

Dell EMC NetWorker

Version 19.1.x

Administration Guide

302-005-688

REV 03

August 2019

Copyright © 1990-2019 Dell Inc. or its subsidiaries. All rights reserved.

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

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

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

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

CONTENTS

Figures	17
Tables	19
Preface	23
Chapter 1	Overview 29
	The NetWorker environment..... 30
	NetWorker Components..... 30
	NetWorker services..... 32
	Processes on NetWorker hosts 33
	Stop and start the NMC Server..... 36
	Stop and start a NetWorker Server, Client, or Storage Node..... 38
	NetWorker user interfaces..... 41
	NMC user interface..... 41
	NetWorker Administration window..... 42
	NetWorker client interface..... 42
	NetWorker character-based interface..... 43
	NetWorker command-line interface..... 43
	Introduction to the NetWorker Management Web UI..... 43
Chapter 2	Getting Started 45
	NetWorker Management Console interface..... 46
	Connecting to the Console window..... 46
	Connecting to the Administration window..... 51
	Opening the Administration window..... 51
	Administration window..... 51
	Editing multiple resources..... 53
	Drag-and-drop functionality..... 53
	Multiple library devices and slots..... 55
	Setting user interaction preferences..... 55
	Monitoring NetWorker Server activities in the Administration window..... 55
	Getting started with a new installation..... 66
	Common NetWorker tasks..... 66
Chapter 3	Backup Target 75
	Label templates..... 76
	Using label templates..... 76
	How the NetWorker server uses volume labels..... 76
	Preconfigured label templates..... 77
	Guidelines for completing Label Template attributes..... 77
	Naming label templates..... 79
	Working with label templates..... 80
	Setting up a label template to identify volumes..... 82

Media pools.....	83
Using media pools.....	83
Storage nodes.....	98
Requirements.....	98
Licensing.....	98
Storage node configuration.....	99
Storage Node Options.....	105
Configuring a dedicated storage node.....	106
Troubleshooting storage nodes.....	106
Disk storage devices.....	107
Example environment.....	108
Considerations for Client Direct clients.....	109
Differences between FTDs, AFTDs, and DD Boost devices.....	111
Device target and max sessions default values and ranges.....	113
Advanced file type devices.....	114
DD Boost and Cloud Tier devices.....	128
Libraries and silos.....	130
Overview of tape device storage.....	130
Support for LTO-4 hardware-based encryption.....	130
Linux device considerations.....	130
Solaris device considerations.....	132
HP-UX device considerations.....	132
AIX device considerations.....	135
SCSI and VTL libraries.....	135
Silo libraries.....	173
NDMP libraries.....	182
NetWorker hosts with shared libraries.....	182
Dynamic drive sharing.....	185
File type devices.....	190
FTD capacity issues.....	190
Full FTD prevention.....	191
Stand-alone devices.....	191
Autodetecting and configuring a stand-alone tape drive.....	191
Adding a stand-alone device manually.....	192
Auto Media Management for stand-alone devices.....	193
Mounting or unmounting a volume in a stand-alone tape drive.....	193
Labeling and mounting a volume in one operation (stand-alone tape drive).....	194
Labeling volumes without mounting.....	195
Mounting uninventoried volumes.....	195
Labeling volumes.....	196
Labeling or re-labeling library volumes.....	196
Verifying the label when a volume is unloaded.....	197
Troubleshooting devices and autochangers.....	197
Additional attributes in the Autochanger resource.....	197
Maintenance commands.....	198
Autodetected SCSI jukebox option causes server to stop responding.....	198
Autochanger inventory problems.....	198
Destination component full messages.....	198
Tapes do not fill to capacity.....	199
Tapes get stuck in drive when labeling tapes on Linux Red Hat platform.....	199
Increasing the value of Save Mount Time-out for label operations....	200
Server cannot access autochanger control port.....	200

	Changing the sleep times required for TZ89 drive types.....	201
	Message displayed when CDI enabled on NDMP or file type device..	202
	Verify firmware for switches and routers.....	202
	Commands issued with nsrjb on a multi-NIC host fail.....	202
	SCSI reserve/release with dynamic drive sharing.....	202
	Recovering save sets from a VTL on a different NetWorker server...	203
Chapter 4	Data Protection Policies	205
	Overview of protection policies.....	206
	Designing data protection policies.....	207
	Default data protection policies in NMC's NetWorker Administration window.....	208
	Overview of configuring a new data protection policy.....	209
	NetWorker resource considerations.....	210
	Strategies for traditional backups.....	210
	Strategies for server backup and maintenance.....	236
	Strategies for cloning.....	241
	Policy Notifications.....	263
	Monitoring policy activity.....	263
	Monitoring cloning.....	265
	Policy log files.....	265
	Starting, stopping, and restarting policies.....	267
	Protection period.....	267
	Starting actions in a workflow for an individual client.....	269
	Modifying data protection Policy resources.....	270
	Policies.....	270
	Workflows.....	271
	Protection groups.....	279
	Actions.....	280
	Configuring NSR Protection Policies from nsradmin	283
	Commands.....	284
	Managing NSR Data Protection Policy.....	284
	Managing Workflows.....	285
	Managing Actions.....	286
	Configuring Filters for Clone and Index Actions	288
	Deleting multiple actions from multiple workflows and policies ...	289
	Querying NSR protection policy workflow and NSR protection policy action resources without specifying RAP resource types .	290
	Managing policies from the command prompt.....	290
	Creating Data Protection Policy resources from a command prompt	291
	Creating Action resources from a command prompt.....	292
	Starting, stopping, and restarting workflows from a command prompt.....	293
	Displaying Data Protection Policy resource configurations.....	295
	Identifying clients that missed the workflow schedule.....	298
	Managing policies using the NetWorker Management Web UI.....	299
	Policies, workflows, and actions in NetWorker Management Web UI	299
	Create a policy using the NetWorker Management Web UI.....	299
	Create a workflow using the NetWorker Management Web UI....	301
	Protection groups.....	303
	Create an action using the NetWorker Management Web UI.....	308

	Create a rule using the NetWorker Management Web UI.....	350
	Troubleshooting policies.....	352
Chapter 5	Backup Options	355
	Overview of resources that support backups.....	356
	Save sets.....	356
	The ALL save set.....	357
	Backup levels.....	359
	Comparing backup levels.....	360
	Backup levels and data recovery requirements.....	361
	Backup levels for the online indexes.....	362
	Synthetic full backups.....	362
	Virtual synthetic full backups.....	369
	Backup scheduling.....	372
	Migration of NSR schedule from 8.x to 19.1.1.....	372
	Scheduling backup cycles.....	373
	Considerations for scheduling backups.....	374
	Methods for scheduling backups.....	374
	Managing Schedules using NetWorker Management Web UI.....	383
	Create a schedule using the NetWorker Management Web UI.....	383
	Edit a schedule using the NetWorker Management Web UI.....	385
	Backup retention.....	386
	Methods for setting retention.....	386
	Assigning a retention policy to a Client resource.....	387
	Editing retention for a save set.....	388
	General backup considerations.....	389
	Renamed directories.....	389
	Raw partitions.....	389
	Access control lists.....	390
	Client parallelism and parallel save streams.....	390
	Maximum path and save set length.....	394
	Open files.....	394
	Data deduplication.....	395
	Directives.....	396
	Types of directives.....	396
	Format of directive statements.....	396
	Order of Execution in the Directive.....	401
	Global directives.....	401
	NetWorker User local directives.....	405
	Creating local directives.....	406
Chapter 6	Backing Up Data	409
	Configuring a Client resource for backups on Windows hosts	410
	Windows backup considerations.....	410
	Windows file system backups.....	421
	Windows Bare Metal Recovery.....	427
	Creating a Client resource with the Client Backup Configuration wizard.....	446
	Mapped drives.....	451
	Configuring a Client resource for backups on UNIX hosts	452
	UNIX/Linux backup considerations.....	452
	Creating a Client resource with the Client Backup Configuration wizard.....	453
	Supported save set configurations for UNIX hosts.....	457

Configuring a Client resource for backups on Mac OS X hosts	459
Mac OS X backup considerations.....	459
Creating a Client resource with the Client Backup Configuration wizard.....	460
Assigning directives to Mac OS X clients.....	463
Configuring Open Directory database backups.....	463
Sending client data to AFTD or Data Domain devices only.....	464
Non-ASCII files and directories.....	465
Configuring checkpoint restart backups.....	465
About partial save sets.....	466
Partial saveset cloning and scanning.....	466
Checkpoint restart requirements.....	466
Configuring checkpoint restart.....	467
Restarting checkpoint-enabled backups.....	468
Recovering data from partial save sets.....	469
Probe-based backups.....	470
Encryption.....	471
AES Encryption.....	471
In-flight encryption.....	473
Compression.....	474
Configuring compression for scheduled backups.....	474
Configuring compression for manual backups.....	474
Configuring Client Direct backups.....	475
Requirements for Client Direct backups.....	475
Configuring Client Direct backups.....	476
Backup command customization.....	477
Creating a custom backup script.....	477
Customizing backups with the pre and post commands.....	483
Client resources.....	483
Create a Client resource with the Client Properties dialog box....	484
Editing a Client resource.....	485
Copying a Client resource.....	486
Changing the hostname of a client.....	487
Deleting a Client resource.....	488
Manual backups.....	488
Performing a manual backup on Windows.....	488
Performing a manual backup from the command prompt.....	489
Performing a manual backup on Mac OS X.....	490
Troubleshooting manual backups.....	490
Verifying backup data.....	491
 Chapter 7	
Cloning, Staging, and Archiving	493
Cloning, staging, and archiving.....	494
Benefits of cloning and staging.....	494
Cloning save sets and volumes.....	495
Deciding when to clone.....	495
Clone retention.....	495
Cloning requirements and considerations.....	496
Cloning example.....	498
Cloning with tape devices.....	498
Cloning with file type and AFTD devices.....	500
Cloning with Avamar.....	500
Cloning with Data Domain (DD Boost).....	501
Controlling storage node selection for cloning.....	501
Recover Pipe to Save	505

	Cloning save sets from the command prompt.....	505
	Updates to group, policy, workflow, and action fields in mediadb for a cloned save set	509
	Staging save sets.....	509
	Staging bootstrap backups.....	510
	Creating a staging resource.....	511
	Editing staging configurations.....	513
	Copying a Staging resource.....	514
	Deleting a staging policy.....	514
	Manual staging from the command prompt.....	514
	Common NetWorker staging commands and issues.....	515
	Archiving data.....	516
	Storage of archived data.....	517
	Enabling archiving.....	518
	Archiving data from Windows.....	519
	Archiving data from UNIX.....	519
	Recovering archived data.....	520
	Troubleshooting NetWorker archiving and retrieval.....	522
Chapter 8	Backup Data Management	525
	Overview of backup data management.....	526
	Viewing volume and save set details.....	526
	Viewing disk volume details.....	526
	Viewing tape volume details.....	528
	Viewing save set details for a volume.....	529
	Viewing save set details from a search.....	533
	Managing volumes.....	537
	Changing the volume mode.....	537
	Changing the volume recycle policy.....	538
	Marking a tape volume as full for offsite storage.....	538
	Removing volumes from the media database and online indexes.....	539
	Changing save set status.....	540
	Changing the save set retention time.....	540
	Removing expired save sets.....	541
	Save set management on tape devices.....	542
Chapter 9	Recovery	543
	Recovering data.....	544
	Recovery roadmap.....	544
	Prerequisites for recovering a NetWorker client or storage node.....	545
	Downloading the NetWorker software and documentation.....	545
	NetWorker recovery overview.....	546
	Recovery types.....	546
	Directed recoveries.....	546
	Local recoveries.....	550
	Recover programs.....	550
	Using the NetWorker User program.....	551
	Using the NetWorker Recovery program.....	551
	Using the Recovery Wizard.....	551
	Using the recover command.....	557
	Scanner recovery.....	557
	Recovering the data.....	557
	Determining the volume for recovering cloned data.....	557
	Recovering access control list files.....	558

	Browsable recovery.....	559
	Save set recovery.....	571
	Using the scanner program to recover data.....	576
	VSS File Level Recovery.....	578
	Recovering deduplication data.....	578
	vProxy recovery in NMC.....	578
	Entering management credentials for the Data Domain resource (instant recovery and User mode file-level restore only).....	579
	File-level restore as an Active Directory user.....	580
	Recovering a virtual machine using the NMC Recovery wizard...	582
	NMC function to collect vProxy log bundle information.....	600
	Recovering file system data on Windows.....	601
	Recovering Windows volume mount points.....	601
	Recovering Windows DHCP and WINS databases.....	601
	Recovering DFS.....	602
	Recovering data on OS-X clients.....	603
	Recovering files and directories from the command prompt.....	603
	Recovering files and directories by using the NetWorker Recover GUI.....	604
	Recovering client files on a different NetWorker server.....	610
	Recovering as NMC User with Non-Admin Privileges	612
	Preparing the NMC Server for Recovery.....	612
	Creating an NMC Group for Recovery.....	613
	Creating an NMC User with recovery privileges.....	613
	Creating a NetWorker User Group for recovery.....	614
	Recover the NMC Server database.....	614
	Prepare for an NMC Server recovery.....	615
	Recovering the NMC Server.....	615
Chapter 10	Special recoveries on Windows hosts	619
	Special windows recoveries Restoring a Windows Domain Controller host....	620
	Active Directory restore information.....	620
	Selecting a restore method.....	620
	Performing a non-authoritative AD restore on Windows Server 2008 R2 SP1, Windows Server 2012, 2012 R2, 2016, and 2019.....	621
	Performing an authoritative AD restore on Windows Server 2008 R2 SP1, Windows Server 2012, 2012 R2, 2016, and 2019.....	621
	Performing a Windows BMR to physical or virtual computers.....	622
	Prerequisites to performing a Windows BMR.....	622
	Gathering configuration information required by a Windows BMR....	622
	Obtaining the Windows BMR image.....	622
	Creating a Windows BMR bootable image.....	623
	Perform the BMR.....	624
	Troubleshooting Windows BMR.....	633
	Additional recovery options.....	637
	Restart required after recovery operation.....	638
	Online recovery of Active Directory, DFSR, or Cluster services.....	638
Chapter 11	Reporting NetWorker Datazone Activities	641
	Enterprise data reporting.....	642
	Enabling or disabling the gathering of report data.....	642
	Data retention and expiration policies.....	643

Restricted report views.....	644
Report categories.....	645
Legacy report categories.....	645
Report modes and types.....	646
Preconfigured reports.....	653
Customizing and displaying report output.....	679
Customizing and saving reports.....	682
Sharing a report.....	683
Command line reporting.....	684
Reporting policy status and backup job status.....	685
Policy completion and failure notifications.....	686
Querying the job status.....	687
Reporting recover job status.....	705
Using nsrrecomp.....	706
Checkpoint-enabled backup reporting.....	706
View the policy reports for checkpoint-enabled client backups...	706
Determine the status of a checkpoint-enabled backup.....	707
SNMP traps.....	707
Receive SNMP traps on Linux.....	709
Receive SNMP traps on Windows.....	710
Configuring NetWorker SNMP notifications.....	710
SNMP v2c trap support for NetWorker Modules.....	716
NetWorker Notifications.....	717
Preconfigured notifications.....	717
Customizing notifications.....	722
Creating a custom notification.....	726
Editing a notification.....	726
Copying a notification.....	726
Deleting a custom notification.....	727
Configuring owner notifications.....	727
Logging event notifications.....	728
Breakthrough logging	728
Front-end Capacity Estimation.....	729
Configuring EMC Secure Remote Services (ESRS).....	729
Troubleshooting ESRS.....	731
 Chapter 12	
NetWorker Server Monitoring	733
Enterprise events monitoring.....	734
Polling interval for system events.....	734
Enabling or disabling event capture for a host.....	735
Event viewing.....	735
Dismissing an event.....	737
Monitoring NetWorker Server activities in the Administration window.....	737
About the Monitoring window.....	738
Customizing the Monitoring window.....	740
Policies/Actions pane.....	741
Sessions window.....	744
Alerts pane.....	745
Devices pane.....	745
Operations window.....	746
Log window.....	748
Recover window.....	749
Monitoring changes to the NetWorker and NMC Server resources.....	752
Disabling or enabling the Monitor RAP Attribute.....	753
Monitoring user access to the NMC server.....	753

	Monitoring NetWorker server activities in the log files.....	753
Chapter 13	NMC Server Management	755
	Enterprise.....	756
	Enterprise components.....	756
	Organizing NetWorker servers.....	756
	Viewing the enterprise.....	757
	Managing various servers in the Enterprise.....	757
	Managing folders in the enterprise.....	759
	Adding or deleting multiple servers by using a hostname file.....	761
	Customizing the Console window and views.....	763
	Using the NMC filters.....	764
	Connecting to the NMC GUI using an ssh connection.....	766
	Backing up the NetWorker environment.....	766
	Configuring an NMC server database backup.....	767
	Performing a manual backup of the NMC server database.....	768
	Using the NMC Configuration Wizard.....	768
	NMC server authentication.....	769
	Configuring the NMC server to manage additional NetWorker servers.....	769
	Changing the NetWorker Authentication Service hostname and port number.....	770
	Modifying user groups for new NetWorker Authentication Service users.....	772
	Adding the NMC service account to the Users user group.....	774
	Enabling two factor authentication for AD and LDAP users.....	775
	Moving the NMC Server.....	776
	Migrating NMC users to the authentication service database.....	777
	Updating the NetWorker User Group resources for migrated NMC users.....	778
	Resetting the administrator password.....	779
	Changing the service port used by the NMC database.....	781
	Changing the service port used by the NMC database on Linux... 781	
	Changing the service port used by the NMC database on Windows.. 782	
	Changing database connection credentials.....	783
	Updating the NMC server IP address/hostname.....	783
	Setting system options to improve NMC server performance.....	784
	Individual User Authentication.....	786
	Displaying international fonts in non-US locale environments.....	787
	NetWorker License Manager.....	787
	Entering an enabler code.....	787
	Deleting an enabler code.....	787
	Entering an authorization code.....	788
	Changing the License Manager server.....	788
	NMC error messages and corrective actions.....	788
	Console troubleshooting notes and tips.....	793
	Troubleshooting an NMC server that is not responding.....	793
	Unable to connect to host: Please check Security setting and daemon logs on the NetWorker client and Console server for more information.....	794
	Username/password validation fails when you use the NMC New Device wizard to configure an AFTD if storage node is UNIX.....	795
	Querying large numbers of save sets in the NetWorker user interface may cause a Java heap space error.....	795

	NMC user interface exits unexpectedly.....	796
	Label and Mount devices page is not displayed in NMC device configuration wizard.....	796
Chapter 14	NetWorker Server Management	797
	Setting up the server.....	798
	License the NetWorker Server.....	798
	Setting the Job inactivity timeout.....	798
	Modifying the retention period for jobs in the jobs database.....	798
	Viewing the migration log file.....	799
	Hostname changes.....	799
	Managing the NSR task resource for nsrclientfix	800
	Parallelism and multiplexing.....	801
	Parallelism.....	801
	Multiplexing.....	804
	Managing server access.....	805
	Resource databases	806
	Viewing resources in the resource database.....	806
	Repairing resource database corruption.....	806
	Indexes.....	807
	Characteristics of the online indexes.....	807
	Automated index activities.....	808
	Checking online indexes.....	808
	Viewing information about the indexes.....	808
	Index save sets.....	809
	Querying the media database.....	810
	Cross-checking client file indexes.....	810
	Refreshing index information.....	810
	Client file index locations.....	811
	Managing the size of the online indexes.....	812
	Internationalization.....	816
	Log file viewer.....	816
	Display issues.....	816
	Creating a Server Backup action.....	817
	Creating an expire action.....	820
Chapter 15	NetWorker Host Management	823
	Controlling access to a NetWorker client.....	824
	NetWorker host management.....	824
	Windows client interface.....	826
	Starting the NetWorker User program on Windows.....	827
	Toolbar buttons.....	828
	Browse window.....	828
	Connecting to a NetWorker server.....	829
	Editing a client NSRLA database.....	829
Chapter 16	Restricted Datazones	831
	Restricted Datazones overview.....	832
	Administrators and users of RDZ.....	832
	Using multiple instances of an RDZ.....	833
	Setting up the RDZ.....	833
	Setting up RDZ Users.....	834
	Setting up an RDZ resource.....	835
	Removing a resource association.....	839

	Backward compatibility.....	839
Chapter 17	Block Based Backup and Recovery	841
	Overview.....	842
	Supported operating systems and configurations.....	843
	Limitations.....	845
	Block based backups.....	846
	Devices for block based backups.....	846
	Installing the lgtobbb package on Linux.....	846
	Configuring block based backups.....	847
	Performing block based backups.....	848
	Verifying block based backups.....	851
	Cloning block based backups.....	851
	Block based recoveries.....	851
	Preparing for block based recoveries.....	851
	Performing block based recoveries.....	852
	Performing block based clone recoveries.....	855
	Troubleshooting block based backup and recovery issues.....	858
Chapter 18	Networking and Connectivity	861
	Name resolution and connectivity.....	862
	Troubleshooting name resolution and connectivity errors.....	863
	Verifying basic connectivity.....	864
	Verifying name resolution.....	866
	Verifying the NetWorker configuration.....	870
	Using multihomed systems.....	871
	Multihomed system requirements.....	871
	Configuring multihomed hosts in a datazone.....	872
	NIC Teaming.....	877
	Using DHCP clients.....	878
	NetWorker TCP/IP keep-alive parameters	878
	Modifying the NetWorker TCP/IP parameters in Linux platform	879
	Modifying the NetWorker TCP/IP parameters in Windows platform	
		879
Chapter 19	Cloud Supportability	881
	CloudBoost appliance as the back up target.....	882
	Support for Azure Stack.....	882
	Cloud service provider support matrix for NetWorker.....	883
Chapter 20	Troubleshooting	885
	Before you contact technical support.....	886
	Determining the version of NetWorker software running on a client..	886
	Displaying diagnostic mode attributes.....	887
	NetWorker log files.....	887
	NetWorker Server log files.....	887
	NMC server log files.....	890
	NetWorker Client log files.....	891
	View log files.....	894
	Raw log file management.....	898
	Configuring logging levels.....	901
	NetWorker Authentication Service logs.....	910

NetWorker Authentication Service log files.....	910
NetWorker Authentication Service server log file management....	911
CLI log file management.....	912
NetWorker functionality issues.....	913
Backup and recovery.....	913
Backups fail to start when the daylight savings time change occurs..	916
Shut down NetWorker services prior to any significant changes to system date.....	916
Clone ID timestamp does not reflect the time the clone was created	916
Memory usage when browsing large save sets.....	916
Memory usage and nsrjobd.....	916
Media position errors encountered when auto media verify is enabled.....	917
The scanner program marks a volume read-only.....	917
The scanner program requests an entry for record size.....	917
Limitations for groups containing a bootstrap.....	917
Index recovery to a different location fails.....	918
Illegal characters in configurations.....	918
Inaccessible object exception error when launching NMC with Java 9.....	918
Error backing up large number of clients.....	918
Hostname aliases.....	919
Directory pathname restrictions.....	919
Backup of a new client defaults to level full.....	919
Non-full backup of Solaris files with modified extended attributes....	920
Client file index errors.....	920
Aborting a recovery.....	921
xdr of win32 attributes failed for <i>directory</i>	921
Cannot create directory <i>directory</i>	921
The All save set and duplicate drive serial numbers.....	921
No disk label errors.....	922
Resolving copy violation errors.....	922
Converting sparse files to fully allocated files.....	922
Backing up large sparse files.....	923
Queries using the mminfo -N command are case-sensitive.....	923
Renamed directories and incremental backups.....	923
Resolving names for multiple network interface cards.....	924
Libraries entering ready state.....	924
Successful save sets listed as failed in the Group Backup Details window.....	924
The NetWorker Server window does not appear on HP-UX.....	925
Backup fails with Win32 error 0x2.....	925
Error displaying workflow details.....	925
Back up of All Save Sets takes a long time to complete.....	926
GSS-API authentication error	926
NetWorker locale and code set support.....	926
Enabling service mode for NetWorker.....	927
Resetting the NMC password for NetWorker 19.1 and later.....	927
No privileges to view NetWorker server from NMC.....	927
Network and server communication errors.....	928
Unapproved server error.....	928
Unapproved server error during client setup.....	928
Server copy violation.....	929

Remote recover access rights.....	929
NetWorker server takes a long time to restart.....	929
Changing the NetWorker server address.....	929
Binding to server errors.....	930
Identifying connection lost to device.....	930
Device Connectivity Check Configurations	931
New.Net and NetWorker software are incompatible.....	933

Glossary

935

FIGURES

1	NetWorker components.....	30
2	Stopping the NetWorker Remote Exec Service.....	39
3	NMC GUI window.....	41
4	Welcome to the NMC Server Configuration Wizard page.....	48
5	Set authentication server service account for the NMC Server page.....	49
6	Specify a list of managed NetWorker servers page.....	50
7	Administration window.....	52
8	Monitoring window.....	56
9	Recover window.....	63
10	Labeling a volume by using a label template.....	76
11	Identifying WORM tapes in the NetWorker Console.....	95
12	Example NetWorker disk backup configuration in a mixed backup environment.....	109
13	Paths for CIFS AFTD.....	110
14	How library sharing works.....	183
15	Dynamic Drive Sharing.....	186
16	Data Protection Policy.....	207
17	Platinum policy configuration.....	208
18	Data protection policy example.....	210
19	Workflow path from a traditional backup action.....	217
20	Visual representation of a workflow.....	235
21	Workflow path from a server database backup action.....	239
22	Workflow path from an NMC server backup action.....	239
23	Visual representation of the Server Protection workflows.....	240
24	Replication using AMS.....	241
25	Workflow path from a clone action.....	249
26	Visual representation of a clone workflow.....	255
27	Example of a policy with separate workflows for backup and cloning.....	256
28	Creating a new policy.....	268
29	Policy properties.....	269
30	Workflow path from a snapshot backup action.....	271
31	Workflow path from a probe action.....	271
32	Workflow path from a server backup action.....	272
33	Workflow path from a check connectivity action.....	272
34	Workflow path from a clone action.....	272
35	Workflow path from a discover action.....	272
36	Workflow path from a generate index action.....	272
37	Workflow path from a VBA checkpoint discover action.....	273
38	Traditional backup workflow.....	273
39	Incremental and cumulative incremental backup levels.....	362
40	Synthetic full backups.....	363
41	Default weekly backup schedule.....	373
42	Staggered weekly backup schedule for multiple groups of clients.....	373
43	Selecting the schedule.....	374
44	Defining the schedule.....	376
45	The Force Backup Level attribute.....	378
46	VSS backup process.....	422
47	Paths for CIFS AFTD.....	477
48	Cloning example.....	498
49	Overview of archive operation.....	517
50	Volume Save Sets window.....	532
51	Change Expiration window.....	541
52	Recovery roadmap.....	544
53	A directed recovery from a remote client	547

