=3 LTO1=4</pre>
Drive Types/Media Type Compatibilities	<p>This option specifies the supported media type for each drive type. The format for creating this list is:</p> <pre><DriveTypeIdentifier>=<SlotTypeIdentifier></pre> <p>Example:</p> <pre>SDLT=SDLTMedia STKR=STKRMedia LTO1=LTO1Media</pre> <p>This configuration ensures that a piece of media is only obtained from the corresponding slots when a particular type of drive is accessed.</p>

Configuring default settings for tape drives

To configure default settings for a tape drive:

- 1 Start a terminal session or command window, and navigate to the “bin” directory under the NetVault installation directory.
- 2 Type the following command:

```
deviceconfig [-servername <FQDN> (-httpport <HTTP port> |  
-httpsport <HTTPS port>)] -username <NetVault user>  
-password <user password>  
-drivemachine <drive machine name> -drivepath <path to drive>
```

These options are described in the following table.

Table 154. deviceconfig options for modifying tape drive settings

Option	Description
-servername	Specifies the FQDN, or IP Address of the NetVault Server. With this option, you must specify either the HTTP or HTTPS port to connect to the remote Web Service. If you omit this option, the local server is used. The deviceconfig utility connects to the local Web Service and automatically detects the local Web Service configuration.
-httpport	Specifies the HTTP port for Web Service. This option is only required if the server name is supplied.
-httpsport	Specifies the HTTPS port for Web Service. This option is only required if the server name is supplied.
-username	Specifies a valid NetVault user name.
-password	Specifies the password for the user account.
-drivemachine	Specifies the name of the NetVault machine to which the drive is attached.
-drivepath	Specifies the name or path of the drive.

- 3 The **deviceconfig** utility displays the **Drive Configuration** screen. You can access the configuration settings by pressing the corresponding option number. The available setting types include the following:
 - **NDMP Configuration:** To configure NDMP Settings, select this option. For more information about these settings, see [NDMP settings](#).
 - **Configuration:** To configure general settings for a tape drive, select this option. For more information about these settings, see [General settings for tape drives](#).
 - **Software Compression:** To configure software compression settings, select this option. For more information about these settings, see [Software compression settings](#).
 - **Performance:** To configure drive performance settings, select this option. For more information about these settings, see [Drive performance settings](#).
 - **Statistics:** To configure statistics collection settings, select this option. For more information about these settings, see [Statistics collection settings](#).
 - **Generic Cleaning:** To configure generic cleaning settings for a tape drive, select this option. For more information about these settings, see [Generic cleaning settings](#).
- 4 To save the changes and quit, press **s**.
(To quit without saving the changes, press **q**.)

Examples

- `deviceconfig -username default -drivemachine Server-A -drivepath "Tape fas3020:nrst2a(QUANTUM SDLT320)"`
- `deviceconfig -username admin -password mypassword -drivemachine Server-A -drivepath "2-0.2.0 (SONY SDX-500C)"`
- `deviceconfig -servername 10.11.25.125 -httpsport 8443 -username admin -password mypassword -drivemachine Client-A -drivepath "2-0.2.0 (SONY SDX-500C)"`

NDMP settings

On the **NDMP Configuration** screen, the following settings are available.

Table 155. NDMP settings

Setting	Description
Allow Mover to do Local Data Transfer if Possible	The default setting is OFF . To change it, press the option number.
Allow Mover to do IPC Data Transfers if Possible	The default setting is OFF . To change it, press the option number.
Allow Mover to do TCP Data Transfer if Possible	The default setting is ON . To change it, press the option number.
Allow Mover to do Direct Backup if Possible	The default setting is ON . To change it, press the option number.
Allow Mover to do Direct Restore if Possible	The default setting is ON . To change it, press the option number.
Emulate NDMP Device	The default setting is ON . To change it, press the option number.

General settings for tape drives

On the **Configuration** screen, the following settings are available.

Table 156. General settings for tape drives

Setting	Description
Device Serial Number	Displays the drive serial number.
End of media warning	This option specifies the amount of media reserved at the end of the tape at which the "end of media" warnings are issued. This value is specified in MB. The default value is 0MB.
Time between polling empty drive	This option specifies the interval at which NetVault polls a standalone drive to detect a tape in the device. This value is specified in number of minutes. The default value is one (1). To turn off polling, set this option to zero (0).

Table 156. General settings for tape drives

Setting	Description
Media block size (KiB)	<p>This option specifies the block size used for read and write operations. The default value is 64KiB.</p> <p>You can change the media block size in increments of 1KiB, but many devices may only accept a value in the multiples of 4KiB or 32KiB.</p> <p>NOTE: The changes to the media block size settings are only applied to blank media items. If you are reusing a media item, blank it first for these changes to take effect.</p> <p>Increasing the block size can reduce the number of times a backup needs to read data and write it to media. However, large media block sizes do not always imply an overall faster backup. The maximum block size is limited by several factors, such as the OS, SCSI adapter, drive make, drive model, and drive type.</p> <p>On Linux and UNIX systems, you can increase the media block size for optimum performance.</p> <p>On Windows, you might be required to change the registry setting MaximumSGList to use block sizes larger than 64KB. Before changing this setting, check that the SCSI bus is only used by the tape devices. If other devices also use the SCSI bus, this registry change might prevent them from working. If you want to apply these changes only to a specific channel on the HBA, consult the hardware vendor.</p> <p>To change the registry setting on Windows, follow these steps:</p> <ol style="list-style-type: none"> 1 Start the Registry Editor. 2 Open the key <code>[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\<HBA Vendor>\Parameters</code> (where <HBA Name> is specific to your SCSI card — for example, QL2200 for a Qlogic 2200 card). 3 Create the Parameters key if not present. 4 Under Parameters, create the Device key if not present. 5 Under the Device key, add the DWORD registry value MaximumSGList if not present. 6 Calculate the hexadecimal value of MaximumSGList: <ul style="list-style-type: none"> On 32-bit systems: $\text{MaximumSGList} = (\text{Maximum Block Size} / 4\text{KiB}) + 1$ <p>For example, if the block size is set to 256KiB, the value for this key is: $(256\text{KiB}/4\text{KiB}) + 1 = 65$</p> <p>The decimal value is 65 and the hexadecimal value is 0x41.</p> <p>You can set the block size to any value from 64KiB through 1012KiB. The maximum value 255 is internally converted to 257 to make a block size of 1 MiB (1024 KiB).</p> On 64-bit systems: <p>On 64-bit systems, the default OS page size is 8KiB. The formula for calculating MaximumSGList is:</p> $\text{MaximumSGList} = (\text{Maximum Block Size} / 8\text{KiB}) + 1$ <p>Thus, the maximum value of 255 corresponds to a maximum media block size of 2MiB.</p> 7 Reboot the system to apply the changes.
Time to wait for plugin to connect	<p>This option specifies the timeout period for the plug-in to connect to NetVault. The job is aborted if connection cannot be established within the specified interval. This value is specified in number of seconds. The default value is zero (0). With the default setting, the job is not timed out.</p>