54	NSR Data Domain Properties.....	579
55	Virtual machine recovery in the NMC Recovery wizard.....	583
56	Select the Virtual Machine to Recover.....	584
57	Select the Target Backup (individual virtual machine).....	585
58	Select the Target Backup (multiple virtual machines).....	585
59	Select the Virtual Machine Recovery method.....	585
60	Choose Disks to Revert.....	587
61	Select Alternate Recovery Sources.....	588
62	Configure the Instant Recovery.....	589
63	Configure the virtual machine recovery.....	591
64	Configure the Virtual Disk Recovery.....	592
65	Configure the Emergency Recovery.....	594
66	Select Alternate Recovery Sources for file level recovery.....	596
67	Mount the save set for file level recovery.....	597
68	Select the files and folders to recover.....	598
69	Connect to Server	605
70	NetWorker Recover window.....	606
71	List of clients available for a NetWorker server.....	607
72	Search browse view.....	608
73	Versions side bar.....	609
74	Group Summary in table view.....	647
75	Group Summary in Bar Chart view.....	648
76	Sample log output.....	714
77	SNMP trap output.....	715
78	ESRS Properties.....	730
79	Monitoring window.....	738
80	Recover window.....	750
81	NetWorker servers worldwide.....	757
82	Using filters to search and view policies.....	764
83	Copying the group DN.....	773
84	Copying the group DN.....	773
85	Add Distinguished Names window.....	775
86	Hosts window.....	825
87	NetWorker User program.....	827
88	Example of the browse window.....	829
89	Restricted Data Zones in NMC.....	834
90	Restricted Datazone User Configuration.....	835
91	Create Restricted Data Zone in the NetWorker Administration Server window.....	836
92	Restricted Data Zone Client Properties	837
93	Restricted Data Zones in Device Properties window.....	837
94	Restricted Data Zones in Create Policy window.....	838
95	New workflow associated with RDZ group.....	838
96	Multihomed environment.....	874
97	Configuring the Aliases attribute for NetWorker Server Client resource.....	875
98	Configuring the Aliases attribute for NetWorker Storage Node Client resource.....	875
99	Storage Nodes attribute for clients in VLAN1.....	876
100	Aliases and Server network interface attributes for VLAN1 clients.....	876
101	Storage Nodes attribute for clients in VLAN2.....	877
102	Aliases and Server network interface attributes for VLAN2 clients.....	877
103	Azure stack backup and disaster recovery.....	882
104	WinPE registry key to troubleshoot recoveries.....	909
105	Settings on Server CLI.....	932
106	Settings on Storage Node CLI.....	932
107	NMC Server settings.....	932
108	NMC Storage Node Settings.....	933

TABLES

1	Revision history.....	23
2	Style conventions.....	25
3	NetWorker Server processes.....	33
4	NetWorker Storage Node processes.....	35
5	NMC Server processes.....	36
6	NetWorker startup commands	40
7	Windows opened from the NMC GUI.....	42
8	Supported operations in the NetWorker Management Web UI.....	43
9	Windows that are launched from the Administration window.....	52
10	Monitoring window panel	57
11	Alerts window icons.....	58
12	Devices status icons	59
13	Operations window icons.....	60
14	Icons in the Log pane.....	62
15	Recovery toolbar options	63
16	Save recover configuration job status.....	64
17	Find options.....	65
18	Key label template attributes.....	77
19	Examples of number sequences for volume labels.....	79
20	Using label template components.....	79
21	Preconfigured media pools.....	83
22	Determining which pool receives backup data.....	86
23	NetWorker hierarchy for resolving media pool conflicts	89
24	WORM supported devices	93
25	WORM/DLTWORM attributes	95
26	Differences between disk devices	111
27	Default values and ranges for target and max sessions attributes	113
28	Determining the major number value.....	133
29	ioscan output when driver is configured.....	133
30	ioscan output when driver is not configured.....	134
31	Tape alert severity.....	162
32	Common jbedit options.....	164
33	Device settings and environment variables	170
34	StorageTek environment variables	177
35	Library resource sleep attributes	183
36	Shared Devices attributes.....	189
37	Schedule icons.....	218
38	Schedule icons.....	221
39	Schedule icons.....	225
40	Schedule icons.....	231
41	nsrcloneconfig file details.....	242
42	Save set criteria.....	245
43	Schedule icons.....	250
44	Schedule icons.....	259
45	Policy status icons.....	264
46	Methods to create an action.....	281
47	Methods to open the Policy Action wizard.....	282
48	Command descriptions.....	284
49	Commands to manage NSR Protection Policy resources.....	285
50	Commands to manage NSR Protection Policy Workflow resources.....	285
51	Commands to manage NSR Protection Policy Action resources.....	286
52	Filter values for clone and index actions.....	289
53	Resource overview.....	356

54	Data in the ALL save set	357
55	File systems excluded from the ALL save set.....	358
56	Special ALL save sets	359
57	Backup levels	360
58	Advantages and disadvantages of backup levels	360
59	mminfo commands for synthetic full backup validation	368
60	Comparison of traditional synthetic full and virtual synthetic full backups.....	369
61	Requirements for virtual synthetic full backups.....	369
62	mminfo commands for VSF backup validation	372
63	Scheduled backup level icons.....	375
64	Preconfigured NetWorker schedules	378
65	Types of schedule.....	381
66	Types of actions.....	382
67	Log files for PSS troubleshooting.....	393
68	Supported wildcards in directives.....	399
69	Preconfigured directives.....	401
70	Backup considerations for Windows features.....	411
71	VSS Save operation attribute values	426
72	DISASTER_RECOVERY:\ components in an incremental backup.....	430
73	Save set configuration for a specific host	438
74	Special ALL save sets	451
75	File systems excluded from the ALL save set.....	457
76	Special ALL save sets	458
77	NetWorker software requirements for checkpoint restart.....	466
78	Example backup script on Windows.....	479
79	NetWorker Server Versions.....	481
80	Job control attribute selections.....	482
81	List of nsrclone options and their descriptions.....	508
82	Staging criteria options.....	511
83	Disk volumes window.....	526
84	Volume details.....	528
85	Save Set details.....	530
86	Query criteria.....	533
87	Save set search results view.....	535
88	Query criteria.....	535
89	VBA save set search results window.....	536
90	General recover requirements	547
91	Volume selection by recovery method.....	557
92	Query criteria.....	560
93	Save set status.....	562
94	Optional browsable recovery options	566
95	Save set information.....	569
96	Optional save set recovery options	572
97	Save set information.....	575
98	Additional recovery options.....	637
99	Data retention policies.....	643
100	Report categories	645
101	Legacy report categories	645
102	Report icons.....	646
103	Report chart formats.....	650
104	NetWorker recovery statistics parameters	667
105	Event parameters	670
106	Host reports	671
107	NetWorker backup statistics parameters	672
108	NetWorker backup status parameters	675
109	Clone report parameters	677

110	Date and time input formats for common locales.....	681
111	Workflow-specific job record attributes.....	691
112	Action job record attributes.....	694
113	Job details for a Workflow	702
114	Job details for a Workflow continued.....	703
115	SNMP attributes and descriptions.....	708
116	Command-line options for nsrtrap	711
117	Preconfigured notifications	717
118	Actions	723
119	Priorities	725
120	Event Viewer messages	728
121	ESRS fields and descriptions.....	730
122	NMC event information.....	735
123	Event priorities	736
124	Monitoring window panel	739
125	Policy status icons.....	741
126	Alerts window icons.....	745
127	Devices status icons	746
128	Operations window icons.....	747
129	Icons in the Log pane.....	748
130	Recovery toolbar options	750
131	Save recover configuration job status.....	751
132	Find options.....	752
133	Viewing the enterprise.....	757
134	NMC windows with filtering capability.....	764
135	NMC server system options	784
136	Error messages or symptoms	788
137	Indexes window information.....	808
138	Index save sets dialog box information.....	809
139	Schedule icons for the expire action	821
140	When to modify the servers file.....	824
141	Summary pane.....	825
142	NetWorker User Groups requirements.....	827
143	NetWorker User toolbar functions	828
144	Supported backup and recovery scenarios.....	843
145	Key options for the block based recover.exe command.....	855
146	Troubleshooting block based backup and recovery issues.....	858
147	Configuring multihomed hosts in NetWorker (continued).....	872
148	TCP/IP parameters.....	878
149	Cloud service provider support matrix.....	883
150	NetWorker Server log files.....	887
151	NMC server log files.....	890
152	Client log files.....	891
153	Message types	896
154	Raw log file attributes that manage log file size.....	898
155	Raw log file attributes that manage the log file trimming mechanism.....	899
156	NetWorker Authentication Service log files.....	910
157	NetWorker Startup commands.....	930
158	Enabling and Disabling DCC based on Server and Storage node settings.....	933

Preface

As part of an effort to improve product lines, periodic revisions of software and hardware are released. Therefore, all versions of the software or hardware currently in use might not support some functions that are described in this document. The product release notes provide the most up-to-date information on product features.

If a product does not function correctly or does not function as described in this document, contact a technical support professional.

 **Note:** This document was accurate at publication time. To ensure that you are using the latest version of this document, go to the Support website <https://www.dell.com/support>.

Purpose

This document describes how to configure and use NetWorker.

Audience

This guide is part of the NetWorker documentation set, and is intended for use by system administrators who are responsible for setting up and maintaining backups on a network. Operators who monitor daily backups will also find this guide useful.

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
03	August 14, 2019	The following topics are added in the guide: • Types of schedule • Migration of NSR schedule from 8.x to 19.1.1 • Resetting the NMC password for NetWorker 19.1 and later The following topics are updated: • Methods for scheduling backups • Defining a schedule for a client • Creating an action • A note is added in the topic “Creating a Server Backup action”
02	June 12, 2019	Technical errors were fixed.
01	May 20, 2019	First release of the document for NetWorker 19.1.

Related documentation

The NetWorker documentation set includes the following publications, available on the Support website:

- *NetWorker E-LAB Navigator*

Provides compatibility information, including specific software and hardware configurations that NetWorker supports. To access E-LAB Navigator, go to <https://elabnavigator.emc.com/eln/elhome>.

- *NetWorker Administration Guide*
Describes how to configure and maintain the NetWorker software.
- *NetWorker Network Data Management Protocol (NDMP) User Guide*
Describes how to use the NetWorker software to provide data protection for NDMP filers.
- *NetWorker Cluster Integration Guide*
Contains information related to configuring NetWorker software on cluster servers and clients.
- *NetWorker Installation Guide*
Provides information on how to install, uninstall, and update the NetWorker software for clients, storage nodes, and servers on all supported operating systems.
- *NetWorker Updating from a Previous Release Guide*
Describes how to update the NetWorker software from a previously installed release.
- *NetWorker Release Notes*
Contains information on new features and changes, fixed problems, known limitations, environment and system requirements for the latest NetWorker software release.
- *NetWorker Command Reference Guide*
Provides reference information for NetWorker commands and options.
- *NetWorker Data Domain Boost Integration Guide*
Provides planning and configuration information on the use of Data Domain devices for data deduplication backup and storage in a NetWorker environment.
- *NetWorker Performance Optimization Planning Guide*
Contains basic performance tuning information for NetWorker.
- *NetWorker Server Disaster Recovery and Availability Best Practices Guide*
Describes how to design, plan for, and perform a step-by-step NetWorker disaster recovery.
- *NetWorker Snapshot Management Integration Guide*
Describes the ability to catalog and manage snapshot copies of production data that are created by using mirror technologies on storage arrays.
- *NetWorker Snapshot Management for NAS Devices Integration Guide*
Describes how to catalog and manage snapshot copies of production data that are created by using replication technologies on NAS devices.
- *NetWorker Security Configuration Guide*
Provides an overview of security configuration settings available in NetWorker, secure deployment, and physical security controls needed to ensure the secure operation of the product.
- *NetWorker VMware Integration Guide*
Provides planning and configuration information on the use of VMware in a NetWorker environment.
- *NetWorker Error Message Guide*
Provides information on common NetWorker error messages.
- *NetWorker Licensing Guide*
Provides information about licensing NetWorker products and features.

- *NetWorker REST API Getting Started Guide*
Describes how to configure and use the NetWorker REST API to create programmatic interfaces to the NetWorker server.
- *NetWorker REST API Reference Guide*
Provides the NetWorker REST API specification used to create programmatic interfaces to the NetWorker server.
- *NetWorker 19.1 with CloudBoost 19.1 Integration Guide*
Describes the integration of NetWorker with CloudBoost.
- *NetWorker 19.1 with CloudBoost 19.1 Security Configuration Guide*
Provides an overview of security configuration settings available in NetWorker and Cloud Boost, secure deployment, and physical security controls needed to ensure the secure operation of the product.
- NetWorker Management Console Online Help
Describes the day-to-day administration tasks performed in the NetWorker Management Console and the NetWorker Administration window. To view the online help, click **Help** in the main menu.
- NetWorker User Online Help
Describes how to use the NetWorker User program, which is the Windows client interface, to connect to a NetWorker server to back up, recover, archive, and retrieve files over a network.

Special notice conventions that are used in this document

The following conventions are used for special notices:

 **NOTICE** Identifies content that warns of potential business or data loss.

 **Note:** Contains information that is incidental, but not essential, to the topic.

Typographical conventions

The following type style conventions are used in this document:

Table 2 Style conventions

Bold	Used for interface elements that a user specifically selects or clicks, for example, names of buttons, fields, tab names, and menu paths. Also used for the name of a dialog box, page, pane, screen area with title, table label, and window.
<i>Italic</i>	Used for full titles of publications that are referenced in text.
Monospace	Used for: • System code • System output, such as an error message or script • Pathnames, file names, file name extensions, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Used for variables.
Monospace bold	Used for user input.
[]	Square brackets enclose optional values.
	Vertical line indicates alternate selections. The vertical line means or for the alternate selections.

Table 2 Style conventions (continued)

{ }	Braces enclose content that the user must specify, such as x, y, or z.
...	Ellipses indicate non-essential information that is omitted from the example.

You can use the following resources to find more information about this product, obtain support, and provide feedback.

Where to find product documentation

- <https://www.dell.com/support>
- <https://community.emc.com>

Where to get support

The Support website <https://www.dell.com/support> provides access to product licensing, documentation, advisories, downloads, and how-to and troubleshooting information. The information can enable you to resolve a product issue before you contact Support.

To access a product-specific page:

1. Go to <https://www.dell.com/support>.
2. In the search box, type a product name, and then from the list that appears, select the product.

Knowledgebase

The Knowledgebase contains applicable solutions that you can search for either by solution number (for example, KB000xxxxxx) or by keyword.

To search the Knowledgebase:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Knowledge Base**.
3. In the search box, type either the solution number or keywords. Optionally, you can limit the search to specific products by typing a product name in the search box, and then selecting the product from the list that appears.

Live chat

To participate in a live interactive chat with a support agent:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Contact Support**.
3. On the **Contact Information** page, click the relevant support, and then proceed.

Service requests

To obtain in-depth help from Licensing, submit a service request. To submit a service request:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Service Requests**.

Note: To create a service request, you must have a valid support agreement. For details about either an account or obtaining a valid support agreement, contact a sales representative. To get the details of a service request, in the *Service*

Request Number field, type the service request number, and then click the right arrow.

To review an open service request:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Service Requests**.
3. On the **Service Requests** page, under **Manage Your Service Requests**, click **View All Dell Service Requests**.

Online communities

For peer contacts, conversations, and content on product support and solutions, go to the Community Network <https://community.emc.com>. Interactively engage with customers, partners, and certified professionals online.

How to provide feedback

Feedback helps to improve the accuracy, organization, and overall quality of publications. You can send feedback to DPAD.Doc.Feedback@emc.com.

CHAPTER 1

Overview

This chapter contains the following topics:

- [The NetWorker environment](#).....30
- [NetWorker services](#)..... 32
- [NetWorker user interfaces](#).....41

The NetWorker environment

The NetWorker® environment provides the ability to protect an enterprise against data loss. As the enterprise grows, so does the complexity and importance of protecting data. The NetWorker software provides the power and flexibility to meet these challenges.


The NetWorker software is a cross-platform, client/server application that provides the ability to remotely manage all NetWorker Servers from a web-enabled, graphical interface.

NetWorker Components

Several components make up the NetWorker environment and provide the ability to protect against data loss.

The following figure illustrates the main components in a NetWorker environment.

Figure 1 NetWorker components



NMC Server

The NetWorker Management Console (NMC) server or Console server is a Java-based application and database server. The NMC Server manages all NetWorker Servers and

Clients. The NMC Server provides reporting and monitoring capabilities for all NetWorker Servers and Clients in the environment. The NMC Server relies on the NetWorker Authentication Service for user account authentication.

Datazone

A NetWorker datazone is composed of a single NetWorker Server, its clients, and storage nodes. You can add additional datazones as backup requirements increase.

NetWorker Authentication Service

The NetWorker Authentication Service provides centralized token-based authentication to components in a NetWorker 19.1 environment. You can configure the NetWorker Authentication Service to use a local user database or external identity providers (LDAP, LDAPS, and AD) for authentication.

NetWorker Server

The NetWorker Server is a collection of processes and programs that are installed on a host that performs NetWorker services. The NetWorker Server also acts as a storage node and can control multiple remote storage nodes.

NetWorker client

A NetWorker client is a physical or virtual computer on which you install the NetWorker client software on. The NetWorker client computer can be any computer in a datazone that contains data you want to back up. The NMC server, NetWorker server, and NetWorker storage node hosts are also NetWorker clients.

NetWorker client resource overview

A NetWorker client resource defines the data that you want to back up on a host. You can create multiple client resources for a NetWorker host, and each resource defines a different dataset.

The NetWorker client software is available for a variety of operating system platforms. Any NetWorker server can backup a NetWorker client, regardless of the platform the client resides on. For example, you can back up a NetWorker client on a Microsoft Windows computer to a NetWorker server on a Linux computer.

NetWorker Storage Node

NetWorker can back up data to local devices on a NetWorker Server or remote devices on a storage node. A storage node controls storage devices such as tape drives, disk devices, autochangers, and silos.

The NetWorker Server is a local storage node. Use a remote storage node to offload most of the data movement in a backup or a recovery operation from the NetWorker Server. A remote storage node improves performance, but it requires high I/O bandwidth to manage data transfer from local clients or network clients to target devices. The operating system of a remote storage node can differ from the NetWorker Server.

NetWorker REST API

The NetWorker REST API is an interface that allows customer to access the NetWorker data protection service and to build client applications that automate NetWorker operations. The *NetWorker REST API Getting Started Guide* describes how

to use NetWorker REST API, and the *NetWorker REST API Reference Guide* provides a full description of the API resources.

Dell EMC Licensing Solution

NetWorker 9.0.x and later servers use the Dell EMC Licensing Solution.

The Dell EMC Licensing Solution is a licensing standard that stores all licensing information for the environment in one license file, which is stored on both the NetWorker server and, if using a served license, the License Server.

All new installations of NetWorker use the Dell EMC Licensing Solution. The chapter "Dell EMC Licensing Solution" in the *NetWorker Licensing Guide* provides information on how to implement the Dell EMC Licensing Solution for new and upgraded installations of the NetWorker software. The "Dell EMC Licensing Solution" chapter also describes the Dell EMC Licensing Server and the use of the license file.

Restricted datazones

Restricted datazones provide NetWorker administrators with the ability to organize a NetWorker environment into a multi-tenancy configuration.

In a multi-tenancy configuration, each restricted datazone contains one NetWorker server and other associated NetWorker resources. Global administrators oversee the setup and management of several restricted data zones and assign tenant administrators with access to a restricted datazone. A tenant administrator can only manage NetWorker resources within an assigned restricted datazone.

Deduplication storage systems

The NetWorker software supports backup data deduplication on Data Domain® storage systems.

The *NetWorker Data Domain Boost Integration Guide* provides detailed information about setting up DD Boost deduplication devices to work with NetWorker.

Virtual environments

The *NetWorker VMware Integration Guide* provides more information on the virtual environment solutions from NetWorker.

NetWorker services

The main services and programs for the NetWorker Server are the NetWorker Storage Node, NetWorker Client, and the NetWorker Management Console (NMC) server.

This section includes information on the NetWorker services, and how to start and stop the services.

For more information about:

- Main NetWorker services—The *NetWorker Command Reference Guide* or the UNIX man pages provides more information.
- Service port requirements when configuring a firewall—The *NetWorker Security Configuration Guide* provides more information.

Processes on NetWorker hosts

Each NetWorker host requires processes to provide configuration and management support of the NetWorker software.

NetWorker Authentication Service

To support the NetWorker Authentication Service feature, one or more tomcat processes start on the NetWorker Server. The tomcat process provides the authentication service with a database server instance, which enables the authentication service to manage tokens and supports user database management.

NetWorker REST API

The NetWorker REST API service is deployed in the same Apache Tomcat container as NetWorker Authentication Service. The NetWorker REST API uses the same set of Tomcat processes to deliver its service.

NetWorker Client

The `nsrexecd` process runs on a NetWorker Client. This process authenticates and manages NetWorker Server remote execution requests and starts the `save` and `savefs` processes on the client to support backup requests.

NetWorker Server

The following table summarizes the processes that support the NetWorker Server software.

Table 3 NetWorker Server processes

Process	Function
<code>nsrctld</code>	The top-level NetWorker Server process that monitors, stops, and starts all NetWorker Server processes.
<code>nsrd</code>	- NetWorker save and recovery daemon. - The master service that controls other services on the NetWorker Server, clients, and storage nodes. - Monitors active save or recover program sessions. - In response to a recover session, <code>nsrd</code> spawns an agent process, <code>ansrd</code>.
<code>nsrmmdbd</code>	- NetWorker save and recover media management database service daemon. - Provides media database management services to the local <code>nsrd</code> and <code>nsrmmmd</code> services and records entries in the media database.
<code>nsrjobd</code>	Monitors NetWorker activity during a backup or recovery operation.

Table 3 NetWorker Server processes (continued)

Process	Function
<code>nsrindexd</code>	Provides an indexing service to read, write, and remove index entries. The <code>nsrd</code> service starts one <code>nsrindexd</code> process on the NetWorker server. The <code>nsrindexd</code> process spawns an additional helper <code>nsrindexd</code> process for each index session. NetWorker uses index sessions to read, write, or delete index entries, for example, when NetWorker saves an index, or when a user performs a file-level or browsable recover. When the read or write operation completes, the helper <code>nsrindexd</code> process closes.
<code>nsrmmgd</code>	- Manages tape library operations. - Provides an RPC-based service that manages all jukebox operations on behalf of the <code>nsrd</code> service. - The <code>nsrd</code> service starts only one instance of <code>nsrmmgd</code> on the NetWorker Server as needed.
<code>nsrlogd</code>	Supports the NetWorker audit log service, which is configured to run on the NetWorker Server by default.
<code>nsrcpd</code>	- Starts automatically when a user accesses the Hosts Task window in the NetWorker Administration interface. - Allows users to distribute and upgrade NetWorker and module software from a centralized software repository across a network.
<code>nsrdispd</code>	Handles RPC-based calls for the <code>nsrd</code> process, from remote third party processes.
<code>nsrdisp_nwbq</code>	Started by <code>nsrdispd</code> to handle NMC Server requests for information from the RAP and media databases on the NetWorker Server.
<code>nsrlmc</code>	- Supports licensing requests. - For the traditional licensing model, <code>nsrlmc</code> requests a license from the <code>lgtolmd</code> process. - For the CLP/ELMS licensing model, <code>nsrlmc</code> requests capacity and update licenses from the ELMS server.

Table 3 NetWorker Server processes (continued)

Process	Function
<code>nsrvmwsd</code>	Provides a web service to manage VMware VM backups that are part of the NetWorker VMware protection feature.
<code>tomcat7</code> (Windows), <code>tomcat</code> (UNIX)	Tomcat web server instance for the NetWorker Authentication Service.
<code>nsrexecd</code>	Authenticates and processes the NetWorker Server remote execution requests and runs the <code>save</code> and <code>savefs</code> programs on the client.

NetWorker Storage Node

The following table summarizes the services that support the NetWorker Storage Node software.

Table 4 NetWorker Storage Node processes

Process	Function
<code>nsrmmmd</code>	- Provides device support, generates mount requests, multiplexes save set data during a multi client backup, and de-multiplexes recover data. It writes the data sent by <code>save</code> to storage media. - Forwards storage information to the <code>nsrmmdbd</code> process on the NetWorker Server, which the NetWorker Server adds to the media database.
<code>nsrsnmd</code>	- Provides an RPC-based service to manage all the device operations that the <code>nsrmmmd</code> process handles on behalf of the <code>nsrd</code> process on the NetWorker Server. - Ensures that the necessary device operations are actually performed when needed by <code>nsrd</code>. - Automatically run by <code>nsrd</code> as required. - Only one <code>nsrsnmd</code> runs on each storage node that has configured and enabled devices.
<code>nsrlcpd</code>	- Provides a uniform library interface to the NetWorker media management daemon, <code>nsrmmgd</code>. - Manages the library subsystem media, slot, drive, and port resources providing control to move and access the resources within the library subsystems.

Table 4 NetWorker Storage Node processes (continued)

Process	Function
	One <code>nsrlcpd</code> starts for each configured tape library.
<code>nsrexecd</code>	Authenticates and processes the NetWorker Server remote execution requests and runs the <code>save</code> and <code>savefs</code> programs on the client.

NMC Server

The following table summarizes the processes that support the NMC Server software.

Table 5 NMC Server processes

Process	Function
<code>nsrexecd</code>	Authenticates and processes the NetWorker Server remote execution requests and runs the <code>save</code> and <code>savefs</code> programs on the client.
<code>gstd</code>	Known as the Generic Services Toolkit (GST), controls other services that are provided by the NMC Server.
<code>httpd</code>	Starts the NMC Console GUI on the client through a web browser.
<code>postgres</code>	A database server that manages information pertaining to NMC Server management. For example, Console reports.
<code>gstdsnmptrapd</code>	- Monitors SNMP Traps on a managed Data Domain system. - Provides the ability to report SNMP Trap events in the NMC Events task. - Started only when SNMP Trap monitoring is configured for the Data Domain system.

Stop and start the NMC Server

To complete some tasks in the NetWorker software, first stop the NetWorker Console service and then start the NetWorker Console service.

Stopping the NMC Server on Windows

Perform the following steps as a Windows administrator to stop the NMC Server service, which also stops the `postgres` and `httpd` processes.

Procedure

1. Right-click **My Computer**, and then select **Manage**.
2. Expand **Services and Applications**, and then select **Services**.
3. Right-click **EMC GST Service** and select **Stop**.

 **Note:** The EMC GST Service stops the EMC GST Database Service and the EMC GST Web Service.

Starting the NMC Server on Windows

Perform the following steps as a Windows administrator to start the NMC Server service, which also starts the `postgres` and `httpd` processes.

Procedure

1. Right-click **My Computer**, and then select **Manage**.
2. Expand **Services and Applications**, and then select **Services**.
3. Verify that the NetWorker Client service is running.

The **NetWorker Remote Exec Service** should have a status of Started. If the service has not started:

- a. Right-click **NetWorker Remote Exec Service**.
- b. Select **Start**.
4. Right-click **EMC GST Service**, then select **Start**.

 **Note:** The EMC GST Service starts the EMC GST Database Service and the EMC GST Web Service.

Stopping the NMC Server on Linux

Perform the following steps as root on the NMC Server to stop the NMC Server process, which also stops the `postgres` and `httpd` processes.

Procedure

1. To stop the NMC Server processes:
 - a. On sysVinit enabled Linux machines, type `/etc/init.d/gst stop`
 - b. On systemd enabled Linux machines, type `systemctl stop gst`
2. To confirm that the `gstd`, `httpd`, and `postgres` process are not running, type `ps -ef | grep lgtonmc`

Starting the NMC Server processes on Linux

Perform the following steps as root on the NMC Server to start the NMC process, which also starts the `postgres` and `httpd` processes.

Procedure

1. To verify that the NetWorker Client process, `nsrexecd` is running, type `ps -ef | grep /usr/sbin/nsr`.

When the client process is running, a message similar to the following appears:

```
root 240 1 0 ? 0:04 /usr/sbin/nsrexecd -s mysrvr
```

If `nsrexecd` is not running, type `/etc/init.d/networker start` to start the process.

2. To start the NMC Server daemon, `postgres`, and `httpd` processes., type `/etc/init.d/gst start`
3. To confirm that the `gstd`, `postgres`, and `httpd` processes have started, type `ps -ef | grep lgtonmc`

When the processes have started, output similar to the following appears:

```
nsrnmcc 7190 1 0 Nov23 ? 00:00:06 /opt/lgtonmc/bin/gstd
nsrnmcc 7196 1 0 Nov23 ? 00:00:00 /opt/lgtonmc/apache/bin/
httpd -f /opt/lgtonmc/apache/conf/httpd.conf
nsrnmcc 7197 7196 0 Nov23 ? 00:00:00 /opt/lgtonmc/
apache/bin/httpd -f /opt/lgtonmc/apache/conf/httpd.conf
nsrnmcc 7212 1 0 Nov23 ? 00:00:00 /opt/lgtonmc/
postgres/bin/postgres -D /nsr/nmc/nmcdb/pgdata
root 18176 18141 0 02:47 pts/0 00:00:00 grep lgtonmc
```

Stop and start a NetWorker Server, Client, or Storage Node

This section describes how to manually stop and start the services for a NetWorker Server, client, or storage node. Attributes exist that allow you to configure a NetWorker Server to not accept any new backup or recover sessions in preparation of a NetWorker daemon shutdown or server restart.

NetWorker Security Configuration Guide provides more information around how to prevent the NetWorker Server from accepting new backup and recover sessions.

Stopping a NetWorker host on Windows


Perform the following steps as a Windows administrator to stop the services on a NetWorker Server, Storage Node, and Client.

Procedure

1. Right-click **My Computer**, and then select **Manage**.
2. Expand **Services and Applications**, and then select **Services**.
3. Right-click **NetWorker Remote Exec Service**, and then select **Stop**.

Note: On a NetWorker Server, the **NetWorker Remote Exec Service** stops the **NetWorker Backup and Recovery** and the **NetWorker Message Queue Adaptor** services. On an NMC Server, the **NetWorker Remote Exec Service** also stops the **EMC GST Service**.

The following figure shows how to stop the **NetWorker Remote Exec Service** service.

Figure 2 Stopping the NetWorker Remote Exec Service

Starting a NetWorker host on Windows

Perform the following steps as a Windows administrator to start the services on a NetWorker server, storage node, and client.

Procedure

1. Right-click **My Computer**, and then select **Manage**.
2. Expand **Services and Applications**, and then select **Services**.
3. Start the appropriate service:
 - NetWorker server: Right-click the **NetWorker Backup and Recover Server** service and select **Start**.
 - **Note:** The NetWorker Backup and Recover Server service also starts the NetWorker Remote Exec Service and the NetWorker Message Queue Adaptor service.
 - NetWorker client or storage node: Right-click the **NetWorker Remote Exec Service** and select **Start**.

Stopping a NetWorker host on UNIX

Perform the following steps as the root user to stop the NetWorker processes on a NetWorker server, storage node, or client.

Procedure

1. To stop the NetWorker processes:
 - a. On sysVinit enabled Linux machines, type `/etc/init.d/networker stop`
 - b. On systemd enabled Linux machines, type `systemctl stop networker`
2. To confirm that the NetWorker processes are not running, type the following command from a prompt:

```
ps -ef | grep /usr/sbin/nsr
```

Starting a NetWorker host on UNIX

Perform the following steps as the root user to start the NetWorker processes on a NetWorker server, storage node, or client.

Procedure

1. Type the appropriate startup command for the operating system, as summarized in the following table.

Table 6 NetWorker startup commands

Operating system	Startup command
Solaris, Linux	<code>/etc/init.d/networker start</code> For systemd enabled Linux machines - <code>systemctl start networker</code>
HP-UX	<code>/sbin/init.d/networker start</code>
AIX	<code>/etc/rc.nsr</code>

2. Type `/etc/init.d/networker status` to confirm that the NetWorker processes that are appropriate to the NetWorker installation type have started.

[Processes on NetWorker hosts](#) on page 33 provides more information.

Stopping the NetWorker processes on Mac OS X

Perform the following steps as a Mac Administrator to stop the NetWorker processes on a Mac OS X host.

Procedure

1. Open the Mac OS-X Terminal application utility.
2. To stop the NetWorker processes, type the following command:

```
sudo launchctl unload /Library/LaunchDaemons/
com.xyz.NetWorker.plist
```

 **Note:** The `launchd` daemon/agent manager controls the NetWorker processes, and NetWorker configures the processes to run continuously on the host in the background. It is not recommended that you manually stop and start NetWorker processes under normal operating conditions.

Starting the NetWorker process on Mac OS X

Perform the following steps as a Mac Administrator to start the NetWorker processes on a Mac OS X host.

Procedure

1. Open the Mac OS X Terminal application utility.
2. Type `launchctl load /Library/LaunchDaemons/com.emc.NetWorker.plist` to start the NetWorker client process.

NetWorker user interfaces

The NetWorker application consists of several user interfaces that provide the ability to configure and use NetWorker features and functionality.

NMC user interface

The NMC server uses `httpd` to provide administrators with a graphical user interface to connect to an NMC server and managed NetWorker servers. The NMC UI can be accessed from any computer in the environment with a supported web browser and Java Runtime Environment (JRE).


The *NetWorker Installation Guide* provides more information about the web browser and JRE requirements for a host that runs the NMC UI. Multiple users can use different browser sessions on different hosts to access the NMC UI simultaneously.

NMC GUI

Use the NMC GUI to manage an NMC server and NetWorker servers.


The following figure illustrates the NMC GUI.

Figure 3 NMC GUI window



The NMC window is the first point of access for NMC and NetWorker tasks. The following table lists the task-based windows that can be opened from the NMC window taskbar.

Table 7 Windows opened from the NMC GUI

Button	Window	Description
	Enterprise	Select a NetWorker server to manage and monitor the server and its backup clients. The Enterprise window provides the ability to open the Administration window for a NetWorker server.
	Reports	Configure and view NMC reports.
	Setup	Control administrative functions: • User management — Add, edit, and delete NMC user accounts, restrict user views of servers. The <i>NetWorker Security Configuration Guide</i> provides information about user management. • License management — Manage NetWorker licenses. The <i>NetWorker Licensing Guide</i> provides information about license management.

NetWorker Administration window

The **NetWorker Administration** window provides you with the ability to manage and configure NetWorker server resources in a GUI. The NMC UI provides you with the ability to open up a **NetWorker Administration** window for each managed NetWorker server.

NetWorker client interface

Manual back up, recovery, and archive operations can be performed from a client. Manual operations are not scheduled. They are client-initiated tasks that are performed when a user wants to back up, recover, or archive one or more files on the NetWorker host immediately. You can schedule backup, recovery, and archive operations in the NMC GUI.

On Windows hosts only, you can use the NetWorker User GUI to perform manual back up, recovery, and archive operations.

On UNIX and Windows hosts, you can use command line utilities to perform manual operations:

- Use the `save` command to perform a manual backup.
- Use the `recover` command to perform a manual recovery.
- Use the `nsarchive` command to perform a manual archive.

The *NetWorker Command Reference Guide* or the UNIX man pages provide more information about these commands.

NetWorker character-based interface

Use the NetWorker character-based interface (`nsradmin`) to perform configuration and management tasks in the NetWorker server resource database (`resdb`) and the NetWorker client resource database (`nsrexec`).

You can start the `nsradmin` interface by typing this command:

```
nsradmin
```

For more information about `nsradmin`, the *NetWorker Command Reference Guide* or the UNIX man pages provides more information.

NetWorker command-line interface

Perform client and server tasks by typing commands at the prompt. The *NetWorker Command Reference Guide* or the UNIX man pages provides information about these commands.

Introduction to the NetWorker Management Web UI

The NetWorker Management Web UI is a web-based management interface that provides support for the following NetWorker operations:

- Managing VMware vCenter servers
- Managing VMware Proxies
- Installing the vCenter Plugin
- Recovering virtual machines
- Monitoring recovery operations
- Creating and updating policies, workflows, and actions
- Creating and updating groups
- Creating and updating rules

The following table provides more information on the functionality available in the NetWorker Management Web UI.

Table 8 Supported operations in the NetWorker Management Web UI

Operation	Description
Protection	VMware vCenter servers • Manage vCenter servers. • Refresh and view the vCenter inventory. • View properties of entities in the vCenter Inventory tree. Groups

Table 8 Supported operations in the NetWorker Management Web UI (continued)

Operation	Description
	• Add, edit, and delete groups • Refresh and view groups Policies • Add, edit, and delete policies, workflows, and actions • Refresh and view policies, workflows, and actions Rules • Add, edit, and delete rules • Refresh and view rules
	VMware vProxies • Manage vproxies. • Monitor progress of vProxy registration.
Recovery	Recover virtual machines. Supports both image-level and file-level recovery.
Monitoring	• View and monitor the progress of virtual machine recovery; includes the list of completed and currently running recover jobs. • View recover logs.

You can log in to the NetWorker Management Web UI by using the NetWorker credentials for authentication.

The *NetWorker VMware Integration Guide* and the *NetWorker Administration Guide* provides more information on how to use the NetWorker Management Web UI to perform the supported tasks.

The *NetWorker Installation Guide* provides more information on how to install the NetWorker Management Web UI.

 **Note:** NetWorker Management Web UI is not backward-compatible with the earlier versions of NetWorker.

Supported browsers

The NetWorker Management Web UI supports the latest versions of the following browsers:

- Microsoft Internet Explorer 11
- Google Chrome
- Microsoft Edge
- Mozilla Firefox
- Safari

CHAPTER 2

Getting Started

This chapter contains the following topics:

- [NetWorker Management Console interface](#)..... 46
- [Connecting to the Administration window](#)..... 51
- [Getting started with a new installation](#).....66

NetWorker Management Console interface

The interface for NetWorker Management Console (NMC), also called the NetWorker Console, consists of both the Console window, and the Administration window.

 **Note:** To start NMC, you must use 64-bit Java. NMC will fail to start if Java 32-bit is used.

Connecting to the Console window

The following sections describe how to connect to the **Console** window:

Before you connect

Ensure that you configure the NetWorker datazone correctly, and that the required daemons are running on the NetWorker Server and the NMC Server.

- **Linux NetWorker Server**—Confirm that the NetWorker daemons have started, by typing the command below, based on the initialization system running on your Linux machine : `/etc/init.d/networker status`. For a NetWorker server, the `nsrctld` daemon starts. The `nsrctld` daemon starts other processes that the NetWorker server requires. Output similar to the following example appears when the daemons are started:

```
+--o nsrctld (29021)
+--o epmd (29029)
+--o rabbitmq-server (29034)
+--o beam (29038)
+--o inet_gethost (29144)
+--o inet_gethost (29145)
+--o jsvc (29108)
+--o jsvc (29114)
+--o nsrd (29123)
+--o java (29135)
+--o nsrmmdbd (29828)
+--o nsrindexd (29842)
+--o nsrdispd (29853)
+--o nsrjobd (29860)
+--o nsrvmsd (29968)
+--o eventservice.ru (29154)
+--o jsvc (29158)
+--o jsvc (29159)
+--o java (29838)
+--o node-linux-x64- (29885)
+--o nsrexecd (29004)
+--o nsrlogd (29899)
+--o nsrsnmd (30038)
```

- **Linux NMC Server:**

1. Type `ps -ef | /usr/sbin/nsrexecd`. Output similar to the following example should appear:

```
root 24959 1 1 13:29 ? 00:00:00 /usr/sbin/nsrexecd
```

If you do not see this output, type `/etc/init.d/networker start`.

2. Type `ps -ef | grep lgtommc` Output similar to the following should appear:

```

nsrnmcc 7190 1 0 Nov23 ? 00:00:06 /opt/lgtommc/bin/gstd
nsrnmcc 7196 1 0 Nov23 ? 00:00:00 /opt/lgtommc/apache/bin/
httpd -f /opt/lgtommc/apache/conf/httpd.conf
nsrnmcc 7197 7196 0 Nov23 ? 00:00:00 /opt/lgtommc/
apache/bin/httpd -f /opt/lgtommc/apache/conf/httpd.conf
nsrnmcc 7212 1 0 Nov23 ? 00:00:00 /opt/lgtommc/
postgres/bin/postgres -D /nsr/nmc/nmcdb/pgdata
root 18176 18141 0 02:47 pts/0 00:00:00 grep lgtommc

```

- **Windows NetWorker Server:**
 1. Confirm that the following services are started: NetWorker Backup and Recover Server, NetWorker Message Queue Adaptor, and NetWorker Remote Exec Service.
 2. If these services are not started, start the NetWorker Backup and Recover Server Service.
- **Windows NMC Server:**
 1. Confirm that the following services are started: EMC GST Database Service, NetWorker Server Service, and NetWorker Server Web Service.
 2. If these services are not started, start the NetWorker Server service.

Connecting to the NMC server GUI

About this task

Complete the following procedure to connect to the NMC Server GUI from an NMC client. By default, the NetWorker Authentication Service uses the local user database for user authentication. Specify the NetWorker Authentication Service administrator account to log in to the NMC Server. The *NetWorker Security Configuration Guide* describes how to configure the NetWorker Authentication Service to use LDAP or AD for user authentication.

Procedure

1. From a supported web browser session, type the URL of the NMC Server:

`https://server_name:https_service_port`
where:

- *server_name* is the name of the NMC Server.
- *https_service_port* is the port for the embedded HTTP server. The default https port is 9000.

For example: `https://houston:9000`

The `gconsole.jnlp` file downloads to the host. When the download completes, open the file.

2. Optional, associate the `jnlp` file with a program.

When you use Mozilla Firefox on Windows, and the `jnlp` extension is not associated with Java, you are prompted to choose the program that opens the `jnlp` file. In the dialog box that appears, select **Open with**, and then select Java (TM) Web Start Launcher. If this application does not appear, browse to the Java folder and select the `javaws.exe` file.


3. On the **Welcome** page, click **Start**.

Note: If the **Start** button does not appear but you see a warning message that states that Java Runtime Environment cannot be detected, click the [here](#) hyperlink.

4. For Internet Explorer only, if a security warning appears, select **I accept the risks and want to run this application**, and then click **Run**.
5. On the **Log in** page, specify the NetWorker Authentication Service administrator username and password, and then click **OK**.
6. On the **Licensing Agreement** page, select **Accept**.
7. On the **Welcome to the NMC Server Configuration Wizard** page, click **Next**.

The following figure shows the **Welcome to the NMC Server Configuration Wizard** page.

Figure 4 Welcome to the NMC Server Configuration Wizard page



8. On the **Set authentication server service account for the NMC server** page, review the setting and click **Next**.

The following figure shows the **Set authentication server service account for the NMC server** page.

Figure 5 Set authentication server service account for the NMC Server page

Console Configuration Wizard

Set authentication server service account for the NMC server

The following service account will be created on the Authentication Server.
If the account exists, it will be updated.

Authentication Server:

Port:

Service Account Name:

< Back Next > Cancel


9. On the **Specify a list of managed NetWorker Servers** page:
 - a. Specify the names of the NetWorker Servers that the NMC Server will manage, one name per line.

Note: If the NMC Server is also the NetWorker Server, specify the name of the NetWorker Server.
 - b. Leave the default **Capture Events** and **Gather Reporting Data** options enabled.

Consider the following options:

- To allow the NMC Server to monitor and record alerts for events that occur on the NetWorker Server, select **Capture Events**.
- To allow the NMC Server to collect data about the NetWorker Server and generate reports, select **Gather Reporting Data**.

The following figure shows the **Specify a list of managed NetWorker servers** page.

Figure 6 Specify a list of managed NetWorker servers page

10. Click **Finish**. The installation starts the default web browser and connects to the NMC server. The **NetWorker Management Console** and **Getting Started** windows appear.
11. In the **Enterprise** window, right-click the NetWorker Server, and then select **Launch Application**.

Note: If you do not specify any NetWorker Servers in the **Specify a list of managed NetWorker servers** page, the NMC **Enterprise** window does not display any NetWorker Servers. To add a host, in the left navigation pane, right-click **Enterprise**, and then click **New > Host**. The **Add New Host** wizard appears.

Connecting to the NMC server after the first time

Use one of the following methods to connect to the NMC server after the initial connection.

- Point the browser to the same URL.
- Double-click the NMC product name in the Java Web Start Application Manager.
- Double-click the desktop button , if one was configured by using the Java Web Start Application Manager.

Connecting to the NMC GUI using an ssh connection

You can use ssh port forwarding to connect to the NMC server and generate reports, from the NMC client.

About this task

Perform the following steps on the NMC client.

Procedure

1. Open an ssh connection from the NMC client to the NMC server with ssh tunnels for ports 9000 and 9001.

For example:

```
ssh -L9000:localhost:9000 -L9001:localhost:9001 -
L5432:localhost:5432 Console_servername -N
```

 **Note:** If you changed the default NMC server ports, specify the correct port numbers.

2. Use `javaws` to connect to the NMC server.

For example:

```
javaws https://localhost:9000/gconsole.jnlp
```

Connecting to the Administration window

The following sections describe how to connect to the Administration window and browse through the interface.

Opening the Administration window

You can add and select a NetWorker server and open the **Administration** window.

About this task


Procedure

1. From the **Console** window, click **Enterprise**.
2. Add one or more NetWorker servers:
 - a. Highlight **Enterprise** in the navigation tree.
 - b. From the **File** menu, select **New>Host**.
 - c. Type the name of the host on which the NetWorker server is running, and click **Next**.
 - d. Select **NetWorker** for the type of application to be managed.
 - e. Click **Finish**.
 - f. Repeat for all NetWorker servers in the network.
3. From the left pane, click a host in the **Enterprise** list.
4. From the right pane, click the application and select **Enterprise > Launch Application**, or double-click the application. The **Administration** window opens as a separate application.

Administration window

NetWorker Servers are managed through the **Administration** window.

The following figure illustrates the NetWorker **Administration** window.

Figure 7 Administration window

You can toggle between the **Administration** window and the NMC UI.

The following table lists the windows that can be launched from the **Administration** window taskbar.

Table 9 Windows that are launched from the Administration window

Butt on	Window	Description
	Monitoring	Monitor various activities that are related to the NetWorker Server. For example, you can monitor the progress of a policy and view any alerts. A portion of the Monitoring window always appears at the bottom of the Administration window, providing information on Log messages and Alerts.
	Protection	Manage NetWorker Server resources such as clients, groups, policies, probes, and schedules. Provide the ability to monitor, start, stop, and restart data protection policies.
	Recover	Manage recover configurations and schedule recover jobs for NetWorker hosts from a centralized location on the NMC Server.
	Devices	Add, configure, and operate single or multiple devices, libraries, and silos for the NetWorker Server.
	Media	Manage the activities and the resources that are related to backup volumes. For example, you can mount a backup volume or create a label template for backup volumes.
	Hosts	View information about known NetWorker hosts such as the NetWorker version, CPU type, and operating system. Manage the NetWorker Client resource database. Perform software upgrades on NetWorker hosts by using Package Manager.

Table 9 Windows that are launched from the Administration window (continued)

Butt on	Window	Description
	Server	Manage NetWorker Server resources such as licenses, notifications, user groups, directives, and restricted datazones.

Editing multiple resources

In the NMC Protection window, you can edit an attribute for multiple resources at the same time.

About this task

For example, if you want the schedule for all clients within a group to change from the default to “Full Every Friday”, perform the following steps:

Procedure

1. Select each client resource row in the window.
2. Place the cursor in the column you want to change (in this case, the Schedule column).

The color of the column changes when the cursor is in the column.

3. Right-click in that column and select from the list of available options. The options include Edit, Add to, and Remove from, depending on the column selected.

Only the columns that appear in the window can be selected for multiple resource editing. To add a column that is not currently in view:

- a. Right-click a table header and select **Add Column** from the drop-down.
- b. Select from the list of available attributes.

Drag-and-drop functionality

Drag-and-drop functionality is available in the Console and Administration interfaces for many tasks.

Drag-and-drop between resource types in the Console window

The drag-and-drop functionality allows multiple resources to be selected and moved from one resource type to another.

In the Enterprise window from the Console interface, you can drag-and-drop to perform the following actions:

- Copy an individual folder in the enterprise hierarchy by selecting the folder, press and holding the **Ctrl** key, and dragging the folder to a new location.
- Move an individual folder in the enterprise hierarchy to a new location by selecting and dragging a folder to a new location.
- Copy an individual host node in the enterprise hierarchy by selecting and dragging the host to a new parent folder.

- Move an individual host node in the enterprise hierarchy by selecting and dragging the host to a new parent folder.
- Copy a selected number of objects in a folder to a new folder in the hierarchy tree or folder contents table. Select an individual folder in the navigation tree to display the contents of the folder, select the contents, while pressing **Ctrl**, drag the contents to a new folder. Select a collection of folders or hosts and drag them to a new folder by creating a copy of the selected contents in a new location.
- Move a selected number of objects in a folder to a new folder in the hierarchy tree or folder contents table. Select an individual folder in the navigation tree to display the contents of the folder, select the contents, and drag the contents to a new folder. Select a collection of folders and or hosts and drag them to a new folder by moving the selected contents to a new location.



Note: Only one object may be selected for drag-and-drop in the navigation tree.

Client and group management in the Administration window

The drag-and-drop functionality allows multiple clients or groups to be selected and moved from one location to another. You can use drag-and-drop functionality in the **Protection** window to do the following:

- Copy selected clients to a new NetWorker group:
 1. In the left navigation pane, expand the server resource, and then expand the **Groups** resource.
 2. Select **Clients** in the directory tree.
 3. Drag-and-drop the client objects from the **Client Summary** table to a group in the directory tree.
- Move selected clients from one NetWorker group to another group:
 1. Select a group in the directory tree.
 2. Move clients from the **Client Summary** table to another NetWorker group.

Library operations in the Devices window

The drag-and-drop functionality allows multiple slots or devices to be managed in the Devices window.

You can use drag-and-drop functionality to manage media from the **Library** window from the Devices task, for instance:

- Mount an individual volume onto a device by selecting a slot in the Slots table and dragging it to a device in the Devices table.
- Mount multiple volumes to available devices as assigned by the NetWorker server. To mount multiple volumes, select multiple slots in the Slots table and drag them anywhere in the Devices table.
- Unmount a volume from a selected device and deposit it back in its designated slot by selecting an individual device from the Devices table and dragging it anywhere in the Slots table. The volume image displays in the corresponding slot.
- Unmount multiple volumes from a selected device and deposit them back in their designated slot by selecting the devices from the Devices table and dragging them anywhere in the Slots table. The volumes display in the corresponding slots.

Copy and paste tabular information to operating system clipboard

Tabular information can be selected and moved to an operating system clipboard by using drag-and-drop functionality. All tables support selection of multiple rows in a

table and the ability to copy and paste the data in the selected rows to the system clipboard. Subsequently, the data in the operating system clipboard can be moved to a target application.