Table 156. General settings for tape drives

Setting	Description
Supports Short Reads	The default setting is ON . To change it, press the option number.
Cleaning Tapes Supported	This option indicates whether the library supports cleaning tapes or not. The default value is usually correct, unless certain library models have a different setting.
Compression	This option allows you to use a tape drive's built-in compression feature. Not all tape drives support hardware-based data compression. The compression ratio varies depending on the type of data and the compression algorithm that the drive uses. Consult the hardware documentation to determine the compression ratios for the drive.

Software compression settings

On the **Software Compression** screen, the following settings are available. The software compression settings only apply to virtual tape drives.

Table 157. Software Compression settings

Setting	Description
Perform software data compression	The default setting is OFF . To perform software compression, change the setting to ON . The data is compressed when it is transferred to the device during backup.
Compression decision switch	<p>The value set for this option determines the minimum level of compression that must be achieved when data is compressed during a backup. For example, if you set the value to 80 percent, one of the following occurs:</p> <ul style="list-style-type: none"> • If the compressed data size is less than 80 percent of the original data size, the data is backed up in its compressed form. • If the compressed data size is more than 80 percent of the original data size, the data is backed up in its uncompressed form. <p>If you specify 80 percent, a file size of a 100MB must be <= 80MB after compression however, the file size can be >80MB (like 81MB, 82MB, etc.), as the compression is done at block level.</p> <p>When the specified compression level is not achieved for any block, in that case NetVault backs up that block in its uncompressed form. Where, some blocks may be compressed and some blocks may remain in uncompressed form. Hence, in some case, the file size after compression can be greater than 80 percent.</p> <p>The extent that data can be compressed depends on the data contents. Also the encrypted data cannot be compressed. With some files, compression may actually result in a file that is larger than the original uncompressed file.</p>
Number of data blocks per compression unit	Type the number of data blocks per compression unit. The default block size is 8KiB.

Drive performance settings

On the **Performance** screen, the following settings are available.

Table 158. Drive performance settings

Setting	Description
Open disk media for Synchronous IO on Unix	<p>The default setting is OFF. If your RAID system can respond instantaneously to synchronous IO writes, change the setting to ON. The ON setting allows multiple streams to generate output at similar rates resulting in similar jobs ending at the same time (but at the expense of overall throughput).</p> <p>Under most circumstances, it is best to leave this option at the default OFF state to achieve best overall performance.</p>
Amount of memory to assign for transfer buffers	<p>The transfer buffer or the shared memory is allocated in blocks of 32KiB. The default value is 8193KiB.</p> <p>Increasing the transfer buffer size can improve backup performance. To calculate the buffer size, use the following formula:</p> $(<\text{Total number of buffers}> \times 32\text{KiB}) + 1 \text{ byte}$ <p>On Linux and UNIX systems, you require sufficient RAM and large Shared Memory segment. Before increasing the transfer buffer size, check the following settings on these platforms:</p> <ul style="list-style-type: none"> • Maximum size of a shared memory segment (SHMMAX) • Minimum size of shared memory segment (SHMMIN) • Maximum number of shared memory identifiers in the system (SHMMNI) • Maximum number of shared memory segments a user process can attach (SHMSEG) • Maximum number of semaphore identifiers in the system (SEMMNI) • Maximum number of semaphores in a set (SEMMSL) • Maximum number of semaphores in the system (SEMMNS) • Maximum number of operations per semop call (SEMOPM) • Semaphore maximum value (SEMVMX) <p>The total allowed shared memory is determined by the formula $\text{SHMMAX} * \text{SHMSEG}$. These values are often limited by the ulimit setting, and the command <code>ulimit -a</code> can be used to view these system settings.</p> <p>On Windows, you require at least 2GB RAM and large virtual memory. You might also have to change the MaximumSGlist setting on the SCSI card.</p> <p>For examples, see the Optimal transfer buffer size.</p>
Number of media blocks to write at a time	<p>We recommend that you do not change the default setting for this option.</p> <p>If you change the value, record it because it is necessary to re-apply the setting if the drive is reinstalled. Restores require the same values that are set at the time of backup and fail if the settings do not match.</p>
Number of media blocks to read at a time	<p>We recommend that you do not change the default setting for this option.</p> <p>If you change the value, record it because it is necessary to re-apply the setting if the drive is reinstalled. Restores require the same values that are set at the time of backup and fail if the settings do not match.</p>
Lock transfer buffer in memory	<p>Use this option to lock the allocated transfer buffer position in memory, and thus increase the potential performance and prevent other processes from using it when NetVault is running.</p> <p>The default setting is ON. Do not change the default setting unless otherwise advised by Quest Technical Support.</p>

Statistics collection settings

On the **Statistics** screen, the following settings are available.

Table 159. Statistics collection settings

Setting	Description
Gather statistics on device performance	The default setting is OFF . To log drive performance statistics, change the setting to ON . This information facilitates troubleshooting. However, it also increases the size of the NetVault Database.
Gather statistics on data channel performance	The default setting is OFF . To log Data Channel statistics, change the setting to ON . This information facilitates troubleshooting. However, it also increases the size of the NetVault Database.
Record drive performance	The default setting is ON . To record drive performance details with each job that uses the drive, use the default setting.
Length of time between sampling drive transfer rates	Specify the interval at which NetVault records the drive transfer rates. The value is specified in number of seconds. The default value is 60 seconds.

i **NOTE:** To apply any changes to these settings, do the following:

- From the NetVault WebUI (**Tape Drive Management** page), restart the Device Manager process (nvdevmgr) associated with the device by setting the device offline, and then back online. For more information about this process, see [Changing the status of a tape drive](#).
— or —
- Restart the NetVault Services on the relevant machine.

Generic cleaning settings

On the **Generic Cleaning** screen, the following settings are available.

Table 160. Generic Cleaning settings

Setting	Description
Length of cleaning cycle	Specify the duration of the cleaning cycle. The value is specified in number of seconds. The default value is 350 seconds,
Issue a load command to start cleaning cycle	The default setting is ON . If the drive requires a load command to initiate a cleaning cycle, use the default setting.
List of ASC/ASQ codes that indicate cleaning has completed	Specify the ASC/ASCQ SCSI codes for cleaning, if necessary.
Use Generic Cleaning	The default setting is OFF . To use generic cleaning properties, change the setting to ON .

NetVault processes

- [About NetVault processes](#)
- [Process description](#)

About NetVault processes

NetVault includes several static and dynamic processes that run on the server and client machines.

The static processes remain active while the NetVault Service is running. These processes are assigned fixed single-digit process ID numbers and generally use the same amount of system resources during their life.

The dynamic processes are initiated and destroyed according to the current NetVault activities. These processes are assigned changing ID numbers and use varying amounts of system resources during their life.

On Windows-based machines, you can view the NetVault processes from the Task Manager. On Linux- and UNIX-based platforms, you can use the following command to view these processes:

```
ps -ef | grep nv
```

Process description

This section briefly describes the functions of various NetVault processes that run on the server and client machines. NetVault includes the following processes:

- [nvpmgr \(Process Manager\)](#)
- [nvcmgr \(Inter-Process Communications Manager\)](#)
- [nvnmgr \(Network Manager\)](#)
- [nvmedmgr \(Media Manager\)](#)
- [nvsched \(Schedule Manager\)](#)
- [nvlogdaemon \(Logging Daemon\)](#)
- [nvavp \(Audit Verification Manager or Auditor\)](#)
- [nvstatsmgr \(Statistics Manager\)](#)
- [nvrepdbmgr \(Report Manager\)](#)
- [nwsrequesthandler \(Web Service Request Handler\)](#)
- [nvconsolesvc \(Console Service\)](#)
- [nvdeploymentmgr \(Deployment Manager\)](#)
- [nvdevmgr \(Device Manager\)](#)
- [nvnmpdevmgr \(NDMP Device Manager\)](#)
- [nvchmgr \(Changer Manager\)](#)

- [nvndmpchgmgr \(NDMP Changer Manager\)](#)
- [nvrasccontroller \(RAS Device Controller\)](#)
- [nvjobmgr \(Job Manager\)](#)
- [nvduplicate \(Duplicate Process\)](#)
- [nvverify \(Verification Process\)](#)
- [nvplgscript \(Plug-in Script Process\)](#)
- [nvswworker \(Web Service Worker Process\)](#)

nvpmgr (Process Manager)

nvpmgr runs on all NetVault Server and Client machines. This process manages all other NetVault processes. **nvpmgr** creates and destroys the transient processes. The Process Manager also manages the allocation of shared memory area for the process table and trace buffers. Although the Process Manager is assigned a static process ID, this process is seen as a Dynamic process because it requires varying levels of system resources.

Process Type: Dynamic

Process ID: 1

nvcmgr (Inter-Process Communications Manager)

nvcmgr supports the inter-process messaging system. This process runs on all NetVault Server and Client machines. On UNIX and Linux OS, the Communications Manager runs as a process. On Windows, this process runs as a thread within the **nvpmgr** process. **nvcmgr** handles communication between various NetVault processes on a local machine.

Process Type: Static

Process ID: 2

nvnmgr (Network Manager)

nvnmgr supports the inter-process messaging system. This process runs on all NetVault Server and Client machines. On UNIX and Linux OS, the Network Manager runs as a process. On Windows, this process runs as a thread within the **nvpmgr** process. **nvnmgr** transmits the inter-process messages to remote clients. The Network Manager also broadcasts availability messages that help determine the status of the clients.

Process Type: Static

Process ID: 3

nvmedmgr (Media Manager)

nvmedmgr runs on the NetVault Server.

This process manages the Media Database, which contains information about the media contents and online backup savesets. The Media Manager also stores and manages the device configuration details. It manages the backup devices through the Device Manager processes. The Media Manager issues high-level instructions for loading and unloading media; the Device Manager processes carry out these instructions. The Media Manager controls the selection of device and media for a job according to the media requests submitted by the Job Manager.

Process Type: Static

Process ID: 4

nvsched (Schedule Manager)

nvsched runs on the NetVault Server. This process manages the job schedules and queues. It also manages the Scheduler Database. The Schedule Manager initiates the Job Manager to start a job instance, and schedules the next instance for recurring jobs; the Job Manager runs the job. This process also updates the **Job Status** page and provides job scheduling data to the Reporting utility.

Process Type: Static

Process ID: 5

nvlogdaemon (Logging Daemon)

nvlogdaemon runs on the NetVault Server. This process starts along with the NetVault Service.

The Logging Daemon manages the log messages generated by various NetVault processes and writes these messages to the NetVault Database. Log messages contain information that can be used for diagnosing and troubleshooting problems. The Logging Daemon also performs periodic disk space checks, and issues alert messages when the space usage level reaches the Warning or Critical Threshold. These checks apply to the NetVault Home, Database, Logs, and Reports directories.

Process Type: Static

Process ID: 7

nvavp (Audit Verification Manager or Auditor)

nvavp runs on the NetVault Server. This process tracks and controls user activities. **nvavp** validates each user request, and depending on the assigned privileges, allows or denies a request.

Process Type: Static

Process ID: 8

nvstatsmgr (Statistics Manager)

nvstatsmgr runs on all NetVault Server and Client machines. This process collects drive statistics, event history, media requests, server capacity, and transfer information for the reporting utility.

Process Type: Static

Process ID: 9

nvrepdbmgr (Report Manager)

nvrepdbmgr runs on the NetVault Server. This process manages the Reports Database.

The Reports Database Manager polls the Statistics Manager at periodic intervals to fetch the collected data, and writes the data to the Reports Database. **nvrepdbmgr** also transmits the information in the Reports Database to the reporting utility and performs periodic purging of the Reports Database.

Process Type: Static

Process ID: 10

nvwsrequesthandler (Web Service Request Handler)

nvwsrequesthandler runs on the NetVault Server. This process runs the Web Service used by the WebUI.