 **Note:** Drag-and-drop operations from the operating system clipboard to a table are not supported.

Multiple library devices and slots

A single operation can be performed on multiple library devices and slots. Multiple rows can be selected in both the Devices and Slots tables simultaneously.

In the Devices table for a library, multiple devices can be selected to perform the following operations:

- Unmount
- Release device (STL only)
- Enable/Disable

In the Slots table for a device, multiple volume operations can be performed for the following operations:

- Mount
- Load without mount
- Withdraw
- Label
- Inventory
- Remove (STL only)

Setting user interaction preferences

Depending on the window button that was selected from the Console window, you can set various user preferences such as the user interface font, font size, parallel windows, and table settings. For the Reports window, there are ways you can enhance the viewing of displayed reports.

About this task

Procedure

1. On the main menu, select **View**.
2. Set the various options available under the selected window button. You may need to click **OK**, depending on the option selection.

Monitoring NetWorker Server activities in the Administration window

The **Monitoring** window in the NetWorker **Administration** application enables you to monitor the activities of an individual NetWorker Server.

About this task

The **Monitoring** window provides the following types of activity and status information:

- Data protection policies, workflows, and individual actions.
- Cloning, recovering, synthetic full backups, and browsing of client file indexes.
- Operations that are related to devices and jukeboxes.


- Alerts and log messages.

You can also perform some management operations from the **Monitoring** window, for example, starting, stopping, or restarting a data protection policy.

Procedure

1. From the **NMC Console** window, click **Enterprise**.
2. In the **Enterprise** view, right-click the NetWorker Server, and then select **Launch Application**.
The **Administration** window appears.
3. To view the **Monitoring** window, click **Monitoring**.

Figure 8 Monitoring window



About the Monitoring window

On the **Administration** window taskbar, select **Monitoring** to view the details of current NetWorker server activities and status, such as:

- Policies and actions.
- Cloning, recovering, synthetic backups, checkpoint restart backups, and browsing of client file indexes.
- Alerts and log messages, and operations that are related to devices and jukeboxes.

While the **Monitoring** window is used primarily to monitor NetWorker server activities, it can also be used to perform certain operations. These operations include starting, stopping, or restarting a workflow.

The **Monitoring** window includes a docking panel that displays specific types of information. Select the types of information you want to view from the docking panel.

A portion of the **Monitoring** window, which is known as the task monitoring area, is always visible across all windows. A splitter separates the task monitoring area from

the rest of the window. You can click and move the splitter to resize the task monitoring area. The arrow icon in the upper right corner of the **Monitoring** window allows you to select which tasks you want to appear in this view.

Smaller windows appear within the **Monitoring** window for each window. Each smaller window, once undocked, is a floating window and can be moved around the page to customize the view. You can select multiple types from the panel to create multiple floating windows that can be viewed simultaneously. The following table describes the various types of information available in the docking panel, and the details each one provides.

Table 10 Monitoring window panel

Window	Information provided
Policies/Actions	The Policies tab provides you with status information about all configure policies and the associated workflows and actions. The Actions tab provides you with status information for all actions. Policies/Actions pane on page 741 provides more information.
Sessions	Allows you to customize whether to display all session types, or only certain session types. The information that is provided depends on which session type you select. For example, if you select Save Sessions , the window lists clients, save sets, groups, backup level, backup start time, duration of the backup, devices, rate, and size. Sessions window on page 57 provides more information.
Alerts	Lists the priority, category, time, and message of any alerts. Alerts pane on page 58 provides more information.
Devices	Lists devices, device status, storage nodes, libraries, volumes, pools, and related messages. Devices pane on page 59 provides more information.
Operations	Lists the status of all library and silo operations, including <code>nsrjb</code> operations that are run from the command prompt. Also lists user input, libraries, origin, operation data, operation start time, duration of the operation, progress messages, and error messages. When displaying Show Details from the Operations window, the length of time that the window is displayed depends on the value that is typed in the Operation Lifespan attribute on the Timers tab of the Properties dialog box for the corresponding library. To access library properties, click Devices in the taskbar. By default, this pane is hidden.
Log	Lists messages that are generated by the NetWorker server, including the priority of each message, the time the message was generated, the source of the message, and the category. Log window on page 62 provides more information.

Sessions window

Use the **Sessions** window to view the sessions that are running on a NetWorker server. You can change the view of this window to display these sessions:

The **Sessions** pane below the **Policies/Actions** pane provides details on individual save, recover, clone, and synthetic full sessions by client.

To view all sessions or to limit the list of sessions by the session type, click the tabs at the bottom of the **Sessions** pane. Session types include:

- Save
- Recover
- Clone
- Browse
- Synthetic Full/Rehydrated Sessions
- All

To change the displayed session types go to **View > Show**, and select the type of sessions to display. To display all sessions currently running on the NetWorker Server, regardless of type, select **All Sessions**.

You can stop a session (backup, synthetic full backup, clone, and recovery sessions) from the **Monitoring** window, even if the session was started by running the `nsrworkflow` command.

To stop a session, right-click the workflow in the pane, and select **Stop** from the list box.

Alerts pane

The **Alerts** pane displays alerts that are generated by a particular NetWorker server or Data Domain system that has devices that are configured on the NetWorker server. The **Alerts** pane includes priority, category, time, and message information.

An icon represents the priority of the alert. The following table lists and describes each icon.

Table 11 Alerts window icons

Icon	Label	Description
	Alert	Error condition detected by the NetWorker server that should be fixed by a qualified operator.
	Critical	Severe error condition that demands immediate attention.
	Emergency	Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority.
	Information	Information about the current state of the server. This icon represents the lowest priority.
	Notification	Important information.
	Waiting	The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.
	Warning	A non-fatal error has occurred.

When items on the **Alerts** pane are sorted by the **Priority** column, they are sorted in alphabetical order based on the label of the icon.

Removing alerts

Remove individual alert messages from the **Events** tables by removing them from the **Events** table. To delete a message in the **Events** table, right-click the message, and select **Dismiss**.

About this task

 **Note:** The alert message remains in the **Log** window in the NetWorker **Administration** program.

Devices pane

The **Devices** pane allows you to monitor the status of all devices, including NDMP devices. If the NetWorker server uses shared and logical devices, the window is adjusted dynamically to present a set of columns appropriate for the current configuration.

The **Devices** pane provides the following information:

- Status of the operation.
- Name of the device.
- Name of the storage node that contains the device.
- For tape devices, the name of the library that contains the device.
- Name of the volume in the device.
- Name of the pool that is associated with the volume.
- Last message generated for the device.
- Whether the operation requires user input.

For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape.

[Entering user input](#) on page 61 provides instructions on how to deal with a user input notification.

If the current server configuration includes a shared device, a **Shared Device Name** column appears on the **Devices** pane. The name of the shared device appears in the **Shared Device Name** column. If other devices for that configuration are not shared devices, then the **Shared Device Name** column is blank for those devices. Only a single device per hardware ID can be active at any particular moment. The information for inactive shared devices is filtered out, and as a result, only one device per hardware ID is presented on the window at any time.

An icon represents the device status. The following table lists and describes each icon.

Table 12 Devices status icons

Icon	Label	Description
	Library device active	The library device is active.
	Library device disabled	The library device is disabled.
	Library device idle	The library device is idle.
	Stand-alone device active	The stand-alone device is active.

Table 12 Devices status icons (continued)

Icon	Label	Description
	Stand-alone device disabled	The stand-alone device is disabled.
	Stand-alone device idle	The stand-alone device is idle.

When you sort items in the **Devices** pane by the **Status** column, NetWorker sorts the devices in alphabetical order based on the label name of the icon.

Operations window

The **Operations** window displays information about device operations. It provides the following information:

- Status of the operation.
- Name of the library.
- Whether the operation requires user input.
For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape. [Entering user input](#) on page 61 provides instructions on how to deal with a user input notification.
- The origin, or source, of the operation.
For example, the interface, nsrjb or the NetWorker server.
- Time the operation started.
- Type of operation.
- Duration of the operation.
- Status messages from the operation.
- Any error messages.

 **NOTICE** Only the last error message of the operation appears in the **Error Messages** column. Move the mouse pointer over the cell containing the last error message to display the entire list of error messages.

The operation status is represented by an icon. The following table lists and describes each of the icons.

Table 13 Operations window icons

Icon	Label	Description
	Failed	The operation failed.
	Queued	The operation is waiting in the queue to run.
	Retry	The operation failed, but may work if you try again.
	Running	The operation is running.
	Successful	The operation completed successfully.
	User Input	The operation requires user input.

When items on the **Operations** window are sorted by the Status column, they are sorted in alphabetical order based on the label of the icon.

Viewing operation details

The **Operation Details** dialog box opens, providing information about the completion of the operation. The **Completion Time** displays the time that the operation finished. The time that it took to complete the operation is the difference between the completion and start times of the operation.

About this task

To save operation details to a file, click **Save** in the **Operation Details** dialog box. When prompted, identify a name and location for the file.

Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation, then select **Show Details**.

Stopping an operation

Certain operations can be stopped from the **Operations** window.

Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation to stop, then select **Stop**.
4. Click **Yes** to confirm the stop.

 **Note:** Operations that were started from a command line program, such as the `nsrjb` command, cannot be stopped from the **Operations** window. To stop these operations, press `ctrl-c` from the window where the command was started.

Entering user input

If the system requires user input, select the labeling operation in slow/verbose mode and the **Supply User Input** icon appears.

Procedure

1. Right-click the operation, then select **Supply Input**.
2. Confirm the requirement to supply input.
 - If **Yes**, and input is supplied, the icon in the **User Input** column disappears.

 **Note:** If two users try to respond to the same user input prompt, the input of the first user takes precedence, and the second user receives an error message.

- If **No**, and input is not supplied, the operation will time out and fail.

Log window

To view the most recent notification logs, click the **Log** window from the docking panel in the **Monitoring** window. The **Log** window provides the priority, time, source, category, and message for each log.

 **Note:** If a particular log file is no longer available, check the log file on the NetWorker server. The log files are located in *NetWorker_install_path* \logs directory.

An icon represents the priority of the log entry. The following table lists and describes each icon.

Table 14 Icons in the Log pane

Icon	Label	Description
	Alert	Error condition that is detected by the NetWorker server that should be fixed by a qualified operator.
	Critical	Severe error condition that demands immediate attention.
	Emergency	Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority.
	Information	Information about the current state of the server. This icon represents the lowest priority.
	Notification	Important information.
	Waiting	The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.
	Warning	Non-fatal error has occurred.

When you sort items on the **Log** pane by using the **Priority** column, NetWorker sorts the icons in alphabetical order based on the name of the label.

Recover window

The **Recover** window displays information about recover configurations that are created with the NetWorker Management Console (NMC) Recovery wizard.

You can use this window to:

- Start the NMC Recovery wizard to create recover configurations or modify saved recover configurations.
- Identify the status of a recover configuration that is created with the NMC Recovery wizard.
- Start and stop a recover job.

The **Recover** window is divided into five sections:


- **Toolbar**—The toolbar is hidden by default. To display the recovery toolbar, select **View > Show toolbar**
- **Summary**

- Configured Recovers
- Currently Running

A splitter separates the **Configured Recovers** section from **Currently running** window. You can click and move the splitter to resize these two windows.

The following table shows an example of the **Recover** window.

Figure 9 Recover window



Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. The following table summarizes the function of each toolbar button.

Table 15 Recovery toolbar options

Button	Function
	Starts the NMC Recover wizard to create recover configurations.
	Displays the Properties window for the saved recover configuration that you selected in the Configured Recover window.
	Deletes the saved recover configuration that you selected in the Configured Recover window.
	Displays online help for the Recover window.
	Displays the Find window at the bottom of the Recover window. The Find window allows you to perform keyword searches for messages that appear in the Logs window.
	Start the recover operation for a selected saved recover configuration. This option is only available for a recover configuration that has a Never run, or Failed status.
	Stop in-progress recover operation that you selected in the Currently Running window.

 **Note:** The **Recover** toolbar does not appear by default. To display the **Recover** toolbar, select **View > Show toolbar**.

Recover Summary

The Recover Summary section displays a high-level overview of recover jobs.

This section includes the following information:

- Total Recovers—The total number of successful recover jobs.
- Since—The number of successful recover jobs since this date.

Configured Recovers

The **Configured Recovers** window displays a list of saved recover configurations in a tabular format. You can sort the information by column. The **Configured Recovers** table displays the following information for each saved recover configuration:

- Status—The job status of a saved recover configuration.
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB.
- Comment
- OS—The operating system of the source host.
- Recover requestor—The Windows or UNIX account used to create the recover configuration.
- Start Time
- End Time
- Start date

Table 16 Save recover configuration job status

Icon	Description
	The last recover attempt failed.
	The last recover attempt completed successfully.
	The recover job has never run.
	The recover job is scheduled to run in the future.
	The recover job has expired.

Currently running

The **Currently Running** window displays a list of in progress recover jobs in a tabular format. You can sort the information by column. The **Currently Running** table displays the following information for each job:

- Status
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB
- Volume
- Comment
- Device
- Size
- Total size
- % complete
- Rate (KB/s)
- Start time
- Duration
- Currently running

Find

The **Find** section appears along the bottom of the **Recover** window, after you select the **Find** button on the **Recover** toolbar. **Find** allows you to search for keywords in the **Configured Recovers** window. The following table summarizes the available find options.

Table 17 Find options

Find option	Description
Find	Highlight the first saved recover configuration that contains the specified keyword.
Prev	Highlight the previous saved recover configuration that contains the specified keyword.
Highlight All	Highlights each saved recover configuration that contains the specified keyword.
Sort Selected	Sorts each highlighted recover configuration in the Configured Recover table so that they appear at the top of the Configured Recover table.
Match case	Make the keyword search case sensitive.

Getting started with a new installation

The following section provides basic information on how to get started with a new installation by configuring the NetWorker datazone and starting the NetWorker Management Console (NMC) Enterprise window and Administration window.

Common NetWorker tasks

There are several common tasks available in the NetWorker Console.

Adding a new host

You can add hosts by using the NetWorker Console.

About this task

Procedure

1. Log in to Console as a NetWorker Administrator.
2. Click the **Enterprise** button  on the taskbar.
3. Right-click **Enterprise** in the navigation tree.
4. Select **New > Host**.
5. In the **Host Name** field, specify the IP address or DNS name of the NetWorker server and click **Next**.
6. On the **Select Host Type** window, select **NetWorker** and click **Next**.
7. On the **Manage NetWorker** window, leave the default options **Capture Events** and **Gather Reporting Data** enabled.
 - Enable the **Capture Events** option to allow the NMC server to monitor and record alerts for events that occur on the NetWorker server.
 - Enable the **Gather Reporting Data** option to allow the NMC server to automatically collect data about the NetWorker server and generate reports on the NMC server.
8. Click **Finish**.

Device configuration

You can configure devices to test the NetWorker software.

Configuring a stand-alone tape device

Procedure

1. Log in to the NMC GUI as an administrator of the NetWorker server.
2. On the taskbar, click the **Enterprise** icon .
3. In the navigation tree, highlight a host:
 - a. Right-click **NetWorker**.
 - b. Select **Launch Application**. The **NetWorker Administration** window appears.
4. On the taskbar, click the **Devices** button .

5. In the navigation tree view, right-click a host and select **Scan for Devices**.
The **Scan for Devices** window appears.
6. On the **Select Target Storage Nodes** window, perform either of the following steps:
 - Select the storage node for the library.
 - Click **Create a new Storage Node**.
7. Select **Start scan**.
NetWorker scans for new devices. The **Log** pane provides the status of the scan operation.
8. On the left pane, select **Devices** and then from the right pane, select the new device.
9. From the **Devices** menu, select **Devices > Device Operations > Label**.
10. In the **Label** window, verify the information and click **OK**.

Configuring a stand-alone advanced file type device

Create a device that is local to the NetWorker server to receive the backup data.

Procedure

1. Log in to the NMC GUI as an administrator of the NetWorker server.
2. On the taskbar, click the **Enterprise** icon .
3. In the navigation tree, highlight a host:
 - a. Right-click **NetWorker**.
 - b. Select **Launch Application**. The **NetWorker Administration** window appears.
4. On the taskbar, click the **Devices** button .
5. From the **File** menu, select **New Device Wizard**.
6. On the **Select the Device Type** window, select **Advanced File Type Device (AFTD)**, then click **Next**.
7. On the **Select Storage Node** window, leave the default values, and click **Next**.
8. On the **Select the Device Path** window, select an empty folder or create a new folder on the NetWorker server, then click **Next**.
9. On the **Configure Device Attributes** window, specify a name for the new device in the **NetWorker Device Name** field, for example: `myaftd`, and click **Next**.
10. On the **Label and Mount Devices** window, leave the default values and click **Next**.
11. In the **Review the Device Configuration Settings** window, review the configuration information, and click **Configure**.
12. Click **Finish**.

Configuring an autochanger or silo

You can configure a new library resource.

About this task

Procedure

1. Log in to the NMC GUI as an administrator of the NetWorker server.
2. On the taskbar, click the **Enterprise** icon .
3. In the navigation tree, highlight a host:
 - a. Right-click **NetWorker**.
 - b. Select **Launch Application**. The **NetWorker Administration** window appears.
4. On the taskbar, click the **Devices** button .
5. From the left pane, select **Storage Nodes**.
6. Right-click the storage node for the device and select **Configure All Libraries**.
7. On the **Provide General Configuration Information** window, leave **SCSI/NDMP** selected and click **Next**.
8. On the **Select Target Storage Nodes** window, perform either of the following steps:
 - Select the storage node for the library.
 - Click **Create a new Storage Node**.
9. Click **Start Configuration**.
10. Click **Finish**.

Labeling media

You can label tapes from the NMC GUI.

About this task**Procedure**

1. Log in to the NMC GUI as an administrator of the NetWorker server.
2. On the taskbar, click the **Enterprise** icon .
3. In the navigation tree, highlight a host:
 - a. Right-click **NetWorker**.
 - b. Select **Launch Application**. The **NetWorker Administration** window appears.
4. On the taskbar, click the **Devices** button .
5. In the navigation tree view, expand **Libraries** and highlight a library, or select **Devices**.
6. In the **Device list**, right-click a device and select **Label**.

Scheduling backups

Perform scheduled backups to automatically backup client data on an ongoing basis. Data protection policies enable you to define the client resources, schedule, and other settings for the backup. The client resources and backup storage resources must also be configured.

Procedure

1. Configure the backup storage resources:

- a. Configure the storage node that will own the backup storage devices.
- b. Configure the backup storage device.
- c. Create a label template for labeling volumes, or use one of the preconfigured label templates.
- d. Create media pools for sorting and storing backup data.

[Backup Target](#) on page 75 provides more information on configuring backup storage resources.

2. Configure one or more client resources for each client computer by using either the **Client Backup Configuration Wizard** or the **Client Properties** dialog box.

When you configure a client resource, you specify backup settings for the client, including:

- The save sets for the client, which define the data to back up on the client.
- Whether to automatically restart failed backups from a known good point, which is called checkpoint restart.
- Whether to bypass the storage node and send backup data directly to AFTD or DD Boost storage devices, which is called Client Direct.
- Directives that control how the NetWorker server processes files and directories during the backup.
- Probe resources for probe-based backups, where the NetWorker server probes the client for a user-defined script before the backup starts.
- Whether to back up each save set for the client by using multiple parallel save streams.
- Backup command customizations.

[Client resources](#) on page 483 provides more information on configuring client resources.

3. Configure a data protection policy for scheduled backups:

- a. Create a group to define the client resources to back up.

The type of group that you create depends on the type of backup that you are performing:

- Create a client group or dynamic client group for a traditional backup or a server backup.
- Create a VMware group to back up virtual machines or VMDKs.
- Create a NAS device group to perform snapshot backups on NAS devices.

- b. Create a policy.

Policies provide a container for the workflows, actions, and groups that support and define the backup.

- c. Within the policy, create a workflow.

Workflows define the start time for a series of actions, the order of actions in a sequence, and the group of client resources to back up.

- d. Create a backup action.

When you create a backup action, you define the following settings:

- The type of backup to perform each day.
- The destination storage node and media pool.
- The retention setting for the backup, which specifies how long to retain the backup data.

e. (Optional) Create other actions for the workflow.

Actions that you may want to include in a backup workflow include:

- Check connectivity to verify connectivity between the NetWorker server and the client computer.
- Probe to probe a NetWorker client for a user-defined script before the backup starts.
- Clone to automatically clone the save sets that result from the backup.

[Data Protection Policies](#) on page 205 provides more information on configuring groups, policies, workflows, and actions.

Viewing failed backups

You can view the details for failed NetWorker backups.

Procedure

1. Log in to the NMC GUI as an administrator of the NetWorker server.
2. On the taskbar, click the **Enterprise** icon .
3. In the navigation tree, highlight a host:
 - a. Right-click **NetWorker**.
 - b. Select **Launch Application**. The **NetWorker Administration** window appears.

4. Click **Monitoring** .

The **Monitoring** window displays four windows panes. The **Log** pane provides a summary of NetWorker server events. The **Policies** pane displays all configured policies on the NetWorker server. To view details information about the status of the actions in a workflow, expand the policy, right-click the workflow, and select **Show Details**.

Using nsrlogin for authentication and authorization

When you configure the NetWorker Authentication Service to use LDAP/AD authentication, you modify the **External Roles** attribute in the **User Group** resource to assign privileges to LDAP and AD users. As a result, NetWorker command line operations and NetWorker module operations might fail due to insufficient privileges. To resolve this issue, use the `nsrlogin` command to contact the NetWorker Authentication Service and authenticate a user. When user authentication succeeds, the NetWorker Authentication Service issues a token to the NetWorker host for the user, which provides CLI operations with token-based authentication until the token expires.

Before you begin

Ensure that the user that the NetWorker Authentication Service validates has the appropriate User Group privileges to run the CLI commands.

About this task

Perform the following steps on a NetWorker Client on which you initiate the CLI commands, or the requesting host.

Procedure

1. To validate a user and generate a token for the user, use the `nsrlogin` command:

```
nsrlogin [-s NetWorker_server] [-H authentication_host] [-P port] [-t tenant] [-d logindomain] -u username [-p "password"]
```

where:

- `-s NetWorker_server`—Specifies the name of the NetWorker Server. Use this option when you use the `nsrlogin` command on a NetWorker host that is not the NetWorker Server.
- `-H authentication_host`—Specifies the name of the NetWorker Authentication Service host. Use this option when you use the `nsrlogin` command on a NetWorker host that is not the NetWorker Server. This option is only required when you do not use the `-s` option.
- `-P port`—Specifies the NetWorker Authentication Service port number. Use this option when you do not use the `-s` option and when the NetWorker Authentication Service does not use the default port number 9090 for communications.
- `-t tenant`— Specifies the tenant name that the NetWorker Authentication Service should use to verify the username and password. When you omit this option, NetWorker Authentication Service uses the Default tenant to verify the user credentials.
- `-d logindomain`—Specifies the domain name that the NetWorker Authentication Service should use to verify the username and password with an external authentication authority. When you omit this option, the NetWorker Authentication Service uses the local user database to verify the user credentials.
- `-u username`—Specifies the username that the NetWorker Authentication Service should validate to generate a token.
- `-p "password"`—Specifies the password that the NetWorker Authentication Service should use to verify the username. If you do not specify the password, the `nsrlogin` command prompts you to provide the password.

For example, to generate a token for user *Konstantin* in the *iddomain* domain and the *idd* tenant, type the following command:

```
nsrlogin -s bu-idd-nwserver2 -d idddomain -u Konstantin -p "1.Password"
```

Authentication succeeded.

When the NetWorker Authentication Service successfully validates the user, the service issues an authentication token to the requesting host.

2. At the command prompt, type the NetWorker command.

If the validated user does not have the appropriate privileges to run the command, an error message appears or the command does not return the expected result. For example, when you try to perform an operation with a user

account that does not have the required privilege, a message similar to the following appears:

```
Permission denied, user must have the 'Operate NetWorker'
privilege'.
```

Results

The CLI command uses the authenticated token, until the token expires. By default the token expiration period is 480 minutes or 8 hours. When the token expires and the user tries to run a CLI command, the command fails with a permissions error and a message similar to the following appears to indicate that the token has expired:

```
Security token has expired
```

To resolve this issue, run the `nsrlogin` command again to generate a new authenticated token.

Note: To revoke the user token and enable the CLI commands to use the **Users** attribute in the Usergroups resources to authenticate users, use the `nsrlogout` command. The `nsrlogout` UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrlogout` command.

Performing a manual backup

Perform a manual backup of a file or folder, to test the NetWorker installation. The procedure to perform a manual backup is different on Windows and UNIX.

Performing a manual backup on Windows

Use the NetWorker User program to perform a manual backup Windows. The NetWorker User program provides a graphical interface to perform manual backups.

About this task

Procedure

1. On a NetWorker client, start the **NetWorker User** program.
2. In the **Change server** window, select or type the name of the NetWorker server.
3. In the **Source** and **Destination client** windows, select the current NetWorker client.
4. Click **Backup**.
5. In the left pane of the **Backup** window, click the appropriate directory folder.
6. Select a file or directory file to back up in one of the following methods:
 - Select the directory or file and click **Mark**. To clear an item, click **Unmark**.
 - Right-click the directory or file.

When you mark a directory or file for backup, a check mark appears next to that item.
7. Click **Start**.

The Backup Status window displays the progress of the backup. When the NetWorker server has successfully finished the backup, this message appears:

```
Backup completion time: 2-15-07 3:27p
```

If the backup fails, then:

- Review the NetWorker `daemon.raw` log file on both the NetWorker server and client hosts. Use the `nsr_render_log` program to review the log file in a readable format. The *NetWorker Command Reference Guide* describes how to use the `nsr_render_log` program.

The location of the `daemon.raw` file is different on Windows and UNIX:

- On Windows, the log file appears in the `C:\Program Files\EMC NetWorker\nsr\logs` directory.
- On UNIX, the log file appears in the `/nsr/logs` directory.
- To determine the cause, refer to the Troubleshooting chapter.
- Review the operating system log files (Application event log on a Windows client) for more information.