Process Type: Static

Process ID: 14

nvconsoleSvc (Console Service)

nvconsoleSvc runs on all NetVault Server and Client machines. NetVault uses this process to get or set configuration properties on remote client machines.

Process Type: Static

Process ID: 15

nvdeploymentMgr (Deployment Manager)

nvdeploymentMgr runs on the NetVault Server. This process manages the push installation tasks.

You can use the push installation method to perform the following tasks:

- Install and upgrade the NetVault Client software simultaneously on multiple machines
- Install and upgrade NetVault plug-ins simultaneously on multiple machines
- Add new clients to the NetVault Server

This process stores all deployment related information in the **netvault_machines** database. The Deployment Manager creates one or more deployment worker processes (**nvdeploymentwkr**) to manage the installation processes for the remote machines.

Process Type: Static

Process ID: 17

nvdevMgr (Device Manager)

nvdevMgr runs on the NetVault Server and Clients that have locally attached devices. This process performs media reads and writes, and handles loading and unloading of media. NetVault creates one instance of the Device Manager process for each configured drive. In SAN environment, an instance runs for each NetVault Client that shares the device.

Process Type: Static (runs while the device is available)

Process ID: Varies

nvndmpdevMgr (NDMP Device Manager)

nvndmpdevMgr runs on the NetVault Server. This process performs media reads and writes, and handles loading and unloading of media for NDMP-based filers. NetVault creates one instance of the NDMP Device Manager process for each configured drive.

Process Type: Static (runs while the device is available)

Process ID: Varies

nvchgMgr (Changer Manager)

nvchgMgr controls the Robotic Arm Changer. This process runs on the NetVault Server and Clients to which the robotic arm changer is connected. NetVault creates one instance for each arm changer.

Process Type: Static (runs while the device is available)

Process ID: Varies

nvndmpchgMgr (NDMP Changer Manager)

nvndmpchgMgr runs on the NetVault Server. This process controls the Robotic Arm Changer for devices attached to NDMP-based filers. NetVault creates one instance for each arm changer.

Process Type: Static (runs while the device is available)

Process ID: Varies

nvrasccontroller (RAS Device Controller)

nvrasccontroller manages the disk-based storage devices. This process runs on the NetVault Server. NetVault creates one instance of the RAS Device Controller process for each configured device.

Process Type: Static (runs while the device is available)

Process ID: Varies

nvjobmgr (Job Manager)

nvjobmgr runs on the NetVault Server and manages the execution of a job.

The Schedule Manager starts the Job Manager. A single instance of Job Manager runs for each job until the completion of the job. The Job Manager reports on the job run status and exit status. This process coordinates with the Data Plug-in and fetches the required information from the NetVault Server. It is also responsible for sending drive and media requests to the Media Manager process.

Process Type: Dynamic

Process ID: Varies

nvduplicate (Duplicate Process)

nvduplicate performs backup duplication. This process runs on the NetVault machine on which the Duplication phase is run.

Process Type: Dynamic

Process ID: None

nvverify (Verification Process)

nvverify performs backup verification. This process verifies the stream length written to the media and ensures that no blocks were dropped during backup. **nvverify** runs on the NetVault machine on which the Verification phase is run.

Process Type: Dynamic

Process ID: None

nvplgscript (Plug-in Script Process)

nvplgscript executes the pre- and post-scripts for a job. This process runs on the target client when you select the pre- and post-script execution options for a job.

Process Type: Dynamic

Process ID: None

nvswworker (Web Service Worker Process)

nvswworker is the Web Service worker process. The **nvswrequesthandler** process starts the worker processes. Teams of these processes are used to improve performance. When a large amount of content is served to the WebUI, ten or more processes can run on the NetVault Server. These processes continue running for a short time after the content is served.

Process Type: Dynamic

Process ID: None

Environment variables

- [Environment variables in NetVault](#)

Environment variables in NetVault

The following is a list of NetVault environment variables that can be used in user-defined scripts.

Table 161. Environment variables

Variable	Description
NETVAULTCLIACCOUNT	<p>Specifies the NetVault user name. The specified account must have privileges to use the CLI utility.</p> <pre>NETVAULTCLIACCOUNT=<User Account Name></pre> <p>This variable must be included in the script to access the CLI utility.</p>
NETVAULTCLIPASSWORD	<p>Specifies the password for the NetVault user account.</p> <pre>NETVAULTCLIPASSWORD=<Password></pre> <p>This variable must be included in the script to specify the password for the user account.</p>
NV_HOME	Returns the NetVault installation directory.
NV_JOBCLIENT	<p>Specifies the target client for a job.</p> <pre>NV_JOBCLIENT=<Name of the NetVault Client></pre>
NV_JOBID	<p>Specifies the Job ID number.</p> <pre>NV_JOBID=<Job ID></pre>
NV_JOBTITLE	<p>Specifies the name of the job.</p> <pre>NV_JOBTITLE=<Job title></pre>
NV_JOB_WARNINGS	<p>Returns TRUE if a job completes with warnings, else returns FALSE.</p> <ul style="list-style-type: none"> • If a backup completes with warnings: <pre>NV_JOB_WARNINGS=TRUE</pre> • If a backup completes successfully: <pre>NV_JOB_WARNINGS=FALSE</pre> <p>This variable can only be used in a post -script. The mail scripts use this variable, but the variable has general applicability.</p> <p>If a backup completes with warnings, the <code>NV_STATUS</code> variable returns <code>SUCCEEDED</code>, while the <code>NV_JOB_WARNINGS</code> variable returns <code>TRUE</code>.</p> <p>The return value is not localized; it is <code>TRUE</code> or <code>FALSE</code> in English.</p>
NV_OUTPUT_FILE	Returns the user-defined output file for reports.
NV_SERVERNAME	<p>Specifies the NetVault Server Name.</p> <pre>NV_SERVERNAME=<Name of the NetVault Server></pre>
NV_SESSIONID	<p>Specifies the Session ID of a job.</p> <pre>NV_SESSIONID=<Session ID></pre>

Table 161. Environment variables

Variable	Description
NV_STATUS	Returns the exit status of a job. It returns either SUCCEEDED or FAILED. <ul style="list-style-type: none">• If a backup job completes successfully or completes with warnings: NV_STATUS=SUCCEEDED• If a backup job fails: NV_STATUS=FAILED This variable can only be used in a post-script. The return value is not localized; it is SUCCEEDED or FAILED in English.
NV_USER_ARG	Specifies the user-defined arguments passed with the pre- or post-scripts.