Performing a manual backup on UNIX

Use the `save` program to perform a manual backup from the command prompt.

For example, to back up `/tmp/myfile.txt` to a server called `jupiter`, type:

```
save -s jupiter /tmp/myfile.txt
```

The UNIX man pages describe how to use the `save` program.

CHAPTER 3

Backup Target

This chapter contains the following topics:

- [Label templates](#)..... 76
- [Media pools](#)..... 83
- [Storage nodes](#)..... 98
- [Disk storage devices](#)..... 107
- [Libraries and silos](#)..... 130
- [File type devices](#)..... 190
- [Stand-alone devices](#)..... 191
- [Labeling volumes](#)..... 196
- [Troubleshooting devices and autochangers](#)..... 197

Label templates

The NetWorker server creates a unique label for each volume by applying a label template. This section describes how label templates and media pools are used to sort, store, and track data on media volumes.


Using label templates

The NetWorker server selects the media pool to which a given set of data is written. A volume is associated with a media pool by its volume label.

The contents of the volume label follow rules that are defined in a specific label template. You then associate a label template with a specific media pool in the Media Pool resource. If you do not associate data with a specific media pool, the NetWorker server uses the preconfigured Default media pool and corresponding Default label template.

The following figure illustrates how a media pool configuration uses its associated label template to label a volume. For the label template name to appear as a choice in the Media Pool resource, you must configure a label template before configuring the associated media pool.

Figure 10 Labeling a volume by using a label template



How the NetWorker server uses volume labels

A volume label is a unique internal code, applied by the NetWorker server, that initializes the volume for the server to use and identifies a storage volume as part of a specific pool. [Using media pools](#) on page 83 provides more information about pools. Labeling a volume provides a unique name for tracking and recognizing the media, as well as references to volume labels in the records stored in the media database. The NetWorker server uses the media database records to determine which volumes are needed for backing up or recovering data.

When NetWorker labels a volume, the label operation performs the following actions:

1. Verifies that the volume is unlabeled.

2. Labels the volume with the name specified in the **Volume Name** attribute by using one of the following:
 - The next sequential label from the label template that is associated with the chosen pool.
If a recyclable volume from the same pool is relabeled, the volume label name and sequence number remain the same, but access to the original data on the volume is destroyed. The volume becomes available for new data.
 - An override volume name that was entered by the user.

Preconfigured label templates

The NetWorker server contains these preconfigured label templates, which correspond to the preconfigured media pools:

- Archive
- Archive clone
- Data Domain Default
- Data Domain Default Clone
- DD Cloud Tier Default Clone
- Default
- Default clone
- Full
- Indexed archive
- Indexed archive clone
- NonFull
- Offsite
- PC archive
- PC archive clone
- Two Sided

Label templates have multiple fields separated by periods. The first field represents the name of the NetWorker server and the final field contains a number to allow for expansion of the media pool. The number range from 001 to 999. For example:

```

mars.001
jupiter.054
jupiter.archive.197
  
```

Guidelines for completing Label Template attributes

There are certain guidelines to keep in mind when completing the attributes for a Label Template resource. The following table describes how to complete the key attributes for this resource.

Table 18 Key label template attributes

Attribute	Guidelines
Name	Keep the label name consistent with the media pool name, so that the label name reflects how the data is organized. For example, a label template

Table 18 Key label template attributes (continued)

Attribute	Guidelines
	named "AcctFull" would identify volumes that belong to a media pool called "Accounting Full." Do not use these characters in label template names: / \ * ? [] () \$! ^ ; ' " ' ~ < > & { } : - . _
Fields	A label template is made up of one or more fields. Each field, or component, provides a layer of specificity to your organizational structure. There can be any number of components, but it is best to keep the template simple with as few as necessary. The label cannot exceed 64 characters. You can use four types of components: • Range of numbers (for example, 001-999) • Range of lowercase letters (for example, aa-zz) • Range of uppercase letters (for example, AA-ZZ) • Character string (for example, Accounting) Each range includes a start value, a dash (-), and an end value. The start value and the end value must have the same number of characters. For example, use 01-99 (not 1-99) or aaa-zzz (not aa-zzz). The order in which you enter each component of the Field attribute is important. The NetWorker Server applies each component in a left-to-right order, starting with the first one entered.
Separator	Choose the symbol to appear between component entries. Use the period, dash, colon, or underscore to separate each component of the label template. If label components do not have separators (for example, AA00aa), the labels can be difficult to read.
Next	Choose the next sequence number to write on the label that the NetWorker Server places on a volume (according to the template). • To force a label to start the label scheme at a particular point, type a start label value. The server continues to generate labels from that point on, according to the rules of the template. • To have the NetWorker Server generate the first label, leave this attribute blank. When the NetWorker Server recycles a storage volume, the volume label does not change as long as the volume remains in the same media pool. That is, if a storage volume labeled "Dev.006" is recycled, it retains the volume label "Dev.006" and does not receive a new label with the next sequence number.

The following table lists examples of number sequences for volume labels.

Table 19 Examples of number sequences for volume labels

Type of components	Fields	Number sequence result	Total number of labels
Range of numbers	001-100	001, 002, 003,...100	100
Character string	SalesFull	SalesFull.001,...SalesFull.100	100
Range of numbers	001-100		
Range of lowercase letters	aa-zz	aa.00,...aa.99,	67,600 (262 times 102)
Range of numbers	00-99	ab.00,...ab.99,	
		ac.00,...ac.99,	
		:	
		az.00,...az.99,	
		ba.00,...ba.99	
		:	
		zz.00,...zz.99	

The label template should allow for expansion of the backup media storage system. For example, it is better to create a template for 100 tapes and not use all of them, than it is to create a template for only 10 tapes and run out of labels. When the server reaches the end of the template numbering sequence, it wraps to the starting value. For example, after zz.99 (used for the 67,600th label), the next label the server uses is aa.00 for label 67,601.

Note: When the NetWorker server recycles a volume, the volume label does not change if the volume remains in the same media pool. That is, if a volume labeled Dev.006 is recycled, it will retain the volume label Dev.006 and will not receive a new label with the next sequence number. The original data on the volume, however, will be overwritten by the new data.

Naming label templates

The NetWorker server is packaged with preconfigured label templates that correspond to the preconfigured media pools. If you choose to create the templates, you can include any number of components in the Fields attribute. However, it is best to keep the template simple with as few components as necessary for your organization.

For example, if you create a label template for an accounting department, you can customize the label template in several ways, depending on the size of the storage system and media device capabilities.

The following table illustrates several ways you can use components to organize labels.

Table 20 Using label template components

Type of organizational structure	Components	Separator	Resulting volume labels
Sequential	AcctFull	period	AcctFull.001

Table 20 Using label template components (continued)

Type of organizational structure	Components	Separator	Resulting volume labels
	'001-100		(100 total labels)
Storage oriented (for example, 3 storage racks with 5 shelves each, each shelf holding 100 tapes)	1-3 1-5 001-100	dash	1-1-001 This label is for the first tape in rack 1 on shelf 1. (1,500 total labels)
Two-sided media (for example, optical devices)	AcctFull 000-999 a-b	underscore	AcctFull_000_a (side 1) AcctFull_000_b (side 2) (2,000 total labels)

Tips for labelling

Naming schemes vary from site to site. One way is to name the volumes with the name of the NetWorker server followed by a three-digit number, for example:

```
jupiter.001
```

Consider that the simpler a convention is, the easier it can be understood by operators and administrators.

The maximum length for a volume name is 63 characters. With advanced file type devices (adv_file), the maximum length is 60 characters.

Each volume should have a physical (adhesive) label attached to it. Since the NetWorker server keeps track of the backups and which volumes they are on, you can name the volumes with any convenient name. For example, you can label your volumes 1, 2, 3, or Monday.1, Tuesday.1, Wednesday.1. You can assign a volume any name as long as each one is unique.

The adhesive label on the volume should match the name generated by NetWorker. For example, if you physically label a volume mars.1, its NetWorker name should also be mars.1.

Working with label templates

This section explains how to create, edit, copy, and delete label templates.

Creating a label template

When creating a label template, consider the labeling guidelines for the Name, Fields, Separator, and Next components.

Procedure

1. In the **Administration** window, click **Media**.

2. In the expanded left pane, select **Label Templates**.
3. From the **File** menu, select **New**.
4. Enter the components for the label template:
 - **Name:** The name of the new label template.
 - **Comment:** Any user-defined description or explanatory remarks about the label.
 - **Fields:** A list of label components.
 - **Separator:** The character to be inserted between label components. If no symbol is selected, the components will have no separators, such as hostarchive[001-999].
 - **Next:** (Optional) Enter the next label to be generated by the template.
5. Click **OK**.

Editing a label template

You cannot change the name of a label template. However, to change an individual label name, delete the existing name in the Next text box, and type a new name.

Procedure

1. In the **Administration** window, click **Media**.
 2. In the expanded left pane, select **Label Templates**.
 3. In the right pane, perform one of the following tasks:
 - To modify multiple attributes in a single configuration resource by using the **Label Template Properties** window, right-click the staging configuration and select **Properties**.
 - To modify a specific attribute that appears in the resource window, place the mouse in the cell that contains the attribute that you want to change, then right-click. The menu displays an option to edit the attribute. For example, to modify the **Comment** attribute, right-click the resource in the **Comment** cell and select **Edit Comment**.
-  **Note:** To modify a specific attribute for multiple resources, press and hold the **Ctrl** key, select each resource, and then right-click in the cell that contains the attribute that you want to change. The menu displays an option to edit the attribute.
4. Make any required changes, then click **OK**.

Copying a label template

Procedure

1. In the **Administration** window, click **Media**.
2. In the expanded left pane, select **Label Templates**.
3. In the right pane, select the label template to copy.
4. From the **Edit** menu, select **Copy**. The **Create Label Template** dialog box appears, containing the same information as the label template that was copied, except **Name** attribute.
5. In the **Name** attribute, type the name for the new label template.
6. Edit any other attributes as appropriate, and click **OK**.

Deleting a label template

You cannot delete a preconfigured label template or a label template that is in use.

Procedure

1. In the **Administration** window, click **Media**.
2. In the expanded left pane, select **Label Templates**.
3. In the right pane, select the label template to delete.
4. From the **File** menu, select **Delete**.
5. When prompted, click **Yes** to confirm the deletion.

Setting up a label template to identify volumes

If you are not using tapes with barcode labels, and the Match Bar Code Labels attribute is not enabled for the Library resource, then every backup volume requires a unique label for identification. The NetWorker server creates a unique label for each volume by applying a label template.

About this task

Procedure

1. From the **Administration** window, click **Media**.
2. In the expanded left pane, select **Label Templates**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a name for the label template.
5. In the **Comment** attribute, type a description for the label template.
6. In the **Fields** attribute, type the label's components. Place each label component on a separate line. The template can use any or all of these components, although at least one range component must be added:
 - Range of numbers—For example, 001-999
 - Range of lowercase letters—For example, aa-zz
 - Range of uppercase letters—For example, AA-ZZ
 - Character string—For example, Accounting
 - Ranges of numbers or letters change incrementally with each new label. For example:
 - First label: Accounting.001
 - Second label: Accounting.002
 - Third label: Accounting.003
7. Select a **Separator** and click **OK**. If no symbol is selected, the components will have no separators (for example, Accounting001).
8. Click **OK**.

Media pools

NetWorker uses media pools and volume labels to sort backup and clone data on media.

Media is a specific collection of volumes to which the NetWorker server writes data. For example, a tape volume or a Data Domain device. A volume is identified with a unique label based on user configurable label templates.

Media pools act as filters that tell the NetWorker server which backup volumes should receive specific data. The NetWorker server uses media pools along with label templates to track what data is on which specific volume. When you use a barcode-enabled tape library, the NetWorker server uses media pools along with the volume barcode Labels to track which data is on a specific volume.

 **Note:** NetWorker does not use media pools for backup and clone operations to deduplication devices.

Using media pools

Action resources contain an attribute that defines the media pool to which NetWorker should send the backup or clone data.

When a backup or clone action starts, the NetWorker server then checks if a correctly labeled volume for that media pool is mounted on a storage device. If a correctly labeled volume is mounted on a storage device, the NetWorker server writes data to the volume. If there is no correctly labeled volume mounted on a storage device, the NetWorker server generates a request to mount a volume that is labeled for the pool, and waits until an operator or an autochanger mounts an appropriate volume.

Preconfigured media pools

NetWorker provides you with the following preconfigured media pools.

Table 21 Preconfigured media pools

Pool name	Description
Archive	Receives archived backup data when you use the <code>nsrchive</code> command and use <code>-b</code> option to specify the pool name. NetWorker does not assign a retention policy to an archived save set, and the save set never expires. When you enable Archive Services on a client resource and you configure the backup action to send data to the Archive pool, NetWorker does not write information about the archive save set to the client file index for the client.
Archive Clone	Receives the clone copy of archived backup data. when you use the <code>nsrclone</code> command with <code>-b</code> option to specify the pool name. NetWorker does not assign an expiration date to the clone copy of an archive save set. NetWorker does not write information about

Table 21 Preconfigured media pools (continued)

Pool name	Description
	the clone save set to the client file index for the client.
Default	Receives backup data in the following configurations: - When you select the Default pool in the Pool attribute of a backup action resource. - When you use <code>save</code> command to run a manual backup and do not use the <code>-b</code> option to specify a specific backup pool. - When NetWorker performs an action on a client and you define the following configuration attributes: - In the Action resource, the option Client Override Behavior is set to Client Can Override. - In the Client resource, you select the Default pool in the Pool attribute.
Data Domain Default	Receives backup data to DD Boost devices only. <i>NetWorker Data Domain Boost Integration Guide</i> provides more information about how to use NetWorker with DD Boost devices.
Data Domain Default Clone	Receives clone data to DD Boost devices only. <i>NetWorker Data Domain Boost Integration Guide</i> provides more information about how to use NetWorker with DD Boost devices.
DD Cloud Tier Default Clone	Receives clone data on DD Cloud Tier devices only. <i>NetWorker Data Domain Boost Integration Guide</i> provides more information about how to use NetWorker with DD Cloud Tier devices.
Default Clone	Receives clone data in the following configurations: - When you select the Default Clone pool in the Pool attribute of a clone action resource. - When you use <code>nsrclone</code> command to run a manual backup and do not use the <code>-b</code> option to specify a specific backup pool.
Indexed Archive	Receives archived backup data. NetWorker does not assign a retention policy to an archived save set, and the save set never expires. When you enable Archive Services on a client resource and you configure the backup action to send data to the Indexed

Table 21 Preconfigured media pools (continued)

Pool name	Description
	Archive pool, NetWorker writes information about the archive save set to the client file index for the client.
Indexed Archive Clone	Receives the clone copy of an indexed archive. NetWorker does not assign an expiration date to the clone copy of an archive save set. NetWorker does not write information about the clone save set to the client file index for the client.

Changes to the Client and Pool resources after migration

NetWorker uses a number of attributes that are defined in multiple resources to determine which pool receives the data that is generated by an action task, and how NetWorker backs up the data. The migration process preserves the values that are defined for the attributes and introduces new attributes in the Action resource.

NetWorker provides the following attributes, which work together to determine how NetWorker manages a backup and determines which device to use to receive the backup data:

- Client resource—**Pools**, **Retention**, **Save set**, and **Level** attributes on the **General** tab of the **Client Properties** window. The migration process retains the values in these legacy attributes.
 - Note:** The **Modify Client** wizard does not display the **Pools**, **Retention**, **Save set**, and **Level** attributes.
- Action resource—**Destination Pool** and **Retention** attributes on the **Specify the Backup Options** and **Specify the Clone Options** wizard windows. The backup levels are defined for the action schedule on the **Specify the Action Information** wizard window.
- Pool resource—**Clients**, **Save sets**, and **Retention policy** attributes on the **Legacy** tab. The values that appear in these attributes were defined in NetWorker 8.1.x and 8.2.x. After the migration completes, the NetWorker 19.1 server retains the values and these legacy attributes become read-only. You cannot modify the values in these fields after migration.

The Action resource includes an attribute that is called **Client Override Behavior**. The value that is selected for this attribute determines which resource attribute has precedence over the attributes in other resources that determine the same behavior. By default, the migration process enables **Legacy Backup Rules** on an Action resource. **Legacy Backup Rules** allow NetWorker to use the values during the pool selection criteria process.

- Note:** By default, the **NetWorker Administration** window does not show the legacy attributes. To view the legacy attributes in the **Client Properties** window, go to the **View** menu and select **Diagnostic Mode**.

Pool selection criteria

It is recommended that you use the configuration settings in an Action resource to determine which pool received backup data. NetWorker provides you with the ability to configure a Pool attribute in the client resource, which can override the value

defined in the Action resource. Additionally, the Pool resource contains 8.2.x legacy attributes that provide you with the ability to define backup data criteria for the pool.

How and when NetWorker uses the attributes values defined in the Pool, Action, and Client resources to determine which backup pool will receive data depends on the value that you select in the **Client Override Behavior** attribute of the Action resource:

- **Client Can Override**—The value in **Pool** attribute of the client resource takes precedence over the **Destination pool** value that is defined in the Action resource. NetWorker does not use the values that are defined in the **Client**, **Save set**, and **Levels** attributes of the Pool resources when deciding which pool receives backup data for a client.
- **Client Can Not Override**—The value defined **Destination Pool** attribute in the Action resource takes precedence over the value that is defined in the Pool attribute of the Client resource. NetWorker does not use the values that are defined in the **Client**, **Save set**, and **Levels** attributes of the Pool resources when deciding which pool receives backup data for a client.
- **Legacy Backup Rules**—Enabled for migrations only. NetWorker uses the values that are defined in the **Client**, **Save set**, and **Levels** attributes of the pool resource to determine which pool receives backup data from a client. The values that are defined in the **Client**, **Save set**, and **Levels** of the pool resource take precedence over the **Destination Pool** value that is defined in the Action resource, and the **Pool** value that is defined in the Client resource.

 **Note:** You cannot modify the legacy attributes in the migrated Pool resources.

The following table summarizes how NetWorker determines which pool receives the backup data, which is based on the configuration of the Action, Client, and Pool resource attributes.

Table 22 Determining which pool receives backup data

Client Override Behavior (Action)	Destination pool (Action)	Pool (Client)	Legacy criteria attributes (Pool)	Pool that receives the data
Client Can Override	Defined	Defined	Defined and criteria matches	Pool defined in Client resource
Client Can Override	Defined	Undefined	Defined and criteria matches	Pool defined in Action resource
Client Cannot Override	Defined	Defined	Defined and criteria matches	Pool defined in Action resource
Legacy Backup Rules	Defined	Undefined	Defined and criteria matches	Pool that matches legacy criteria
Legacy Backup Rules	Defined	Defined	Defined and criteria matches	Pool that matches legacy criteria
Legacy Backup Rules	Defined	Undefined	Undefined or no matches	Default

Example 1 Client Can Override is enabled

A Protection group contains two clients, *SQL_clnt* and *Exchange_clnt*. The workflow that is associated with the protection group contains a backup action.

Example 1 Client Can Override is enabled (continued)

- Backup action configuration:
 - Destination Pool=*App_backups*
 - Schedule=Daily full backup
 - Client Override Behavior=Client Can Override
- The Pool attribute that is defined for the *SQL_clnt* client resource is *SQL_backups*.
- The Pool attribute for *Exchange_clnt* is *Exchange_backups*.
- The Full level is enabled in the Levels attribute of a pool resource named *Backups*.

In this example, NetWorker sends the backup data for *Exchange_clnt* to *Exchange_backups*, the pool that is defined in the backup action. NetWorker sends the backup data for *SQL_clnt* to the pool defined in the client resource, *SQL_backups*.

Example 2 Example: Client Can Override is enabled

A Protection group contains two clients, *SQL_clnt* and *Exchange_clnt*. The workflow that is associated with the protection group contains a backup action.

- Backup action configuration:
 - Destination Pool=*App_backups*
 - Schedule=Daily full backup
 - Client Override Behavior=Client Can Override
- The Pool attribute that is defined for the *SQL_clnt* client resource is *SQL_backups*.
- The Pool attribute for *Exchange_clnt* is not defined.
- The Full level is enabled in the Levels attribute of a pool resource named *Backups*.

In this example, NetWorker sends the backup data for *Exchange_clnt* to *App_backups*, the pool that is defined in the backup action. NetWorker sends the backup data for *SQL_clnt* to the pool defined in the client resource, *SQL_backups*.

Example 3 Client Cannot Override is enabled

A Protection group contains two clients, *SQL_clnt* and *Exchange_clnt*. The workflow that is associated with the protection group contains a backup action.

- Backup action configuration:
 - Destination Pool=*App_backups*
 - Schedule=daily full backup
 - Client Override Behavior=Client Cannot Override
- The Pool attribute that is defined for the *SQL_clnt* client resource is *SQL_backups*.
- The Pool attribute for *Exchange_clnt* is *Exchange_backups*.

Example 3 Client Cannot Override is enabled (continued)

- The Full level is enabled in the Levels attribute of a pool resource named *Backups*.

In this example, NetWorker sends the backup data for *SQL_clnt* and *Exchange_clnt* to *App_backups*, the pool that is defined in the backup action.

Example 4 Legacy Backup Rules is enabled

A Protection group contains two clients, *SQL_clnt* and *Exchange_clnt*. The workflow that is associated with the protection group contains a backup action.

- Backup action configuration:
 - Destination Pool=*App_backups*
 - Schedule=daily full backup
 - Client Override Behavior= Legacy Backup Rules
- The Pool attribute that is defined for the *SQL_clnt* client resource is *SQL_backups*.
- The Pool attribute for *Exchange_clnt* is not defined.
- The Full level is enabled in the Levels attribute of a pool resource named *Backups*.

In this example, NetWorker sends the backup data for *SQL_clnt* and *Exchange_clnt* to *Backups*, the pool that matches the level Full backup criteria.

Example 5 Legacy Backup Rules is enabled

A Protection group contains two clients, *SQL_clnt* and *Exchange_clnt*. The workflow that is associated with the protection group contains a backup action.

- Backup action configuration:
 - Destination Pool= *App_backups*
 - Schedule=daily full backup
 - Client Override Behavior= Legacy Backup Rules
- The Pool attribute that is defined for the *SQL_clnt* client resource is *SQL_backups*.
- The Pool attribute for *Exchange_clnt* is not defined.
- The manual level is enabled in the Levels attribute of a pool resource named *Backups*.

In this example, NetWorker sends the backup data for *SQL_clnt* and *Exchange_clnt* to the *Default* pool because a pool does not exist with legacy attributes that match the configuration for the backup data.

Matching the pool criteria with Legacy Backup Rules enabled

After a migration and configuring media pools, data generated by an action might match the criteria for more than one media pool configuration. For example, if you configure one media pool to accept data from a client that is called *mnd.corp.com*, and you configure another media pool to accept data from all full backups, NetWorker uses

other criteria to determine which pool of volumes receives the data from a full backup of the mnd.corp.com client.

The NetWorker Server uses the following media pool selection criteria:

1. Groups attribute (highest precedence)
2. Clients attribute
3. Save sets attribute
4. Levels attribute (lowest precedence)

When data matches the attributes for two media pools, for example, Client and Level, the data is written to the media pool specified in the Client attribute. For example, in the case where the data from the client matched the criteria for two different media pools, the data is routed to the media pool that accepts data from the mnd.corp.com client.

The following table details the hierarchy that the NetWorker Server uses to determine media pool selection when a conflict arises. For example, the media pool criteria for Groups takes precedence over the media pool criteria for Clients, Save sets, and Levels. If data does not meet the criteria for any customized pool, NetWorker writes the data to the Default media pool.

Table 23 NetWorker hierarchy for resolving media pool conflicts

Precedence	Groups attribute	Clients attribute	Save sets attribute	Levels attribute
Highest	x	x	x	x
	x	x	x	
	x	x		x
	x	x		
	x		x	x
	x		x	
	x			x
	x			
		x	x	x
		x	x	
		x		x
		x		
			x	x
			x	
Lowest				x

Working with media pools

This section explains how to edit, copy, delete, and create media pools.

Creating a media pool

Perform the following steps to create a new media pool.

Before you begin

Perform either of the following:

- If the Match Bar Code Labels attribute is not used for the Library resource, create a label template for the media pool.
- Determine a preconfigured label template to use for the media pool.

About this task

Procedure

1. In the **Administration** window, click **Media**.
2. In the left pane, select **Media Pools**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a name for the media pool.
A media pool is associated with a label template. Use a name that clearly associates the media pool with the corresponding label template.
5. In the **Comment** attribute, type a description of the media pool.
6. Leave the **Enabled** attribute selected.
7. For the **Pool Type** attribute, select the media pool type.
 - Backup—Select this type to configure the pool to receive backup data.
 - Backup clone—Select this option to configure the pool to receive a clone copy of backup data.
 - Archive—Select this type to configure the pool to receive archive data.
 - Archive clone—Select this option to configure the pool to receive a clone copy of archive data.
8. In the **Label Template** attribute, select the matching label template.
9. In the **Data Source** attribute, select the backup groups that are eligible to back up to this media pool.
10. (Optional), on the **Selection Criteria** tab, configure the following options:
 - **Devices**—Select the devices on which NetWorker can mount volumes for this pool.
 - **Media type required**—Select which device type NetWorker can use to label volumes for this pool. You cannot use this attribute when you select an option in the **Media type preferred** attribute.
 - **Media type preferred**—Select the device type that NetWorker should use first to label a volume for this pool. You cannot use this attribute when you select an option in the **Media type required** attribute.

 **Note:** When you do not configure the **Media type required** or **Media type preferred** attribute, you can write data across several volumes of different media types (for example, magnetic disk and tapes), if the volumes mounted on the storage devices have the appropriate label associated with the media pool.

11. On **Configuration** tab, configure the following options:

Attribute	Definition
Auto Media Verify	Select this attribute to perform automated media verification while data is written to a volume labeled for this media pool. Auto media verification provides more information.
Max parallelism	Increase the value to define the maximum number of simultaneous save streams that NetWorker writes to each device in the pool. The default value for this attribute is 0, which means that the attribute has no effect on other parallelism settings. When you set the Max parallelism attribute to 1, a prolonged delay might occur between the backup of save sets. To resolve this issue, increase the Max parallelism attribute for the pool resource. However, when you increase the pool parallelism value, the time to recover data on the volume increases. Note: For AFTD and DD Boost devices, the Max nsrmmmd count attribute value for a device affects the Max parallelism attribute. For example, consider an AFTD device (AFTD_1) that has a Max sessions attribute value of 20 and a Max nsrmmmd value of 4. Now suppose a backup pool with a Pool parallelism attribute of 1 selects AFTD_1. The total number of save sessions that NetWorker can start for AFTD_1 is 4, one for each nsrmmmd process. Tape and FTD devices can only spawn one nsrmmmd process at a time, so if the previous example used a tape device, then the total number of save sessions would be 1.
Recycle from other pools	Select this option to enable NetWorker to use expired volumes that are labeled for other media pools in this pool that have the Recycle to other pools attribute enabled, when the NetWorker server does not have access to blank volumes or volumes eligible for reuse and assigned to this pool.
Recycle to other pools	Select this option to enable NetWorker to use expired volumes that are labeled for this media pool in other pools that have the Recycle from other pools attribute enabled, when the NetWorker server does not have access to blank volumes or volumes eligible for reuse and assigned to the other pool.
Recycle start	Defines the time to start the automatic relabel process each day. By default this attribute is empty and the automatic relabeling of recyclable volumes is not done. Use the format <i>HH:MM</i> . Automatically relabeling volumes in a media pool provides more information.
Recycle interval	Defines the interval between two starts of the automatic relabel processes. The default value is 24:00. Use the format <i>HH:MM</i> .

Attribute	Definition
Max volumes to recycle	Defines the maximum number of recyclable volumes that NetWorker can relabel during each automatic relabel process. The default value is 200.
Recycle start now	Select this attribute to start the automatic relabel process of recyclable volumes for this pool immediately after you create the pool. The default value is No.
Store index entries	For archive pools only. Select this attribute to configure an archive pool that creates client file index entries for the archive save sets. Clear this option to configure an archive pool that will not create client file index entries for the archive save sets.
Worm pool/ Create DLTWORM	Supported WORM and DLTWORM tape drives provides more information about how to create Worm pools.

12. Optionally, on the **Restricted Data Zones** tab, from the restricted datazone list, select the restricted datazone in which to add the pool.
13. Click **OK**

If any of the settings for a new media pool match an existing media pool, this message appears:

```
Pool(s) pool_name has overlapping selection criteria.
```

If this message appears, review the media pool configuration and modify any overlapping criteria.

14. If you did not select a label template when you create the media pool, a message appears that tells you that NetWorker creates a label template for the media pool, click **OK**.

Auto media verification

If the Auto Media Verify attribute is enabled, the NetWorker server verifies data written to tape volumes from this media pool. This attribute does not apply to AFTD, file type and Data Domain devices.

Data is verified by repositioning the tape volume to read a portion of the data previously written to the media. The data read is compared to the original data written. This feature does not verify the entire length of the tape.

If the data read matches the data written, verification succeeds.

Media is verified when the following occurs:

- A volume becomes full while saving and it becomes necessary to continue on to another volume.
- A volume goes idle because all save sets being written to the volume are complete.

When a volume fails verification, it is marked full so that the server will not select that volume for future saves. The volume remains full until it is recycled or a user marks it not full. If a volume fails verification while the server is attempting to switch volumes, all save sets writing to the volume are terminated.

Auto media verification should not be used to verify the integrity of the data written to the entire tape. To fully verify the data written to the tape, either restore the tape contents or clone the data.

Automatically relabeling volumes in a media pool

About this task

Automatically relabeling a recyclable volume provides the following benefits:

- You can relabel volumes outside of the backup window without the need for a scripted solution.
- NetWorker has access to appendable volumes at the time of a backup or clone, which results in faster backup and clone completion times.

Eligible volumes will not be relabeled if the volume is loaded in a device that is:

- Disabled
- In use by an `nsrmmd` process (for example, during a restore operation)
- In read-only mode
- Busy

When NetWorker automatically relabels a volume, message to the following appears in the `daemon.raw` file on the NetWorker server:

```
"num_of_volumes volumes will be recycled for pool pool_name in
jukebox jukebox_name."
```

Supported WORM and DLTWORM tape drives

NetWorker supports write-once, read-many (WORM) tape drives and media. It is able to recognize the WORM abilities of tape drives and the presence of WORM media in those drives. It also supports the creation of DLTWORM (formerly DLTIce) tapes in drives that are DLTWORM capable.

The following table describes the WORM devices that are supported by the NetWorker software. For a complete listing of supported devices, refer to the *NetWorker Hardware Compatibility Guide*.

Table 24 WORM supported devices

Device	Description
HP LTO Ultrium 3 and higher	Unique to HP Ultrium-3 and higher: • Inquiry VPD page 0xb0, byte 4 bit 0 indicates WORM capable • Read attribute # 0x0408 bit 7 to indicate WORM media present
Quantum SDLT600, DLT-S4, and DLT-V4 (SCSI and SATA)	Any drive with product inquiry data of “*DLT*” tape drive that reports WORM capability the way these drives do (“Quantum” not required in the vendor inquiry data): • Inquiry data VPD page 0xc0, byte 2, bit 0 to indicate WORM capable • Read attribute # 0x0408 bit 7 to indicate WORM media present

Table 24 WORM supported devices (continued)

Device	Description
Sony AIT-2, AIT-3, AIT-4, and SAIT	Any drive with “Sony” in the vendor inquiry data that reports WORM capability like these drives do: - Mode sense page 0x31, byte 5 bit 0 indicates WORM capable - Mode sense byte 4 bit 6 indicates WORM tape present
IBM 3592	Unique to IBM 03592: - Mode sense page 0x24, byte 7 bit 4 indicates WORM capable - Mode sense page 0x23, byte 20 bit 4 indicates WORM tape present
STK 9840A/B/C, 9940B, T10000	Any drive with STK as the vendor data that reports WORM capability like these: - Standard inquiry data byte 55 bit 2 indicates WORM capable - Request sense data byte 24 bit 1 indicates WORM tape present
IBM LTO Ultrium 3 and higher, and Quantum LTO Ultrium 3 and higher	These drives use the SCSI-3 method to report WORM capabilities, so there is not a match against any of the inquiry data. Any drive that does not match the inquiry data patterns listed above will have the SCSI-3 method applied to them: - Inquiry data VPD page 0xb0, byte 4, bit 0 indicates WORM capable - Mode sense page 0x1d, byte 2 bit 0 indicates WORM tape present Byte 4, bits 0,1: label restrictions include - 00 indicates no overwriting allowed - 01 indicates some labels can be overwritten - Byte 5, bits 0,1: filemark overwrite restrictions - 0x02: any filemark at EOD can be overwritten except for the one closest to the beginning of the tape - 0x03: any filemark at EOD can be overwritten

The WORM and DLTWORM attributes determine whether or not the NetWorker software will back up to a write once-read many (WORM) tape. You can apply these tape attributes to any pool.

Note: Various Quantum drive models (SDLT600, DLT-S4, and DLT-V4) have the ability to create WORM tapes from ordinary blank DLT tapes supported by that particular drive. You cannot recycle an existing NetWorker tape to create a DLTWORM volume without first having bulk-erased the tape. When the DLTWORM attribute is set, labeling one of these drives into a WORM pool causes the Quantum drive to make the current tape a WORM tape.

Savegroups that belong to pools that have either the WORM or DLTWORM attribute set, are considered to be WORM savegroups.

How to identify WORM media

Since WORM media cannot be reused, the tapes are uniquely identified as such so that they are only used when required. As shown in this figure, a (W) is appended to the volume names displayed in the **Volumes** window. If a volume is both read-only and WORM, an (R) is appended to the volume name.

Figure 11 Identifying WORM tapes in the NetWorker Console

Volume Name	Barcode	Used	% Used	Mode	Expiration	Location	Pool
000000	000000	0 KB	0%	appen			Default
000016	000016	0 KB	0%	appen			Default
000017	000017	0 KB	0%	appen			Default
000018	000018	0 KB	0%	appen			Default
000024	000024	0 KB	0%	appen			Default
000134(W)	000134	0 KB	0%	appen	manual	rd=aurora.A...	WORM
all2_worm.001(W)		0 KB	0%	appen	manual		worm
ameba.003		0 KB	0%	appen			Default
arch_talks_backup_12_05_2005		194 GB	97%	appen	manual		Default
berferd.001		152 GB	95%	appen	manual		worm
bobs_first_tape(R)		34 GB	17%	recyc	expired		Default
dltv4_worm_001		0 KB	0%	appen		9	Default
dltworm.001		0 KB	0%	appen	manual		worm
fatdat.ameba.001(R)		1356 MB	2%	recyc	expired	9	Default
NEGHVAR.DDS.001(R)		269 MB	0.1%	recyc	expired		Default
not_the_worm.001		15 GB	15%	appen			Default
not_worm_new.001		0 KB	0%	appen		9	Default
salt_via_polarbear.001		1091 MB	full		9/12/06		Default

Note: Since WORM tapes can only be used once, attempting to relabel a WORM tape always results in a write protection error. With the exception of pool selection and relabeling, the NetWorker software treats WORM tapes exactly the same as all other types of tape.

Determining WORM and DLTWORM capability Procedure

1. In the **Administration** window, click **Devices**.
2. Select the drive, right-click, and select **Properties**.
3. Click the **Information** tab and observe the WORM capable and DLTWORM capable attribute settings. NetWorker automatically sets these attributes and, consequently, they are read-only and cannot be changed.

Note: The WORM capable and DLTWORM capable attributes are dimmed out when the device in use is WORM capable but does not support DLTWORM (not a Quantum DTL-type drive).

Configuring WORM and DLTWORM support

The following table describes WORM and DLTWORM attributes.

About this task

Table 25 WORM/DLTWORM attributes

Attribute	Description
WORM pools only hold WORM tape	By default, the NetWorker software only allows WORM tapes into WORM pools. Deselecting this option lets you add new

Table 25 WORM/DLTWORM attributes (continued)

Attribute	Description
	(non-WORM) tapes to a WORM pool. This is useful when you need WORM functionality but do not have WORM tapes available.
WORM tapes only in WORM pools	By default, NetWorker only lets you label WORM tapes into WORM pools. Clear this option when: You do not want to segregate WORM tapes within WORM pools. A volume is needed to complete a group and a non-WORM tape is unavailable.
WORM capable	This attribute indicates that this drive supports the use of WORM media.
DLTWORM capable	This attribute indicates that this drive can create DLTWORM tapes from a blank tape.
WORM pool	This pool should hold WORM tapes (depending on the setting of “WORM pools only hold WORM tape” in the server).
create DLTWORM	If selected, before the NetWorker software labels a tape in a drive capable of creating DLTWORM volumes, NetWorker will try to convert the tape into a DLTWORM tape. If that conversion fails, the labeling for that tape will fail. If a tape drive in a pool where this attribute is set cannot create DLTWORM tapes, (that is, the tape drive is not a Quantum SDLT600, DLT-S4 or DLT-V4 tape drive, this attribute is simply ignored. Refer to the Quantum web site for information on which tapes can be converted to DLTWORM tapes. Not all firmware revisions for all of these devices support WORM operation. Check the tape drives website to make sure that your drive has up-to-date firmware.

Procedure

1. In the **Administration** window, click **Media**.
2. In the left pane, select **Media Pools**.
3. In the right pane, select the appropriate pool.
4. Right-click and select **Properties**.
5. Click the **Configuration** tab and select one of these WORM tape handling attributes:

- WORM pools only hold WORM tapes
 - WORM tapes only in WORM pools
6. Click **OK** when finished making the necessary selections.

 **Note:** If you attempt to assign a non-WORM capable drive to a WORM pool an error message is generated.

Editing a media pool

Perform these steps to edit an existing media pool.

About this task

 **Note:** You cannot change the name of a media pool. Preconfigured media pools cannot be modified.

Procedure

1. In the **Administration** window, click **Media**.
 2. In the left pane, select **Media Pools**.
 3. In the right pane, perform one of the following tasks:
 - To modify multiple attributes in a single configuration resource by using the **Media Pool Properties** window, right-click the staging configuration and select **Properties**.
 - To modify a specific attribute that appears in the resource window, place the mouse in the cell that contains the attribute that you want to change, then right-click. The menu displays an option to edit the attribute. For example, to modify the **Comment** attribute, right-click the resource in the **Comment** cell and select **Edit Comment**.
-  **Note:** To modify a specific attribute for multiple resources, press and hold the **Ctrl** key, select each resource, and then right-click in the cell that contains the attribute that you want to change. The menu displays an option to edit the attribute.
4. Make any required changes, then click **OK**.

Copying a media pool

Perform these steps to create a copy of a pool resource.

Procedure

1. In the **Administration** window, click **Media**.
2. In the left pane, select **Media Pools**.
3. In the right pane, select the media pool.
4. From the **Edit** menu, select **Copy**. The **Create Media Pool** dialog box appears, containing the same information as the media pool that was copied, except for the **Name** attribute.
5. In the **Name** attribute, type a name for the new media pool.
6. Edit any other attributes as appropriate, and click **OK**.

Deleting a media pool

About this task

You can delete a media pool only if the media database does not contain information about active volumes that are labeled for the media pool. You cannot delete a preconfigured media pool.

Procedure

1. In the **Administration** window, click **Media**.
2. In the left pane, select **Media Pools**.
3. In the right pane, select the media pool.
4. From the **File** menu, select **Delete**.
5. When prompted, click **Yes** to confirm the deletion.

Storage nodes

Storage nodes (including the NetWorker server) are host computers with attached storage devices. A storage node has the physical connection and ownership of the attached devices, but the NetWorker server maintains the client file index and media database. With the NetWorker software, client data can be routed directly to a storage node's storage devices without the data first going to the NetWorker server. A storage node may be a client of the NetWorker server, although this is not a requirement. However, the storage node must have the NetWorker client software installed.

From the NetWorker server, typical storage tasks can be performed, such as:

- Mounting and labeling volumes for the storage node devices.
- Configuring NetWorker resources associated with the storage nodes.

Only users who have the Configure NetWorker privilege can add to or change the configuration of the NetWorker server, media devices, and libraries. The *NetWorker Security Configuration Guide* provides more information.

Requirements

To operate the NetWorker software with storage nodes, certain requirements must be met.

- On UNIX systems, this software must be installed on the storage nodes. The packages must be installed in the following order:
 1. NetWorker client software
 2. NetWorker storage node software
- On Windows systems, the Storage Node Option must be installed. The Storage Node Option installs both the NetWorker client and storage node software.

Licensing

The *NetWorker Licensing Guide* provides information on NetWorker licensing support for storage nodes.

Storage node configuration

The following sections provide the procedures for configuring a NetWorker storage node.

Configuring the Linux host as a storage node

Configure a storage node host to manage the data protection activities on a host that is not the NetWorker server.

Procedure

1. Ensure that the storage node software and required enabler codes have been installed on the host.
2. In the NetWorker server **Administration interface**, click the **Devices** view.
3. From the navigation tree, right-click **Storage Nodes**, and select **New**.
The **Create Storage Node** window appears, with the **General** tab displayed.
4. Set the **Identity** attributes:
 - a. In **Name**, specify the hostname of the NetWorker storage node.
 - b. In **Type of Storage Node**, select **SCSI**.
5. In the **Status** attributes, review or set the storage node status:
 - a. **Storage node is configured** indicates whether a device has already been configured on this storage node.
 - b. **Enabled** indicates whether the storage node is available for use:
 - **Yes** indicates available state.
 - **No** indicates service or disabled state. New device operations cannot begin and existing device operations may be canceled.
 - c. **Ready** indicates whether the storage node is ready to accept device operations.
6. Set the **Device Management** attributes:
 - a. In **Max active devices**, set the maximum number of devices that NetWorker may use from this storage node in a DDS environment.
 - b. In **AFTD allowed directories**, for AFTD devices, type the pathnames of directories on the storage host where AFTDs are allowed to be created.
 - c. In **mmnds for disabled devices**, select the nsrmmd (data mover) option:
 - To start nsrmmd processes for disabled devices, select **Yes**.
 - To *not* start nsrmmd processes for disabled devices, select **No**.
 - d. In **Dynamic nsrmmds**, for AFTD or DD Boost devices, select whether nsrmmd processes on the storage node devices are started dynamically.
 - **Selected (Dynamic mode)**: NetWorker starts one nsrmmd process per device and adds more only on demand, for example, when a device's Target sessions is reached.
 - **Unselected (Static mode)**: NetWorker runs all available nsrmmd processes.
In environments where unattended firewall ports must be restricted for security reasons, the storage node settings for mmnds for disabled

devices and Dynamic nsrmmds unselected (static mode) offer more control. These storage node settings cause all available nsrmmd firewall ports to be attended by running nsrmmd processes.

7. Select the **Configuration** tab.
8. In **Scanning**, set the attributes for SCSI library target devices on this storage node:
 - a. In **Device Sharing Mode**, select an option:
 - Server Default uses the NetWorker server setting for device sharing.
 - Maximal Sharing allows sharing of all devices.
 - No Sharing disables device sharing.
 - b. In **Search all LUNs**, select an option:
 - For NetWorker to detect all LUNs (logical unit numbers), select **Yes**. Detection can be time consuming.
 - For NetWorker to stop searching at the first available LUN, select **No**, the default setting.
 - c. In **Use persistent names**, choose whether NetWorker uses persistent device names specific to the storage host operating system when performing device discovery and autoconfiguration operations.
 - d. In **Skip SCSI targets** field:
 - If the storage node type is set to SCSI, list any SCSI targets to exclude from backup operations, one per line.
 - The format is bus.target.lun where the target and LUN fields are optional.
 - You can exclude a maximum of 63 targets.
9. For AFTD or DD Boost devices, configure the following settings in **Advanced Devices**:
 - In **Server network interface**, type the unique network interface hostname of the NetWorker server to be used by the storage nodes.
 - In **Clone storage nodes**, list by priority the hostnames of the storage nodes to be used for the save or “write source” side of clone operations originating from this storage node as the “read source.” The clone operation selects the first storage node in this list that has an enabled device and a functional nsrmmd process.
 - If the **Clone storage nodes** attribute does not contain a value, then the device operations use the value that is defined in the **Clone storage nodes** attribute for the Storage Node resource that was created for the NetWorker server.
 - If the **Clone storage nodes** attribute for the storage node resource is empty, then device operations use the values that are defined in **Storage nodes** attribute for the client resource that was created for the NetWorker server.

In backup-to-disk environments, it is possible for a single backup volume to be shared by multiple storage devices on different storage nodes. This can result in an ambiguous clone write source.
10. Click **OK**.

Modifying the timeout attribute for storage node operations

An attribute named **nsrmmd Control Timeout**, which is set during NetWorker server configuration, configures the amount of time a NetWorker server waits for a storage node request to be completed. If the timeout value is reached without completion of the request, the operation stops and an error message is logged. The default value assigned to **Nsrmmmd Control Timeout** is five minutes.

Procedure

1. In the server's **Administration interface**, click the **Configuration** button.
2. Select **View > Diagnostic Node**.
3. Right-click the NetWorker server in the left pane and select **Properties**.
4. Select the **Media** tab.
5. Modify the attributes as appropriate and click **OK**.

Configuring timeouts for storage node remote devices

Timeouts that determine how long to wait for mount requests on a storage node remote device before the save is redirected to another storage node are set in the **Properties** window of a device.

About this task

The Storage Node Devices area of the tab includes these attributes related to storage node timeouts:

- Save Mount Timeout
- Save Lockout

Save Mount Timeout and Save Lockout attributes change the timeout of a save mount request on a remote device.

If the mount request is not satisfied within the time frame specified by the Save Mount Timeout attribute, the storage node is locked out from receiving saved data for the time specified by the Save Lockout attribute.

The default value for Save Mount Timeout is 30 minutes. The default value for Save Lockout is zero, which means the device in the storage node continues to receive mount requests for the saved data.

 **Note:** The Save Mount Timeout applies only to the initial volume of a save request.

To modify the Save Mount Timeout and Save Lockout attributes, perform the following steps.

Procedure

1. In the server's **Administration interface**, click the **Devices** button.
2. Select **View > Diagnostic Node**.
3. Right-click the remote device and select **Properties**.
4. Select the **Advanced** tab.
5. Modify the attributes as appropriate and click **OK**.

Balancing the load on the storage node

The Save Session Distribution feature allows you to configure how NetWorker distributes save sessions between the storage nodes.

 **Note:** This feature is not available for clone and recover operations.

You can apply this feature to all NetWorker clients or to selected clients. This feature has two options:

- **Max sessions**—Distributes save sessions that are based on the setting in the **Max sessions** option in the storage node device resource. This is the default distribution method. The **Max sessions** option is more likely to concentrate the backup load on fewer storage nodes.
- **Target sessions**—Distributes save sessions that are based on the setting defined in the **Target sessions** option in each storage node device resource. The **Target sessions** option is more likely to spread the backup across multiple storage nodes.

When you select the **Max sessions** option, the NetWorker server distributes the save sessions for a client among eligible storage nodes as follows:

1. Identifies the available storage nodes in the NetWorker client's storage node affinity list.
2. Uses an available device on the first storage node in the list that is working below its Target sessions level.
3. When all devices on the first storage node are running at their target sessions level but some are running below their max sessions level, then NetWorker uses the least loaded device.
4. Continues until all available devices on all storage nodes in the client's storage node affinity list are in use.

When you select the **Target sessions** option, the NetWorker server distributes save sessions among eligible storage nodes as follows:

1. Identifies the available storage nodes in the storage node affinity list for the client.
2. Uses an available device on the first storage node in the list that is working below its Target sessions level.
3. When all devices on the first storage node are running at their target sessions levels, continue to the next storage node even if some devices are running below their max sessions level.
4. When all devices on all eligible storage nodes are running at their target sessions level, use the least loaded device that is running below its max session value.
5. Continues to send data to the least loaded device that is running below the max session value, until all devices on all available storage nodes are running at their max session levels.

Note the following performance considerations for storage node load balancing:

- Depending on the configuration of the backup environment, there is a potential to shorten the backup times by using the device **Target sessions** option rather than the device **Max sessions** option. However, using the device **Target sessions** option with the checkpoint restart feature can result in slower recovery times because a single save set is more likely to be spread across multiple storage nodes.
- It is recommended to use the default values for **Max sessions** as lowering these values can impact performance.

- Each NetWorker client has a storage node affinity list. The Save sessions distribution feature can only distribute a backup session for a client to multiple storage nodes when the client resource has two or more storage nodes in its storage node affinity list. The storage node affinity list is specified on the **Globals (2 of 2)** tab in the **NetWorker Client Properties** window.

Configuring the storage node affinity list for a client

Storage node affinity is a feature that determines which NetWorker servers and storage nodes receive the data from a client. Define the storage node affinity list in the Storage Nodes attribute of the Client resource.

About this task

For most Client resources, the default setting for the Storage Nodes attribute is *nsrserverhost*, which represents NetWorker server host. To configure the NetWorker server to direct the data for a client to a storage node device, modify the Storage Nodes attribute and specify the name of the storage node in the Storage Nodes attribute of the Client resource on a line above the default *nsrserverhost* entry.

If you create the Client resource for a storage node after you create the remote device on the storage node, the default setting of the Storage Nodes attribute is the storage node and the NetWorker server.

To modify the Storage Nodes attribute for a client, perform the following steps:

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In left navigation pane, expand **Clients**, right-click the appropriate client, and select **Properties**.
3. On the **Globals (2 of 2)** tab in the **Storage Nodes** attribute, specify the hostname of the storage node, and then click **OK**.

Results

The NetWorker software directs the client data to the first storage node in the affinity list with an enabled device, capable of receiving the data. The NetWorker software sends additional saves to the next storage node in the storage node affinity list that is based on criteria that are specified in [Balancing the load on the storage node](#) on page 102.

Specifying storage node load balancing

By default, NetWorker balances client backups across storage nodes that are based on the Max sessions attribute for each device on the storage node. If you choose to balance storage node loads by Max sessions, you can override this setting for selected clients.

About this task

Procedure

1. On the **Administration** window, click the **Server** button.
2. From the **View** menu, select **Diagnostic Mode**.
3. Right-click the NetWorker server in the left pane, and select **Properties**.
4. On the **General** tab, select a value from the **Save session distribution** list:
 - If you select **Target sessions**, then the NetWorker server balances the backups for all NetWorker clients across the storage nodes, based on device target session value. The NetWorker server ignores the value that is defined in Save session distribution attribute for each NetWorker client.

- If you select **Max sessions**, then you can still override this value for selected NetWorker client resources by setting the Save session distribution attribute in the client resource.

5. Click **OK**.

Overriding the save session distribution method for selected clients

If you selected Max sessions as the Save session distribution method for the NetWorker server, you can use the following procedure to override the setting for selected clients.

Procedure

1. On the **Administration** window, click the **Protection** button.
2. In the left navigation pane, expand **Clients**.
3. Right-click the appropriate client and select **Properties**.
4. On the **Globals (1 of 2)** tab, select **Target sessions** from the **Save session distribution** list.
5. Click **OK**.

Multiplexing

Multiplexing is the ability to write multiple data streams simultaneously to the same storage device. It is often more efficient for the NetWorker server to multiplex multiple save sets to the same device. There are also times when limiting the number of data streams to a particular device improves performance of the NetWorker environment.

Use the Target sessions, Max sessions, and Pool parallelism attributes to increase or limit the number of data streams that NetWorker writes to a device.

Target sessions

Use the Target sessions attribute on the Configuration tab of the Device resource to define the optimal number of backup sessions to assign to an active device.

Target sessions is not a hard limit; to set a hard limit for the number of sessions that a particular device can accept, use the Max sessions attribute.