Network ports used by NetVault

- [Ports used or required](#)

Ports used or required

The following table provides a list of network ports used by NetVault.

Table 162. Ports used by NetVault

Port	Protocol	Usage	Comments
80	TCP	HTTP Listen port for incoming Web Service connections.	
135	TCP	RPC port used during the client push installation process.	
3306	TCP	Port used to make a TCP/IP connection to the MySQL Server.	Configured on the NetVault Client on which the Plug-in <i>for MySQL</i> is installed. The default port number is 3306. If a non-default port is configured for client connections on the MySQL Server, verify that the same port is configured on the NetVault Client. To run multiple instances of MySQL on the same machine, a different port is configured for each instance.
5432	TCP	Listener port for PostgreSQL Database.	Configured on the NetVault Client on which the Plug-in <i>for PostgreSQL</i> is installed. The default port number is 5432. If the PostgreSQL Database is configured to listen on a non-default port, verify that the same port is configured on the NetVault Client.
8443	TCP	HTTPS Listen port for incoming Web Service connections.	
10000	TCP	Port for sending NDMP messages (NDMP control channels).	Configured on the NetVault Server on which the Plug-in <i>for NDMP</i> is installed. The default port number is 10000. You can change it, if necessary.
20031	UDP	UDP messaging.	Required on server and clients.
20031	TCP	Port used by Network Manager (nvnmgr).	

Table 162. Ports used by NetVault

Port	Protocol	Usage	Comments
User-defined listen ports for devices	TCP	Ports to listen on for device requests.	Configured on the NetVault Server and SmartClients. Requirement: Two ports per drive.
User-defined connect ports for devices	TCP	Ports to connect to remote storage devices.	Configured on clients that connect to remote storage devices. Requirement: Two ports per drive
User-defined listen ports for message channels	TCP	Ports to receive messages during data transfers.	Configured on NetVault Server and Clients. Requirement: Three ports per client. To run two or more plug-ins simultaneously on a client, NetVault requires two ports per plug-in and an extra port per client. For example, to run two plug-ins simultaneously, NetVault requires $(2 * 2) + 1 = 5$ ports for a client.
User-defined connect ports for message channels	TCP	Ports for sending messages during data transfers.	Configured on NetVault Server and Clients. Requirement: Three ports per client. To run two or more plug-ins simultaneously on a client, NetVault requires two ports per plug-in and an extra port per client. For example, to run two plug-ins simultaneously, NetVault requires $(2 * 2) + 1 = 5$ ports for a client.
User-defined connect ports for inter-machine setup	TCP	Ports for establishing initial contact (broadcast channels) while adding a NetVault Client, and later to ascertain its availability.	Requirement: Two ports per client.
User-defined listen ports for NDMP data channels	TCP	Ports to listen on for NetVault devices operating as NDMP movers.	These ports are used for data transfers between NDMP filer and storage device. These ports are required on the NetVault Server or Client to which the device is attached.
49152 through 65535	TCP	Dynamic ports that are used during the client push installation process.	

Troubleshooting

- [Common errors](#)
- [Safe Mode in NetVault](#)

Common errors

This section describes some common errors and their solutions. It includes the following topics:

- [NetVault Service fails to start on Windows](#)
- [NetVault Service fails to start after the machine is restarted](#)
- [NetVault Service starts, but stops immediately on Linux](#)
- [Login fails after any change in the server IP address](#)
- [Unexpected behavior of NetVault WebUI](#)
- [WebUI does not run in Compatibility View in Internet Explorer](#)
- [NetVault installer fails during push installation](#)
- [VSS-based backup fails](#)
- [Modifying TCP/IP socket buffer size on Windows](#)
- [Restores using Data Copy savesets fail on clients running NetVault 10.0.1](#)
- [Restore fails on Itanium platforms if the index is larger than 2GB](#)
- [After upgrade, Data Copy and Consolidated backup job on Linux fails](#)
- [After upgrade, console error is displayed on WebUI pages](#)
- [Deployment task hangs on target Linux machine during push installation.](#)
- [Unable to add package store with hostname.](#)
- [Deployment task fails due to network configuration issues.](#)
- [Domain user is unable to login NetVault Server if the workstation attribute is set.](#)
- [Domain user is unable to login NetVault Server on Debian 9.](#)
- [Adding the target machine as a client fails, after successful push installation.](#)
- [Unable to install, uninstall or navigate catalog search page after manually uninstalling NetVault Client Host.](#)
- [Unable to install, uninstall catalog search on client after NetVault Server migration with the same or different server name](#)
- [External Azure AD user cannot add an external Azure AD user to NetVault Server](#)
- [Failed to verify target Windows machine from a Linux-based NetVault Server](#)
- [NetVault is unable to send reports as an email attachment in PDF format on RHEL 5.x platform](#)
- [Restore fails on NetVault Database backup](#)
- [When using RDA for backups, only four streams are used at once](#)

- [Unable to create large VTL on Linux](#)
- [Browsing a folder with a large number of files times out](#)

NetVault Service fails to start on Windows

Description

The NetVault Service fails to start on a Windows-based NetVault Server.

Symptom

Check the Windows Event Viewer to see if it displays the following message:

PDT FATAL: lock file "postmaster.pid" already exists

Solution

NetVault 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 Server.

NetVault Service fails to start after the machine is restarted

Description

After restarting the machine, the NetVault Service sometimes fails to start on a Windows-based NetVault Server.

Symptom

Check the Windows Event Viewer to see if it displays the following message:

FATAL: could not create any TCP/IP sockets " for a PostgreSQL source

Solution

NetVault 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. Select any one instance of this process, and click **End Process** to remove all instances of **postgres32.exe**. Then, start the NetVault Service.

NetVault Service starts, but stops immediately on Linux

Description

On a Linux-based machine, the NetVault Service starts, and then stops immediately.

Symptom

No error messages are displayed.

Solution

This issue can occur if the Postgres service cannot resolve the host name **localhost**, and fails to start.

Check the `/etc/hosts` file, and if the file does not contain an entry for **localhost**, add the entry.

Login fails after any change in the server IP address

Description

Login from WebUI fails after any change in the NetVault Server IP address.

Symptom

When you try to log in, the NetVault WebUI displays a message that the server is not accessible.

Solution

After any change in the IP address of the NetVault Server (for example, due to reassignment at reboot by DHCP), you must clear the browser cache before logging in to the NetVault WebUI. Otherwise, the login may fail with a message that the server is not accessible. Alternatively, you can assign a static IP address to the NetVault Server.

Unexpected behavior of NetVault WebUI

Description

NetVault WebUI displays improper strings or names and other object-related issues on any browser.

Symptom

NetVault WebUI behaves unexpectedly displaying various improper names, strings, incorrect button labels, and so on.

Solution

To resolve the issue, clear the browser cache and history and then login again to the NetVault WebUI. Perform this activity after you upgrade your NetVault system.

WebUI does not run in Compatibility View in Internet Explorer

Description

The NetVault WebUI does not run in Compatibility View in Internet Explorer.

Symptom

If Compatibility View is enabled for the site in Internet Explorer, the following error message is displayed when you try to access NetVault: "Compatibility view is not supported. Use a different browser or remove this site from the list of websites using Compatibility View to properly view the application."

Solution

Remove the site from the list of websites using Compatibility View or use a different browser.

NetVault installer fails during push installation

Description

Push installation fails for a target client.

Symptom

The following error message appears in the task log:

NetVault installer for the core package failed with error 1.

Solution

There are several reasons why an installation can fail. The information in the installation log file can help you diagnose and troubleshoot the issue. You can find the log file in the system temporary directory. (The `TEMP` environment variable for system account determines the location of this directory; this path is typically set to `%windir%\Temp`.)

Depending on the stage at which the installation process fails, the installer may create a file named `netvault_{GUID}_install.log` or `bitrock_installer.log` (or `bitrock_installer_nnn.log`).

If you are unable to resolve the issue, contact Quest Technical Support.

VSS-based backup fails

Description

When performing a VSS-based backup, if the VSS writer fails to generate a snapshot, the job fails.

Symptom

The log messages show the following errors:

- Failed to add extra items to the selection tree.
- Failed to prepare the selection tree and backup options.

Solution

These messages indicate a problem in the VSS subsystem. NetVault cannot complete the backup job if the VSS subsystem is not functioning properly. Before you run the job again, we recommend that you do the following:

- Delete the existing shadow copies. (You can use the **Vssadmin** or **DiskShadow** utilities to delete the shadow copies.)
- Restart the applicable VSS writer service.

Modifying TCP/IP socket buffer size on Windows

Description

Cannot modify TCP/IP socket buffer size on Windows.

Symptom

Slow performance occurs on Windows when data is transferred over a LAN. Changing the window size using any network tool (for example, IPerf) does not change the buffer size.

Solution

On Windows, the registry settings must be modified to increase the TCP/IP buffer size. For more information about these settings, see the KB article <https://support.microsoft.com/en-us/kb/823764> from Microsoft.

Restores using Data Copy savesets fail on clients running NetVault 10.0.1

Description

When restoring a Data Copy saveset on a client running NetVault 10.0.1 or an earlier version, the job fails.

Symptom

The restore job reports the following error:

Failed to get index.

The trace logs contain the following messages:

New style index.

Corrupt index file, version number invalid.

Solution

On NetVault Server 10.0.5 and later, the Plug-in *for Data Copy* creates index version 4, which is incompatible with earlier versions of NetVault. A client running NetVault 10.0.1 or an earlier version cannot read these index files, causing the restore job to fail.

To restore Data Copy savesets generated from NetVault Server 10.0.5 or later, the client must be upgraded.

Restore fails on Itanium platforms if the index is larger than 2GB

Description

On Itanium platforms, restore fails for backups with indexes larger than 2GB in size.

Symptom

On Itanium platforms, when you try to restore a backup with an index larger than 2GB in size, the job fails with the error message "Failed when sorting items to restore."

Solution

To correct this issue, increase the stack size:

- On the standard UNIX platforms, edit the **ulimit** setting from the CLI to increase the stack size.

The following is a list of available options:

- `ulimit -a`
Displays all settings for the environment.
- `ulimit -s`
Displays the current stack size setting.
- `ulimit -s unlimited`
Sets an unlimited stack size.
- `ulimit -s <n>`
Sets stack size to the specified value.
- `man ulimit`
Provides information about the **ulimit** command.

After changing the value, run the `ulimit -a` command to ensure that the setting has been changed.

- On HP-UX Itanium platforms, use the bash command **kmtune** or **kctune** (depending on the operating system) to access to the kernel stack size information. The stack variables are **maxssiz** for 32-bit applications and **maxssiz_64bit** for 64-bit applications.

The following is a list of available options:

- `kmtune -l -q maxssiz`

Displays the setting information for a 32-bit application.

- `kmtune -l -q maxssiz_64bit`

Displays the setting information for a 64-bit application.

- `kmtune -u -s maxssiz=<n>`

Sets the new stack size to the specified value for a 32-bit application.

- `kmtune -u -s maxssiz_64bit =<n>`

Sets the new stack size to the specified value for a 64-bit application.

- `kmtune -u -s maxssiz+<n>`

Increases the stack size by the specified value for a 32-bit application.

- `kmtune -u -s maxssiz_64bit +<n>`

Increases the stack size by the specified value for a 64-bit application.

After upgrade, Data Copy and Consolidated backup job on Linux fails

Description

After upgrading NetVault to version 11.4 or later, Data Copy and Consolidated backup job (Primary job of Windows client with Active Directory) running on NetVault Server on Linux machine fails.

Symptom

After upgrading NetVault from to version 11.4 or later, primary backup index gets corrupted. Hence, rerunning older Data Copy and Consolidated backup job fails. Also, new Data Copy and Consolidated backup job from the older full backup set fails.

Solution

To correct the issue, after upgrading NetVault to version 11.4 or later, rerun the older full backup job, and then rerun the older or new Data Copy and Consolidated backup job.

After upgrade, console error is displayed on WebUI pages

Description

After upgrading NetVault to version 13.0, console error is displayed on various WebUI pages.

Symptom

After upgrading NetVault to version 13.0, console error is displayed on various WebUI pages due to changed 'localStorage' and 'sessionStorage' functions.

Solution

To correct the issue, after upgrading NetVault to version 12.0 or later, clear browser cache and then reload WebUI page.