The Target sessions attribute aids in load balancing devices by determining when the NetWorker software should write save streams to a device.

When a save session starts, the following actions occur:

- If a device is already receiving the number of backup sessions determined by the target sessions value, the NetWorker server uses the next underutilized device for the backups.
- If all available devices are receiving the number of backup sessions determined by their target sessions value, the NetWorker server overrides the set value and uses the device with the least activity for the next backup session.

Because it is often more efficient for the NetWorker server to multiplex multiple save sets to the same device, rather than write each save set to a separate device, the NetWorker server attempts to assign to each device a number of save sets, up to the value of target sessions, before assigning a save set to another device.

NOTICE When the NetWorker software assesses how many devices need to be involved in multiple save streams assignments with the same storage node, the device with the lowest target session value is used as a reference.

Max sessions

The Max sessions attribute on the Configuration tab of the Device resource defines the maximum number of save sessions for a device. The max sessions value is never

less than the target sessions value. It is recommended to use the default values for Max sessions as lowering these values can impact performance.

Bootstrap backup on a storage node

When the NetWorker server backup action performs a backup of the bootstrap save set, the data writes to a device that is local to the NetWorker server. You cannot back up bootstrap data to a remote device, but you can clone or stage the bootstrap to a remote device. When you recover a bootstrap save set, you must recover the data from a local device.

Staging bootstrap backups

You can direct bootstrap backups to a disk device such as an AFTD or FTD device.

However, if you stage a bootstrap backup to a volume on another device, NetWorker reports the staging operation as complete although the “recover space” operation has not started, and the bootstrap remains on the original device. Therefore, if the staged bootstrap is accidentally deleted, you can recover the bootstrap from the original disk. The *NetWorker Server Disaster Recovery and Availability Best Practices Guide* describes how to recover a bootstrap from the original disk.

Also, if the bootstrap data is not staged from the original disk, the data on the original disk is subject to the same retention policies as any other save set backup and is, therefore, deleted after the retention policy has expired.

Troubleshooting storage node affinity issues

If a backup fails because of a problem related to the storage node affinity, a message similar to the following might appear:

```
no matching devices; check storage nodes, devices or pools
```

Possible causes for this error message include:

- No enabled devices are on the storage nodes.
- The devices do not have volumes that match the pool required by the backup request.
- All devices are set to read-only or are disabled.

For example, if the client has only one storage node in its Storage Node list, and all devices on that storage node are disabled, fix the problem and then restart the backup.

Complete one of the following actions to fix the problem:

- Enable devices on one of the storage nodes in the storage node list for the client.
- Correct the pool restrictions for the devices in the storage node list.
- Configure an additional storage node that has enabled devices that meet the pool restrictions.
- Set one of the devices to read/write.

Storage Node Options

Storage node options secures the environment by restricting unattended firewall ports.

If Dynamic nsrmmnds is selected, then the NetWorker adds devices only on demand. Dynamic mode can improve bandwidth and performance, but the firewall ports may be

left unattended by running processes. If Dynamic nsrmmnds is set to No, then the NetWorker starts all available nsrmmnd processes. Static mode along with the “mmds for disabled devices” option provides greater control on the nsrmmnd firewall ports. They ports are attended by running processes.

MMDS for Disabled Devices	Dynamic nsrmmnds	Functionality
No	No	This is a static configuration. The number of ports that are to be open is equal to the number of disk devices on the storage node, that is, the “max nsrmmnd count”.
Yes	No	This is a “static” configuration. The number of ports that are to be open is equal to the number of disk devices on the storage node, that is, the number of nsrmmnds. In this configuration, the ports are not utilized efficiently. A few of them remain unused. When the disabled devices are re-enabled, then all the open ports might be used.
Yes/No	Yes	This is a “dynamic” configuration. This configuration ensures that the nsrmmnds are started on demand by the server.

 **Note:** You must ensure that the sum of storage node devices nsrmmnds count should not exceed 655.

Configuring a dedicated storage node

All devices created on storage nodes, except the devices for the NetWorker server include the Dedicated Storage Node attribute. A dedicated storage node can only back up data that originates from the storage node host. When you configure a storage node as a dedicated storage node, you require a Dedicated Storage Node license.

About this task

After you create a storage node, perform the following steps to configure the storage node as dedicated.

Procedure

1. On the **Administration** window, click **Devices**.
2. In the left navigation pane, expand **Storage Nodes**, right-click the storage node, and then select **Properties**.
3. On the **Configuration** tab, in the **Dedicated Storage Node** option, select **Yes**.
4. Click **OK**.

Troubleshooting storage nodes

This section provides troubleshooting information about storage nodes.

Storage node affinity errors

A storage node affinity problem may exist when a backup fails with an error message similar to the following:

```
No matching devices; check storage nodes, devices or pools
```

This error message can appear for the following reasons:

- All the devices in the storage node are disabled.
- Each device in the storage node contains a volume that does not match the pool that the backup request requires.
- All the devices in the storage node are set to read-only.

To resolve this error:

- Enable devices on one of the storage nodes.
- Correct the pool restrictions for the devices that are listed in the **Storage Nodes** attribute of the Pool resource.
- Add another storage node that has enabled devices and meets the pool restrictions to the Storage Nodes attribute of the Pool resource.
- Write-enable one of the devices.
- Adjust the **Save Mount Timeout** and **Save Lockout** attributes for in the Device resource for the storage node.

Storage node timeout errors

If the `nsrd` process starts on the NetWorker server and detects that a setting for the `NSR_MMDCONTROL` variable exists, a message similar to the following appears:

```
NSR_MMDCONTROL env variable is being ignored
use nsrmmd control timeout attribute instead
```

If you receive this message, perform the following steps.

1. Shut down the NetWorker services.
2. Remove the environment variable setting for `NSR_MMDCONTROL`.
3. Restart the NetWorker services.
4. Use NMC to connect to the NetWorker server.
5. Adjust the value of the **nsrmmd Control Timeout** attribute in the Storage Node resource to the value that was assigned to the `NSR_MMDCONTROL` variable, or to a value that best meets the current requirements. [Modifying the timeout attribute for storage node operations](#) on page 101 provides more information.

Disk storage devices

NetWorker software supports a variety of different backup to disk (B2D) methods. These methods all use disk files that NetWorker creates and manages as storage devices. These devices can reside on a computer's local disk or a network-attached disk. NetWorker supports FTD, AFTD, and DD Boost device types. This section does not cover disk-based devices that emulate other device types, such as virtual tape libraries (VTLs).

FTD

A file type device (FTD) is a basic disk device type that has been available for many years. FTDs have limited use and support and this chapter describes them for legacy purposes only.

AFTD

Advanced file type devices (AFTDs) support concurrent backup and restore operations and require the NetWorker DiskBackup Option (DBO) license. AFTDs are supported for the following configurations:

- A local disk on a NetWorker storage node.
- A network-attached disk device that is NFS-mountable to a NetWorker storage node running a Linux or UNIX operating system.
- A network-attached disk device that is CIFS-mountable to a NetWorker storage node running on Windows.

The Client Direct feature enables NetWorker clients to back up directly to AFTDs over a CIFS or NFS network, bypassing the storage node. For Client Direct backups, the storage node manages the devices but does not handle the backup data unless the Client Direct workflow is not available.

DD Boost devices

DD Boost devices reside on Data Domain storage systems that have the DD Boost features enabled. These devices are similar to AFTDs except they store backup data in a highly compressed and deduplicated format. The DD Boost API accesses the DD Boost devices over a network. NetWorker can perform DD Boost backups through either the NetWorker storage node workflow or the Client Direct file access workflow.

The Client Direct workflow enables NetWorker clients with distributed segment processing (DSP) and network access to deduplicate their own backup data and send the data directly to the DD Boost devices. This method bypasses the storage node and frees up network bandwidth. The storage node manages the devices but does not handle the backup data workflow if the Client Direct workflow is available.

If Client Direct backup is not available, NetWorker automatically routes the backup through the storage node where it is deduplicated and sent to the DD Boost devices for storage. Restore operations work similarly. If Client Direct is not available for a restore, then NetWorker performs a traditional storage node recovery.


This guide does not cover DD Boost operations. The *NetWorker Data Domain Boost Integration Guide* provides details on DD Boost devices

Example environment

The following figure shows various backup-to-disk options deployed in a mixed operating system environment.

- Linux/UNIX Storage Node A writes its backups to either of the following:
 - The AFTD through an NFS connection to Disk Device 1.
 - The AFTD on Local Disk 1.
- Windows Storage Node B uses a CIFS connection to back up to the NAS AFTD on Disk Device 2.
- Data Domain system C writes its backups to a DD Boost device on Local Disk 2.

Figure 12 Example NetWorker disk backup configuration in a mixed backup environment.



Considerations for Client Direct clients

Client Direct backups enable clients to bypass the storage node and back up directly to storage devices. The storage node manages the devices but does not handle the backup data. Device configuration for Client Direct clients depends on the type of storage device and how it is connected to the storage nodes.

A Client Direct backup reduces bandwidth usage and bottlenecks at the storage node, and provides highly efficient backup data transmission.

If a Client Direct backup is not available, a traditional storage node backup occurs instead.

Requirements for Client Direct backups

Ensure that the environment meets the following requirements to perform Client Direct backups:

- NetWorker clients on UNIX/Linux or Microsoft Windows can perform non-root and cross-platform Client Direct backups to AFTDs. The AFTD can be managed by either a UNIX/Linux or a Windows storage node, and can be either local or mountable on the storage node.

To perform non-root and cross-platform Client Direct backups to AFTDs, the NetWorker server and the storage node software must be version 8.1 or later.

- If an NFS server provides the AFTD storage for Client Direct backups, then the NFS server must permit access by using the NFSv3 protocol with AUTH_SYS (AUTH_UNIX) authentication. The NFS server also must not restrict access to clients by using only privileged ports.
- If you enable checkpoint restart for a client, then Client Direct backups are supported only to AFTDs, and not to DD Boost devices. If a client is enabled for checkpoint restart and a Client Direct backup is tried to a DD Boost device, then the backup reverts to a traditional storage node backup instead.

For Client Direct backups to AFTDs, checkpoint restart points are made at least 15 seconds apart. Checkpoints are always made after larger files that require more than 15 seconds to back up.

- Archive operations are not currently supported for Client Direct backups.

Configuring Client Direct backups

Procedure

1. Ensure that the clients that perform Client Direct backups have a network connection and a remote network protocol to reach the storage device.
Windows clients can use a CIFS or NFS path, although a CIFS path generally yields better performance. UNIX clients must use an NFS path.
2. Specify the complete path for the destination device in the **Device access information** attribute on the **General** tab of the **Device Properties** dialog box for the destination device.

Keep in mind the following points when you specify the path:

- If the storage device is directly attached to a Windows storage node, then the storage node uses a different path than the Client Direct clients. If the storage device is not directly attached to any storage node, then the path is the same for all storage nodes and Client Direct clients.
- The device access information path should include multiple access paths to cover local and remote use cases.
- To specify an NFS path, use the `NFS_host:/path` format regardless of whether the AFTD is local to the storage node or mountable on the storage node. Non-root UNIX/Linux NetWorker clients require this NFS format for Client Direct access.
- For Windows Client Direct backups, specify a CIFS path instead of an NFS path. A CIFS path generally yields better performance.
- If you are setting up an AFTD on a Windows storage node, specify the CIFS path first. For example:


```
\\fileserver\aftd1
fileserver:/aftd1
```

- If you are setting up a UNIX/Linux storage node, specify the NFS path first. For example:

```
fileserver:/aftd1
\\fileserver\aftd1
```

The following figure shows an example set of paths for a CIFS AFTD.

Figure 13 Paths for CIFS AFTD



3. If an NFS server provides the AFTD storage for Client Direct backups, then specify the username and password that is required to access the NFS server

for the AFTD in the **Remote user** and **Password** attributes on the **Configuration** tab of the **Device Properties** dialog box for the device.

4. Ensure that the **Client direct** attribute is enabled on the **General** tab of the **Client Properties** dialog box for each Client Direct client.

Client Direct backups are enabled by default.

Select **View > Diagnostic Mode** in the Administration interface to access the **Client direct** attribute in the **Client Properties** dialog box.

Differences between FTDs, AFTDs, and DD Boost devices

The following table lists the functional differences between traditional file type devices (FTDs), AFTDs, and DD Boost devices.

The *NetWorker Data Domain Boost Integration Guide* provides details on DD Boost devices.

Table 26 Differences between disk devices

Function or operation	File type device (FTD)	Advanced file type device (AFTD)	DD Boost device
Create a device	Device Property Window Select the media type: file. UNIX/Linux storage node: local or NFS only. Windows storage node: local path only. CIFS is not supported for FTDs.	- Device Configuration Wizard - Device Property Window Select media type: adv_file. UNIX/Linux storage node: local or NFS only. Windows storage node: local or CIFS using UNC path or using NFS: Remote user, Password.	- Device Configuration Wizard - Device Property Window Select media type: Data Domain
Storage location	Specified in the Name attribute.	Specified in the Device Access Information attribute.	Specified in the Device Access Information attribute.
Concurrent save set operations	No.	Yes.	Yes.
Concurrent AFTD recovery operation limitations on page 127 provides more information about performing concurrent recovery operations from an AFTD.			

Table 26 Differences between disk devices (continued)

Function or operation	File type device (FTD)	Advanced file type device (AFTD)	DD Boost device
Reclaiming or recovering space	The <code>nsrim</code> program removes both aborted and expired save sets, once every 24 hours, by the Expiration action, at the time defined in the Server backup workflow (if you have set volume recycle to Auto).	- Aborted save sets immediately removed. - The <code>nsrim</code> program removes expired save sets from the media database once every 24 hours, by the Expiration action, at the time defined in the Server backup workflow (if you have set volume recycle to Auto). NetWorker removes space on the AFTD as specified in the Reclaim Space Interval of the staging policy.	- Reclaims only data that is unique, not required by other existing backups. - NetWorker does not immediately remove aborted save sets, but marks them recyclable. A restarted save can be deduplicated. Otherwise, NetWorker removes the aborted save set during the next Recover Space operation.
Volume default capacity for devices	If the file type device was used before setting the Volume Default Capacity attribute, the data for that file type device must be staged or cloned to another device.	Does not apply.	Does not apply.
AFTD Percentage Capacity	Does not apply.	A setting determines the capacity that NetWorker software should stop writing to an AFTD: spans from 1% to 100%.	Does not apply.
When file system or volume is full	- Waiting message is displayed if no writable volume available or until volume becomes available. - Volume marked full and is no longer available for backups until the volume becomes appendable.	- Message is displayed stating file system requires more space. - The <code>nsrim</code> program invoked to reclaim space for expired save set on AFTD. - Notification is sent by email stating device is full. - Device waits until space become available. The volume is never marked as full.	Backup to a DD Boost device fails and stops when full.
Save set continuation	Yes.	No. Save sets that start on an AFTD must be	No. Save sets that start on a DD Boost device

Table 26 Differences between disk devices (continued)

Function or operation	File type device (FTD)	Advanced file type device (AFTD)	DD Boost device
		completed on the same device.	must be completed on the same device.
Data format in device	Open Tape Format (OTF).	Save stream (uasm) format (uses less space).	Deduplicated
Client Direct backup: the storage node manages the devices for the NetWorker clients, but the clients send their backup data directly to the devices via network access, bypassing the storage node.	No.	Yes. Clients send their own backup data directly to the storage devices. If Client Direct backup is not available, a traditional storage node backup is performed. NetWorker archive operations are not supported for Client Direct backup.	Yes. Clients use DD Boost DSP functionality to deduplicate their own backup data before sending it directly to the storage devices. If Client Direct backup is not available, a traditional storage node backup is performed. NetWorker archive operations are not supported for Client Direct backup.

Device target and max sessions default values and ranges

There are default values and ranges for device target and max sessions in the NMC **NetWorker Administration** window.

The following table lists the default values for target and max sessions values.

Table 27 Default values and ranges for target and max sessions attributes

Device type	Default target sessions	Default max sessions	Recommended sessions*	Range
AFTD (traditional storage)	4	32	1 - 32	1 - 1024
AFTD (including Data Domain CIFS/NFS)	4	32	1 - 10	1 - 1024
CloudBoost	10	80	1 - 10	1 - 200
Data Domain (DD Boost)	20	60	1 - 10	1 - 60
DD Cloud Tier	20	60	1 - 10	1 - 60

Table 27 Default values and ranges for target and max sessions attributes (continued)

Device type	Default target sessions	Default max sessions	Recommended sessions*	Range
NDMP	4	512	1 - 32	1 - 1024
FTD (traditional)	4	32	1 - 16	1 - 1024
ProtectPoint	20	120	1 - 10	1 - 1024
VTL/Tape (traditional)	4	32	1 - 16	1 - 512
VTL/Tape (Data Domain / Deduplicated)	4	32	1 - 1	1 - 512
* The recommended session values are guidelines only and are subject to bandwidth, data type, and device capabilities.				

Advanced file type devices

Advanced file type devices (AFTDs) overcome the main restrictions of traditional file type device (FTD) storage. AFTD storage is designed for large disk storage systems that use a volume manager to dynamically extend available disk space if the disk runs out of space during backup.

The *NetWorker E-LAB Navigator* provides a list of supported volume managers.

Memory requirements for AFTD backups

The physical memory requirements for a NetWorker storage node and Client Direct client depends on the peak AFTD usage.

The following is the list of physical memory requirements for AFTD:

- Allowing for other types of devices and services on a typical storage node, a storage node should have a minimum of 8 GB of RAM to host AFTDs.
- AFTD clients require a minimum of 4 GB of RAM at the time of backup to ensure optimum performance for Client Direct backups. Client Direct backups require client access to the AFTDs on either a CIFS or NFS network.
- Each AFTD requires an initial 24 MB of RAM on the storage node and Client Direct client. Each AFTD save session requires an additional 24 MB. To run 10 sessions requires 24 + 240 MB. The default max sessions of 60 sessions per AFTD requires 24 + 1440 MB.

Required AFTD DFA device settings for Hyper-V environments

For Hyper-V environments, when creating a NetWorker AFTD DFA device on an NTFS or ReFS volume, Microsoft requires certain settings.

If the NetWorker AFTD DFA device is created on an NTFS volume, virtual hard disk (VHD/VHDx) files must be uncompressed and unencrypted. If the NetWorker AFTD DFA device is created on an ReFS volume, virtual hard disk (VHD/VHDx) files must not have the integrity bit set.

Create and configure an AFTD

You can create an AFTD by using either the Device Wizard or the device properties window.

Creating an AFTD by using the Device Wizard

About this task

If you are creating an AFTD to use the client direct feature, see [Considerations for Client Direct clients](#) on page 109 for information about specifying network path information when creating the AFTD.

Procedure

1. In the NMC **Enterprise** view, double-click the NetWorker managed application to launch its window.
2. In the **NetWorker Administration** window, select the **Devices** view.
3. Verify that the path to the storage directory that will contain the AFTDs is allowed.
 - a. In the navigation tree, select **Storage Nodes**.
 - b. Right-click the storage node that you will use, and select **Properties**.
 - c. In the **AFTD allowed directories** list, verify or type the path of the storage directory that will contain the AFTDs.

AFTDs can be created and accessed only by these listed paths. If this list is left empty, there are few restrictions as to where a device path can be created.
 - d. Click **OK**.
4. In the navigation tree, right-click **Devices**, and select **New Device Wizard**.
5. In the **Select the Device Type** window, select **AFTD** and click **Next**.
6. In the **Select Storage Node** window, specify the path to the storage directory that will contain the AFTDs.
 - a. In the **Storage Node** list, select the storage node that you will use.
 - b. If the directory for the intended AFTDs is on a different storage node or a remote storage system, select **Device storage is remote from this Storage Node** and type the **Network Path** of the remote host directory that will contain the devices.

For example, if the storage node is a Microsoft Windows system and you use a CIFS AFTD on a remote storage system host, this path could be something like the following:

```
\\dzone1_storhost2.lss.corp.com\share-1
```

This storage path is not a device. It is the directory location in which the shared devices are to be created.

7. In **Browse or Manual**, select which option you will use to specify the pathnames of the devices:
 - **Browse Storage Node or network path.** The next wizard step will prompt you to browse and add the devices.
 - **Manually enter local or remote device paths.** Select this to skip the browse step and manually type unique names for the devices you want to add:
 - For remote devices, type the device paths relative to the Network Path that you specified for the storage directory. For example:
cifsaftd-1
cifsaftd-2
 - For local devices, type the absolute paths to these devices. For example:
C:\cifsaftd-1
C:\cifsaftd-2

[Multiple devices for a single volume configuration](#) on page 122 provides details for shared volumes.
8. If the storage host is remote from the storage node, in the **Authentication** area, type the appropriate **Username** and **Password** to access the storage directory.
9. Click **Next**.
10. If you selected the **Browse** option in the previous window:
 - a. In the **Select the Device Path** window, verify that the storage node shows the path of a storage directory.
 - b. Add devices to the storage directory by clicking **New Folder** and typing unique device names. For example:
cifsaftd-1
cifsaftd-2
 - c. Select the new devices to add and click **Next**.
11. In the **Configure Device Attributes** window, specify the attributes. If you added multiple devices in the previous window, select each device individually and specify its attributes:
 - a. In **NetWorker Device Name**, type a unique name for the AFTD device.
For example, for a device on the NetWorker server host storage node:
aftd-1

If you configure the device on a storage node host that is not the NetWorker server host, it is a “remote device” and this attribute must be specified with **rd=** and a colon (:) in the following format (for Microsoft Windows):
rd=remote_storagenode_hostname:device_name
For example:
rd=dzone1_storhost2:aftd-1
 - b. (Optional) Add a comment in the **Comment** field.

c. If Client Direct backup will be used, follow the details in [Considerations for Client Direct clients](#) on page 109.

d. In **Target Sessions** specify the number of sessions that a nsrmmd data mover process on the device will handle before another device on the host will take the additional sessions. Use this setting to balance the sessions among nsrmmd processes.

If another device is not available, then another nsrmmd process on the same device will take the additional sessions.

Typically, set this attribute to a low value. The default value is 4 for AFTDs. It may not be set to a value greater than 60.

e. In **Max Sessions** specify the maximum number sessions the device may handle. If no additional devices are available on the host, then another available storage host takes the additional sessions, or retries are tried until sessions become available.

The default value is 32 for AFTDs, which typically provides best performance. It cannot be set to a value greater than 60.

 **Note:** The **Max Sessions** setting does not apply to concurrent recover sessions.

f. Click **Next**.

12. In the **Label and Mount device** window, if you select the **Label and Mount** option, specify the attributes for:

- **Pool Type**.
- **Pool** to use.

13. On the **Review the Device Configuration** page:

- a. Review the settings.
- b. Click **Configure**.

14. On the **Check results** page:

- a. Review whether the devices were successfully configured or if any messages appeared.
- b. Click **Finish**.
- c. To change any of the settings, click **Back** to the correct wizard page.

Creating an AFTD by using the Properties window (Linux and UNIX)

Procedure

1. Create one directory for each disk (or partition) to be used for an AFTD.

AFTDs require a directory (folder) to be created in the disk file system that the NetWorker server or storage node recognizes as the device name (and the destination for the data).

 **NOTICE** Do not use a temporary directory for AFTDs. The data could be overwritten.

2. In the **NetWorker Administration** window, click the **Devices** view.
3. Verify that the path to the storage directory that will contain the AFTDs is allowed.

- a. In the navigation tree, select **Storage Nodes**.
 - b. Right-click the storage node that you will use, and select **Properties**.
 - c. In the **AFTD allowed directories** list, verify or type the path of the storage directory that will contain the AFTDs.

AFTDs can be created and accessed only by these listed paths. If this list is left empty, there are few restrictions as to where a device path can be created.
 - d. Click **OK**.
4. In the navigation tree, right-click **Devices** and select **New**.
- The **Create Device** window opens, with the **General** tab selected. The **Identity** area might show a default device name in the **Name** field.
5. In the **Identity** area, set the following attributes:

- a. In the **Name** attribute, type the name of the directory that you created for the AFTD.

For example:

`aftd-1`

If you configure the device on a separate storage node host that is not the NetWorker server host, it is a remote device and this Name attribute must be specified with `rd=` in the following format:

`rd=remote_snode_hostname:device_name`

For example:

`rd=snode-1:aftd-1`

- b. (Optional) Add a comment in the **Comment** field.
- c. In the **Device Access Information** attribute, provide complete paths to the device directory. You can provide alternate paths for the storage node and for Client Direct clients, for example:

For non-root or cross-platform Client Direct access:

For non-root or cross-platform Client Direct access to an AFTD, do not specify an automounter path or a mounted path. Instead, specify the path in the `host:/path` format, even if the AFTD is local to the storage node.

For example:

`NFS_host:/path`

where:

- `NFS_host` is the hostname of the NFS file server
 - `path` is the NFS-mountable path that is exported by the file server
- This format is required to allow Client Direct access for Windows or non-root UNIX clients.

 **Note:** Non-root Client Direct access to an NFS AFTD is supported only with the NFSv3 protocol and AUTH_SYS authentication on the NFS host. For Client Direct access to an AFTD when the backup client is able to run as root on the AFTD host, provide a mount point or automounter path.

 **Note:** For example, for an NFS-mounted device:

```
/mnt/aftd-1
/net/storho-1/snode-1/aftd-1
```

where:

- **aftd-1** is the storage device directory name
 - **storho-1** is the storage system hostname
 - **snode-1** is the storage node hostname
- The first path enables the storage node to access the device via its defined mount point. The second path enables Client Direct clients to use the automounter path to directly access the device, bypassing the storage node.

d. In the **Media Type** field, select **adv_file**, for the AFTD.

[Considerations for Client Direct clients](#) on page 109 provides additional details for Client Direct configurations.

[Multiple devices for a single volume configuration](#) on page 122 provides additional details for shared volumes.