Deployment task hangs on target Linux machine during push installation.

Description

During push installation of NetVault Client on the target Linux machine the deployment tasks keeps running without installing the package.

Symptom

The task logs displays the following log message:

Request sent to launch installation agent on the remote machine

Solution

Perform the following actions before running the task again on the target Linux machine, as the user configuration is incorrect to get sudoers privilege:

- 1 Login the target machine with the same user that is used for push installation.
- 2 Run the following command to list the files, it should not prompt for password.

```
sudo ls
```

However, if the password is prompted, check configuration of `sudoers` file, and set the **NOPASSWD** flag.

Unable to add package store with hostname.

Description

Adding package store with hostname fails with the following error:

Error: Failed to connect to store: 'Error while accessing package store. Check for mount permission"

Symptom

On Linux-based NetVault Server an error message is displayed to check the mount permission.

Solution

Ping the share directory with the **hostname**, if the share directory does not responds to the ping, you need to check the host configurations and the DNS configuration.

Deployment task fails due to network configuration issues.

Description

Deployment task fails on the remote target machine due to network configuration issues.

Symptom

Task logs displays an error message:

Not able to find any valid IP address for NetVault server. Make sure valid n/w configurations are in place.

Solution

Check the network configuration and then run the command `hostname -i` on the target machine. Check if all the valid IP address are listed and not the `Loopback` address.

Domain user is unable to login NetVault Server if the workstation attribute is set.

Description

Domain user is unable to login NetVault Server with correct credentials.

Symptom-

Following log messages are displayed on Linux and Windows based NetVault Server, respectively:

- Authentication error: Failed to add or update domain user record in NetVault Database as cannot fetch its information from domain controller.
- Authentication error: The username or password entered for domain user is incorrect. Check Domain Controller for more details

Solution

On domain controller, check the workstation attribute of the domain user.

Domain user is unable to login NetVault Server on Debian 9.

Description

Domain user is unable to login NetVault Server on Debian 9.

Symptom

NetVault Server displays following error in the **nvavp-*.log** log file:

Domain user validation failed due to access restrictions

Solution

Even after successful authentication, authorization is denied due to **Group Policy Object Settings**.

Set below parameters to resolve this issue in **/etc/sss/sss.conf** configuration file:

```
access_provider = ad
auth_provider = ad
ad_gpo_access_control = permissive.
```

Adding the target machine as a client fails, after successful push installation.

Description

Push installation succeeds on the target machine, however adding the machine as client fails with warning.

Adding a new client on the target machine fails with a warning message, after successful push installation of a remote client target machine.

Symptom

While adding the machine as a client on the target machine, following warnings are displayed in task logs,

- Could not get NetVault name for <target-machine>
- Cannot proceed adding the machine(s) as client

Solution

NetVault Server is unable to contact NetVault Client on the target machine due to firewall configurations. Configure the firewall on the target machine to add the machine as a client. For more information, see, *Quest NetVault Administrator's Guide*.

Unable to install, uninstall or navigate catalog search page after manually uninstalling NetVault Client Host.

Description

If catalog search is installed on NetVault Client host and this client host is manually uninstalled or not accessible, you are unable to install, uninstall or navigate catalog search page.

Symptom

Cannot navigate catalog search page, or it appears as catalog search is installed on the client host.

Solution

Check if catalog search is not available on the client host. If catalog search is not available, change the parameters and their values in the **catalog.cfg** configuration file. You can now re-install catalog search on the selected NetVault Client host. For more information, see, [Managing catalog search](#).

To change the parameters and their values in catalog.cfg file

1 Open the **catalog.cfg** file to edit, available in **config** directory.

- In the [Catalog] section, change the following:

```
Enabled=false
```

- In the [Catalog:Worker Pool Size] section, change the following:

```
Type= Range
```

```
Value= 5
```

```
Width= 4
```

```
Range= 1 20
```

```
Label= Maximum simultaneously active catalog workers
```

```
Label Id= 14744
```

- In the [Timeouts:Catalog Index Saveset] section, change the following:

```
Type= Range
```

```
Range= 5, 600
```

```
Width= 4
```

```
Label= Polling interval (in seconds) to check for backups to catalog
```

```
Label Id= 14745
```

```
Value= 10
```

- In the [CatalogVersion] section, change the following:

```
NVVersion=
```

```
NVBuildLevel=
```

2 **Save and Close** the **catalog.cfg** file

Unable to install, uninstall catalog search on client after NetVault Server migration with the same or different server name

Description

After performing successful NetVault migration you are unable to install catalog search on selected client. After installing catalog search successfully, catalog search still shows as uninstalled or displays an error "Catalog Search is installed by other NetVault Server".

Symptom

After performing successful NetVault migration you are unable to install catalog search on selected client. After installing catalog search successfully, catalog search still shows as uninstalled or displays an error "Catalog Search is installed by other NetVault Server".

Solution

Manually validate catalog search installation in **<NetVault Server Installation\Quest\NetVault\catalog>** folder, if catalog search is installed, you need to uninstall catalog search from the following NetVault installation directory.

External Azure AD user cannot add an external Azure AD user to NetVault Server

Description

An external Azure AD user is unable to add another external Azure AD user to the NetVault Server.

Symptom

When adding another external Azure AD user, the existing external Azure AD user receives the error "User name not present in domain."

Solution

The external collaboration settings in the Azure portal must have the proper permissions. To allow guest users to add other users, complete the following steps:

- 1 Log in to the Azure portal.
- 2 Click **Azure Active Directory**.
- 3 Navigate to **User settings**.
- 4 Click **Manage external collaboration settings**.
- 5 Set **Guest users permissions are limited to No**.
- 6 Click **Save**.

Failed to verify target Windows machine from a Linux-based NetVault Server

Description

When trying to verify a Windows target machine from a Linux-based NetVault server during the creation of a deployment task, the verification fails, even with valid credentials.

Symptom

While verifying the target machines, the following errors displayed:

- Verification throws the following error message: Could not connect to remote target.
- After a verification failure, proceeding with the client installation results in a failure with the same error.

Solution

This issue occurs when the kerberos on the Linux Server machines is not properly configured. Check `/etc/krb5.conf` for a valid `default_realm`.

This issue also occurs when the kerberos ticket is expired for the user. To renew the ticket, run the following command on the Linux NetVault server:

```
kinit user@domain.com
```

Successful execution of the command should result in successful verification for the specified user.

NetVault is unable to send reports as an email attachment in PDF format on RHEL 5.x platform

Description

The ability to generate a PDF requires components that are not available on RHEL 5.x platforms.

Symptom

When attempting to send a report as a PDF from a RHEL 5.x platform, the PDF does not generate.

Solution

To send the report as an email attachment in HTML format, complete the following steps:

- 1 In the Navigation pane, click **Change Settings**.
- 2 Under **User Interface**, click **Reporting**.
- 3 In the **Reporting** dialog box, under **Misc**, select HTML for the **Default mail format** setting.
- 4 Click **Apply**.

Restore fails on NetVault Database backup

Description

On Windows 2012, the Plug-in *for Databases* fails to restore a NetVault Database backup that is stored in a NetVault SmartDisk device.

Symptom

When restoring from a DAV device, the restore fails, but only when NetVault is run as a service.

Solution

If the NetVault Database backup is stored in NetVault SmartDisk, stop the NetVault Service and run the service in “user space” (that is, `nvpmgr debug`) before starting the restore job.

When using RDA for backups, only four streams are used at once

Description

When using RDA to back up 20 or more clients, in device activity, only four streams are used at once. Updating the setting in the configuration file at /config/dellrda config file from 32 streams to 512 streams and restarting NetVault services, yields the same result.

Symptom

When multiple backup jobs that target the same disk-based storage device are scheduled to start at the same time or within a small window of time, the following issues occur:

- The target device is underused regarding the number of streams.
- Some jobs remain in the “Waiting for media” state even if the device can run parallel jobs.
- The completion time for small jobs increases by a few seconds or minutes.

Solution

These issues occur because the Media Manager uses a “**quiesce time**” setting to prevent assigning too many simultaneous jobs to a disk-based backup device. By default, the **quiesce time** is set to 10 seconds. Thus, after a job is assigned to a disk-based storage device, the Media Manager waits for 10 seconds before assigning the next job to the same device. You can change the default setting for **quiesce time** in the **mediamgr.cfg** file.

To change quiesce time for disk-based storage devices:

- 1 Open the **mediamgr.cfg** file in a text editor.

You can find this file in <NetVault home>\config on Windows and <NetVault home>/config on Linux.

- 2 Change the value for the following entry:

```
[Media Requests:RAS quiesce time]
Type=Range
Range=0,1000
Value=10
```

The default value is 10. To start the jobs at the same time, set the value to 0 (Value=0).

- 3 Save the file.

Note the following:

- This setting applies to all disk-based storage devices (DR Series systems, EMC Data Domain Systems, and NetVault SmartDisk devices).
- For shared NetVault SmartDisk devices, this value cannot be set to 0.

Unable to create large VTL on Linux

Description

When creating a Large_VTL on a Linux machine, the NetVault WebUI takes too long to respond and the session expires.

Symptom

On a Linux NetVault Server, when creating a Large_VTL, if the WebUI times out due to inactivity or the user closes the WebUI while the creation is still in progress, the VTL creation does not occur.

Solution

In the WebUI, add a VTL using the **Re-add previously generated virtual device** option.

i | **NOTE:** This solution works only when the VTL entry is present in the "diskdevice.cfg" file under "libraries" section and when the directory is present in the mentioned path.

Browsing a folder with a large number of files times out

Description

During a restore, browsing and expanding a folder with a large number of files, such as 10,000 or more, fails.

Symptoms

After several minutes of browsing, the restore job times out, the WebUI does not respond, or the error, "The remote machine: [MACHINE_NAME] is unavailable," appears.

Solution 1

Use Catalog Search to find the item that you want to restore.

Solution 2

Manually increase the web services physical client long timeout settings by completing the following steps:

- 1 In the Navigation pane, click **Change Settings**.
- 2 On the NetVault Server Settings page, under **Services**, click **Web Service**.
- 3 In the **Web Service** dialog box, configure the **Physical Client Long Timeout** setting to the maximum amount of 60,000 seconds.

Safe Mode in NetVault

The following error conditions can force NetVault to run in Safe Mode:

- A key process such as Media Manager or Schedule Manager fails to start.
- The PostgreSQL database is unavailable or NetVault is unable to connect to one or more databases.
- The NetVault Server is low on disk space.

When these errors occur, NetVault continues to run in a restricted mode, allowing you to use the WebUI to diagnose and troubleshoot the problems.

This file keeps the cause of NetVault Server to run in Safe Mode.

In safe mode, you can access the following pages:

- **Problem Summary:** This page lists the issues that caused NetVault to enter safe mode. Depending on the error type, the **Fix** option is provided to assist you in resolving the error.
- **View Logs:** You can use this page to examine the log messages, which can contain additional information about the error condition.

You can use the following options that are available on the Safe Mode pages:

- **Use the Fix option to resolve the issue:** The following table describes how you can use the **Fix** option to resolve the error.

Error type	Available options
Key process failures	When NetVault enters safe mode due to this error, no options are provided to resolve the issue.
Database connection failures	<p>The Fix option checks the connection status, and displays the following information:</p> <ul style="list-style-type: none"> • Status of the PostgreSQL Service (whether it is running or not) • Connection status for Logs, Auditor, Scheduler, and Media databases <p>The Fix option also provides a link to restart the PostgreSQL Service.</p>
Low disk space issues	<p>The Fix option shows the disk space usage percentage, and provides the following links to free up disk space:</p> <ul style="list-style-type: none"> • A link to delete diagnostic trace files, if there are any. This option also disables tracing on the server. • A link to delete online backup indexes. You can specify the amount of disk space that you want to reclaim in the Index file cleanup dialog box.

- **Restart NetVault in normal mode:** Once NetVault enters safe mode, it remains in this mode whether the error is resolved or not. To exit safe mode, restart NetVault in normal mode.

You can use the **Restart in normal mode** button or links on the page to restart the application.

- **Search Knowledge Base articles:** You can use the **Knowledge Base** link to search knowledge base articles for solution.
- **Restart NetVault with diagnostic tracing enabled:** You can use the link on the page to restart NetVault in normal mode with diagnostics tracing enabled.

i | NOTE: Diagnostic tracing should only be enabled under the guidance of Quest Technical Support.

- **Contact Quest Technical Support:** If you are unable to resolve the issue, click this link to report the issue.

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.

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.