6. In the **Status** area, ensure that the **Auto Media Management** tape feature is not enabled.
7. In the **Cleaning** area, leave the options for cleaning at their default (disabled) settings, so that automatic cleaning is not invoked.
8. Select the **Configuration** tab.
9. In the **Save Sessions** area, set the number of concurrent save sessions (streams) and the number of **nsrmmd** (data mover) processes the device may handle:
 - **Target Sessions** is the number of sessions that a **nsrmmd** process on the device will handle before another device on the host will take the additional sessions. Use this setting to balance the sessions among **nsrmmd** processes.

If another device is not available, then another **nsrmmd** process on the same device will take the additional sessions.

Typically, set this attribute to a low value. The default values are 4 for AFTDs and 6 for DD Boost devices. It may not be set to a value greater than 60.

[Multiple devices for a single volume configuration](#) on page 122 provides details on volume sharing.
 - **Max Sessions** is the maximum number sessions the device may handle. If no additional devices are available on the host, then another available storage host takes the additional sessions, or retries are attempted until sessions become available.

The default values are 32 for AFTDs and 60 for DD Boost devices, which typically provides best performance. It cannot be set to a value greater than 60.

The **Max Sessions** setting does not apply to concurrent recover sessions.
 - **Max nsrmmd count** limits the number of **nsrmmd** processes that can run on the device. Use this setting to balance the **nsrmmd** load among devices. The default value is 4.

To modify this value, first adjust the sessions attributes, apply, and monitor the effects, then update max nsrmmd count.

At least one nsrmmd process is reserved for restore or clone operations.

10. In the **Local Backup** area, leave **Dedicated Storage Node** at **No** (the default).
11. In the **Remote Host** area, if an NFS path is specified in the **Device Access Information**, then type a **Remote User** name and **Password**.

The remote username is the name of the user on the NFS server. It is recommended that you also specify the numeric user id (UID) of that user. Do this by appending a colon (:) and the UID after the username, for example, `user_name:4242`.

Note: If the device username is changed after labeling, manual action may be required to change the owner of all files and directories in the AFTD. NetWorker will try to perform this automatically during the next operation, however the ability to do so depends on the security configuration of the file server where the AFTD storage resides.

12. Click **OK** when the configuration is complete.
13. If a new password for an AFTD is provided, unmount and re-mount the device to ensure that the change takes effect.

Creating an AFTD by using the Properties window (Windows)

You can configure an AFTD on a storage node running Microsoft Windows.

About this task

Procedure

1. Create one directory for each disk (or partition) to be used for an AFTD.
AFTDs require a directory (folder) to be created in the disk file system that the NetWorker server or storage node recognizes as the device name (and the destination for the data).

NOTICE Do not use a temporary directory for AFTDs. The data could be overwritten.

2. In the **NetWorker Administration** window, click the **Devices** view.
3. Verify that the path to the storage directory that will contain the AFTDs is allowed.
 - a. In the navigation tree, select **Storage Nodes**.
 - b. Right-click the storage node that you will use, and select **Properties**.
 - c. In the **AFTD allowed directories** list, verify or type the path of the storage directory that will contain the AFTDs.
AFTDs can be created and accessed only by these listed paths. If this list is left empty, there are few restrictions as to where a device path can be created.
 - d. Click **OK**.

4. In the navigation tree, right-click **Devices** and select **New**.

The **Create Device** window opens, with the **General** tab selected. The **Identity** area might show a default device name in the **Name** field.

5. In the **Identity** area, set the following attributes:

- a. In the **Name** attribute, type the name of the directory that you created for the AFTD.

For example:

`aftd-1`

If you configure the device on a separate storage node host that is not the NetWorker server host, it is a remote device and this Name attribute must be specified with `rd=` in the following format:

`rd=remote_snode_hostname:device_name`

For example:

`rd=snode-1:aftd-1`

- b. (Optional) Add a comment in the **Comment** field.

- c. In the **Device Access Information** attribute, provide complete paths to the device directory. You can provide alternate paths for the storage node and for Client Direct clients, for example:

- For an AFTD on the storage node's local disk, which it shares via CIFS:

`E:\aftd-1`

`\\snode-1\aftd-1`

The first path enables the storage node to access the device via its local drive. The second path enables Client Direct clients to access the device directly, bypassing the storage node.

- For a CIFS-mounted AFTD, specify the complete paths of the directory that is created by using the Universal Naming Convention (UNC), for example:

`\\CIFS_host\share-point-name\path`

- d. In the **Media Type** field, select **adv_file**, for the AFTD.

[Considerations for Client Direct clients](#) on page 109 provides additional details for Client Direct configurations.

[Multiple devices for a single volume configuration](#) on page 122 provides additional details for shared volumes.

6. In the **Status** area, ensure that the **Auto Media Management** tape feature is not enabled.
7. In the **Cleaning** area, leave the options for cleaning at their default (disabled) settings, so that automatic cleaning is not invoked.
8. Select the **Configuration** tab.
9. In the **Save Sessions** area, set the number of concurrent save sessions (streams) and the number of `nsrmmd` (data mover) processes the device may handle:
- Target Sessions** is the number of sessions that a `nsrmmd` process on the device will handle before another device on the host will take the additional sessions. Use this setting to balance the sessions among `nsrmmd` processes.
- If another device is not available, then another `nsrmmd` process on the same device will take the additional sessions.

Typically, set this attribute to a low value. The default values are 4 for AFTDs and 6 for DD Boost devices. It may not be set to a value greater than 60.

[Multiple devices for a single volume configuration](#) on page 122 provides details on volume sharing.

- **Max Sessions** is the maximum number sessions the device may handle. If no additional devices are available on the host, then another available storage host takes the additional sessions, or retries are attempted until sessions become available.

The default values are 32 for AFTDs and 60 for DD Boost devices, which typically provides best performance. It cannot be set to a value greater than 60.

The **Max Sessions** setting does not apply to concurrent recover sessions.

- **Max nsrmmd count** limits the number of `nsrmmd` processes that can run on the device. Use this setting to balance the nsrmmd load among devices. The default value is 4.

To modify this value, first adjust the sessions attributes, apply, and monitor the effects, then update max nsrmmd count.

At least one `nsrmmd` process is reserved for restore or clone operations.

10. In the **Local Backup** area, leave **Dedicated Storage Node** at **No** (the default).
11. In the **Remote Host** area, if a network path is specified in the **Device Access Information**, then type a **Remote User** name and **Password**.
12. Click **OK** when the configuration is complete.
13. If a new password for an AFTD is provided, unmount and re-mount the device to ensure that the change takes effect.

AFTD device target and max sessions

The default settings for AFTD target sessions and max device sessions typically provide optimal values for AFTD performance:

- Device target sessions is 1
- Device max sessions is 32 to avoid disk thrashing

If required, both device target, and max session attributes can be modified to reflect values appropriate for the environment.

 **Note:** The Max Sessions setting does not apply to concurrent recover sessions.

Multiple devices for a single volume configuration

In some environments, a configuration of multiple devices that share a single NetWorker storage volume can result in performance gains. For example, a read or write request can be sent to the storage node that is closest to the requestor. However, for some use cases and environments concurrent read/write operations to a single volume from many storage nodes could result in disk thrashing that impacts performance.

Multiple devices can be created on separate storage nodes or on the same storage node. Each device must be created separately, have a different name, and must correctly specify the path to the storage volume location.

For example, if you create three devices, one on the NetWorker server host named “dzone1” (that uses the server’s local storage node) and two remote devices (rd) on

remote storage nodes, the Name attributes for the three devices, each created separately, might be specified by different aliases as follows:

aftd-1a

rd=dzone1-sn2:aftd-1b

rd=dzone1-sn3:aftd-1c

The Device Access Information for each of these aliases would specify a single directory that must be specified as a valid complete path. For example, if a directory is named “aftd-1” on the storage host named “storho1,” the path might be specified as follows:

- If the storage node uses an automounter:
/net/storho1/dzone1/aftd-1
- If the storage node uses an explicit system mountpoint, you might specify one of the following paths:
 - /mnt/storho1/dzone1/aftd-1
 - /mnt/dzone1/aftd-1
 - storho1:/dzone/aftd-1

AFTD concurrent operations and device formats

The following operations can be performed concurrently on a single storage node with an AFTD:

- Multiple backups and multiple recover operations
- Multiple backups and one manual clone operation
- Multiple backups and one automatic or manual staging operation

It might be required to increase the server parallelism value to complete the concurrent operations with an AFTD device when the number of simultaneous save sessions reaches the maximum value for server parallelism.

For example, if server parallelism is set to 4, and there are 4 simultaneous saves going to an AFTD, set the server parallelism to 5 to complete a concurrent clone/stage operation from this AFTD while the four saves are in progress.

 **Note:** Starting with NetWorker 8.0, multiple clone sessions can be run from a single AFTD or DD Boost device if each clone is written to a dedicated tape device. However, the number of clone sessions that can be run is limited by the value in the device’s max nsrmmd count attribute. [Create and configure an AFTD](#) on page 115 provides more information.

Labeling and mounting an AFTD

If there are multiple volumes in the pool, you can select an available volume to associate with the device.

Procedure

1. Right-click the AFTD storage device and then select **Label**.

The **Label** dialog box appears.

2. In the **Pools** field, select the media pool to be used for the device.

A label for the storage device is generated and displays in the **Volume Label** field. The label name is based on the label template for the selected pool.

It is recommended to use a pool that is dedicated to AFTD backup devices only.

NOTICE If an existing volume is re-labeled, a warning is issued. The data that is previously stored on the volume is lost and this action cannot be undone. Mounting the volume without labeling provides access to previous data.

3. Select **Mount after labeling** and then click **OK**.

Insufficient AFTD disk space

When an AFTD runs out of disk space, the current backup is interrupted and the following message displays:

```
Waiting for more available space on filesystem device-name
```

Immediately following the message, the action that is associated with the "Filesystem Full — Recover adv_file Space" notification occurs. By default, the action for this notification uses the `nsrim` command to delete expired save sets. If enough space is cleared, the backup continues. If the recycle setting for the volume is manual, then the expired save sets are not removed from the volume.

The AFTD deletes expired save sets depending on the retention policy and the recycle setting. If sufficient storage space is not available after 10 minutes from when the expired savesets begin deletion, the associated "Filesystem Full—Waiting for adv_file Space" notification action occurs. By default, an email notification is sent to the root user on the NetWorker server on UNIX and Linux, and a message is logged in the media log file in `NetWorker_install_path\logs` on Windows.

When the notification is sent, and the message is logged in the media log file, the backup stops until space is available for the backup to continue. You can create customized notifications to change and expand how the NetWorker software behaves when an "AFTD Filesystem Full" notification occurs. Custom notifications can also run custom scripts and other programs to expand the capacity of existing AFTDs.

The chapter "Reporting NetWorker Datazone Activities" provides more information about how to configure notifications.

Creating a custom notification to extend disk space

While the NetWorker default *Filesystem Full — Recover adv_file Space* notification works by removing its expired save sets, a custom notification could be configured to expand disk or file system space in other ways.

Procedure

1. In the server's **Administration** interface, click **Server**.
2. Right-click **Notifications** and select **New**.
3. In the **Name** field, type a unique name for this custom notification.
For example: **First adv_full** notice.
4. In the **Event** field, clear all choices except **adv_file**.
5. In the **Priority** field, clear all choices except **Waiting**.
6. In the **Action** field, type the full path of the custom script that is configured to expand disk space.
For example: `/mybin/my_first_custom_script`.
7. Click **OK**.

Creating a custom notification for insufficient disk space

The NetWorker default *Filesystem Full — Waiting for adv_file Space* notification works by sending an email notification. A custom notification could be configured to do whatever the user indicates. The wait time after the default notification is approximately 10 minutes.

Procedure

1. In the server's **Administration** interface, click **Server**.
2. Right-click **Notifications** and select **New**.
3. In the **Name** field, type a unique name for this second custom notification.
For example: Second adv_full Notice.
4. In the **Event** field, clear all choices except **adv_file**.
5. In the **Priority** field, clear all choices except **Critical**, **Emergency**, and **Alert**.
6. In the **Action** field, type the full path of the custom script to be run.
For example: `/mybin/my_second_custom_script`.
7. Click **OK**.

AFTD load balancing

You can adjust the target and max sessions attributes per device to balance the data load for simultaneous sessions more evenly across available devices. These parameters specify the maximum number of save sessions to be established before the NetWorker server attempts to assign save sessions to another device.

For AFTDs, all volumes, depending on the selection criteria (pool settings), choose the AFTD with the least amount of data written to it, and join sessions based on the device's target and max sessions. If the number of sessions being written to the first device exceeds the target sessions setting, another AFTD is considered for new backup sessions and is selected from the remaining suitable AFTDs. The AFTD that is selected will be the AFTD with the least amount of NetWorker data written to it. The least amount of data written is calculated in bytes (not by percentage of disk space used) and only bytes that were written by NetWorker are counted.

To ensure that a new session always writes to the AFTD with the least amount of data written to it, you can set each AFTD device's max sessions attribute to 1. However, setting the max sessions attribute to 1 may not be practical. Alternatively, set the target sessions attribute to 1. In this way, load balancing will occur on a best efforts basis.

Space management for AFTD

A configurable setting for determining at what capacity the NetWorker software should stop writing to an AFTD spans from 1 to 100%. Setting the value to 0 or leaving the attribute empty in the AFTD Percentage Capacity attribute is equivalent to a setting of 100%. This means that the entire capacity of the file system can be used for the AFTD volume.

When set, the AFTD Percentage Capacity attribute is used to declare the volume full and to calculate high/low watermarks. When the percentage capacity attribute is modified, mount and re-mount the volume for the new settings to take effect.

The level watermark is calculated based on the percentage of restricted capacity, not on the full capacity of the file system.

In the Console Administration interface, the AFTD Percentage Capacity displays in the **Configuration** tab of the **Properties** window of a device, when Diagnostic Mode is enabled.

To enable Diagnostic Mode, select **View > Diagnostic Mode**.

NOTICE If your device uses compression or deduplication, you can still use the AFTD Percentage Capacity attribute however, the device will be marked as having reached its threshold prematurely. In this case, there will be more unused space on the disk than expected. This is because the threshold limit is based on the amount of data being protected without accounting for the effect of compression or deduplication.

AFTD operation verification

The AFTD can be deployed in varying environments with local disks, and with NFS-mounted or CIFS-mapped disks. The configuration of this feature affects its operation. Ensure that the AFTD is fully operational in the production environment before deploying it as part of regularly scheduled operations.

As part of the validation process, test these operations:

- Backup
- Recover
- Staging
- Cloning
- Maximum file size compatibility between the operating system and a disk device
- Use of a volume manager to increase the file system size while the file system is in use
- File system behavior when the disk is full

Some versions of NFS and CIFS drop data when a file system becomes full. Be sure to use versions of NFS, CIFS, and operating systems that fully support full file systems. On some disk devices, the volume labeling process can take longer than expected. Labeling time depends on the type of disk device used and does not indicate a limitation of the NetWorker software. The upper limits of save set size depend on either the upper limits supported by the operating system or the file size specified by the disk device's vendor.

NOTICE Do not edit device files and directories. This can cause unpredictable behavior and make it impossible to recover data.

Deactivate and erase an AFTD

You can deactivate an AFTD device so it does not interfere with normal backup operations.

Converting a device to read-only

Conversion of a device to read-only prevents the use of the device for backup operations. The device can still be used for read operations, such as restore and clone.

Procedure

1. In the **NMC** window for the NetWorker server, click the **Devices** view and select the **Devices** folder in the navigation tree.
2. In the **Devices** table, right-click the device to be converted to read-only, and select **Unmount**.

3. Right-click this unmounted device and select **Properties**.
4. In the Device Properties window, select **Read only**, and click **OK**.
5. Right-click the device and select **Mount**.

Disabling a device

Disabling a device prevents further operation of the device. The device may be re-enabled to restore old data, which is retained but not active.

Procedure

1. In the **NMC** window for your NetWorker server, click the **Devices** view and select the Devices folder in the navigation tree.
2. In the **Devices** table, right-click the device to be disabled and select **Unmount**.
3. Right-click this unmounted device and select **Enable/Disable** to disable.
4. Inspect the **Enabled** column of the table to verify that the device is disabled.

Deleting a device

The procedure for deleting a device includes an option for also erasing the volume (access path) that stores the device's data. The volume can be erased only if no other device in the system shares the volume.

About this task

Procedure

1. In the NetWorker server **Device** view, click **Devices** in the navigation tree.
2. In the **Devices** table, right-click the device to be removed and select **Delete**.
A confirmation window appears.
3. In the confirmation window:
 - To delete the device from the NetWorker configuration only, without erasing the device's data, click **Yes**.
 - To delete the device and erase the device's data and volume access path, select the **Permanently erase all data and remove media and index information for any selected AFTDs or Data Domain devices** option, and click **Yes**.
4. If the device is mounted or the device is a member of a pool, then a second confirmation window displays the details of the device and pool. To confirm the device unmount, the removal of the device from the pool, and the deletion of the device, click **Yes**.



Note: If the volume that you want to erase is shared by another device, then an error message displays the name of the other device. You must delete all other devices that share the volume until the last one remaining before you can erase the volume.

Concurrent AFTD recovery operation limitations

AFTD concurrent recovery currently has the following limitations:

- Not available to the Windows recover interface (winworkr). Use the `recover` command. The *NetWorker Command Reference Guide* or the recover man page provides more information.
- Not available to nonfile recoveries, such as NDMP and NetWorker database modules.

- Perform concurrent recoveries from the command line by using the recover command, either by using multiple -S options to identify multiple save sets, or running multiple recover commands concurrently.

When you recover data from an AFTD, NetWorker recovers the save sets concurrently. You can recover multiple save sets to multiple clients simultaneously and you can clone save sets from an AFTD to two different volumes simultaneously.

Changing the AFTD block size

The maximum potential block size for backups to an AFTD device can be adjusted. Larger block sizes for backups can improve backup speed under certain conditions. This is especially noticeable on remote AFTD devices that are not local to the storage node, for example, AFTDs that are connected with CIFS or NFS.

About this task

Changes to the maximum potential block size value for an AFTD device take effect only after the AFTD device is labelled. The minimum allowable block size is 128 kilobytes and the maximum block size is 256 kilobytes.

If you have an AFTD device that is performing backups slowly, try marking the device as read-only and create a new AFTD device with a block size between 128-256 kilobytes.

 **NOTICE** Changing the block size and re-labeling an existing AFTD has the potential to destroy data if the data is not staged to another location.

Procedure

1. In the server's **Administration interface**, click **Devices**.
2. Select **View > Diagnostic Mode**.
3. Select **Devices** in the navigation tree. The Devices detail table appears.
4. Double-click the device in the devices table and select the **Advanced** tab.
5. In the **Device block size** attribute, select a value from 128 to 256.
6. Click **OK**.
7. Relabel the AFTD device for the new setting take effect.

DD Boost and Cloud Tier devices

DD Boost and Cloud Tier devices are covered separately in the *NetWorker Data Domain Boost Integration Guide*

Creating a DD Boost device

Procedure

1. In NMC, click **Devices**.
2. In the left panel, right-click **Devices** and select **New Device Wizard**.
3. On the **Select the Device** page, select **Data Domain** and click **Next**.
4. On the **Data Domain Preconfiguration Checklist** page, click **Next**.
5. On the **Specify the Data Domain Configuration Options** page:
 - a. Under **Data Domain System Name**:
 - Select **Create a New Data Domain System**.
 - In the text box, type the IP address of the Data Domain system.

b. In the **Data Domain DDBoost Username** field, type the username of the Data Domain user.

c. In the **Data Domain DDBoost Password** field, type the password of the Data Domain user.

In DDOS 6 and later, the default password expires in 90 days. To receive system alert reminders for password change, run the command `user password aging set <boost_user> max-days-between-change 99999`.

d. Specify the required values in the other fields.

e. Click **Next**.

6. On the **Select the Folder to Use as Devices** page:

a. Click **New Folder** to create a folder for the device.

b. Select the newly created folder.

c. Specify the required values in the other fields.

d. Click **Next**.

7. On the **Configure Pool Information** page:

a. Under **Pool Type**, select one of the following pool types:

- Backup
- Backup Clone

b. Under **Pool**, perform one of the following tasks to select the pool:

- Select **Create and use a new pool**, and type the pool number in the text box.
- Select **Use an existing pool**, and select the pool from the drop-down list box.

c. Specify the required values in the other fields.

d. Click **Next**.

8. On the **Select Storage Nodes and Fibre Channel Options** page:

a. Select the storage node.

b. Specify the required values in the other fields.

c. Click **Next**.

9. On the **Select SNMP Monitoring Options** page, specify the required field values, and click **Next**.

10. On the **Review the Device Configuration Settings** page, review the configuration settings, and click **Configure**.

11. On the **Device Configuration Results** page, click **Finish**.

Libraries and silos

NetWorker supports SCSI libraries, NDMP libraries, and ACSLS silos. In a fibre channel environment you can configure library and device sharing between storage node hosts.

Overview of tape device storage

This chapter contains information on the creation, configuration, and management of tape devices. Tape devices may be configured as stand-alone devices or configured as part of a traditional tape library or virtual tape library (VTL) storage system.

The libraries and devices available to a NetWorker server are listed in the Devices view of the NetWorker Administrator window. The details and settings of a particular device can be viewed by right-clicking the device and selecting Properties. The full range of property attributes can be viewed by selecting View > Diagnostic Mode. A description of the various attributes is provided by the Field Help button.

As with other Console functions, you can view and work with only those NetWorker servers for which you have access permission.

NetWorker software supports many different types of tape libraries, also called autochangers or jukeboxes. The general categories of libraries are SCSI, NDMP, and silo.

Support for LTO-4 hardware-based encryption

NetWorker supports the use of LTO-4 hardware-based encryption, when controlled by management utilities that are provided with the LTO-4 hardware, or by third-party key management software. EMC does not test or certify these key management utilities. The NetWorker application can read from and write to LTO-4 devices that use hardware-based encryption. The use of this encryption is transparent to NetWorker. The NetWorker application does not perform encryption or manage the key management process. For example, NetWorker does not provide the ability to turn encryption on or off or manage the encryption keys.

Linux device considerations

Review this section for information about using devices on Linux hosts.

Configure Linux operating system to detect SCSI devices

Proper configuration of the SCSI subsystem is required to get full use of SCSI devices and allow the operating system to detect SCSI devices that are attached to the computer. If the device is configured with multiple LUNs, set the kernel parameter `Probe all LUNs of each SCSI Device` to Yes.

The Linux Documentation Project website provides more information on configuring the Linux SCSI subsystem. For information on the SCSI device, contact the manufacturer.

The inquire command and the Scan for Devices operation do not detect more than 128 tape devices

By default, the Linux `st` kernel module only configures up to 128 SCSI tape devices (`/dev/nst`).

When the number of SCSI tape devices exceeds the kernel value `ST_MAX_TAPES`, the following error may appear in the `/var/log/messages` operating system log file:

```
st:Too many tape devices (max. 128)
```

The `inquire` command or the **Scan for Devices** option in NMC only displays the maximum number of `st` devices (`/dev/nst`) defined by the `ST_MAX_TAPES` value.

To resolve this issue, edit and recompile the `st` module of the Linux kernel to increase the maximum number of allowable `st` devices that are created by the OS to exceed the default value. The Linux documentation provides details on how to reconfigure, rebuild, and install the kernel.

Configuration requirements for the inquire command

Depending on the specific OS requirements and the configuration of the NetWorker server or storage node, you may need to create device files so that the `inquire` command can detect all devices.

For example, on a NetWorker server with Red Hat Linux, if devices `sg0` through `sg15` exist, create device file `sg16` by using the `mknod` program as follows:

```
mknod /dev/sg16 c 21 17
```

The operating system vendor documentation provides more information on creating devices.

Solaris device considerations

Review this section for information about using devices on Solaris hosts.

Support for tape devices not supported by Solaris

If Sun Microsystems does not directly support a device for use with the operating system on the storage node, obtain a `st.conf` file from the device manufacturer.

The inquire command and Solaris 10

On Solaris 10, the `inquire` command does not show library information after you configure the library for NetWorker.

HP-UX device considerations

Review this section for information about using devices on HP-UX hosts.

Autochanger installation on an HP-UX system

The following sections explain how to install and configure Hewlett-Packard drivers.

Selecting SCSI addresses for the autochanger

Determine which SCSI address is assigned to each SCSI bus, and select the SCSI addresses to be allocated to the autochanger drives and controller.

About this task

To select unused SCSI addresses for an autochanger, log in as root on the NetWorker server or storage node, and type the `ioscan -f` command.

Use a SCSI address within the range of 0 to 6. The primary hard disk is usually on SCSI address 6.

NOTICE For some devices, such as the HP Model 48AL autochanger, select one SCSI address for the entire autochanger. The 48AL uses a different SCSI logical unit number (LUN) for the device (LUN 0) and robotics (LUN 1). The SCSI LUN appears as the last digit of the `H/W Path` field in the `ioscan` output.

Installing the SCSI pass-through driver

The following procedure describes how to use SAM terminal mode to install a GSC, HSC, or PCI pass-through driver.

Procedure

1. Select **Kernel Config** and press **Enter**.
2. Select **Drivers** and press **Enter**.
3. Select the `SCSI_ctl` driver by selecting **SCTL** from the list.

If the current state is `in`, go to step 9. Otherwise, select any unreserved name for the device. For example, do not select a name such as `/dev/null`.

4. From the **Actions** menu, select **Add Drivers to Kernel**, and press **Enter**.
5. From the **Actions** menu, select **Create a New Kernel**, and press **Enter**.

A confirmation message appears.

6. Specify **Yes**, and press **Enter**.

The **Creating Kernel** message appears, followed by the **Move Kernel** message.

7. Select **OK**, and press **Enter**.

The system reboots.

8. Verify that the `spt` was successfully installed by typing the following command:

```
ioscan -kfn
```

9. Verify that the driver has claimed the autochanger.

If the autochanger has been claimed, `CLAIMED` should appear under the `S/W State` header. If not, verify that the installation completed correctly.

10. If the device entry was defined by the operating system, use the OS-defined entry and continue to verify the installation.

Determining the major number

To determine the value for `majornum`, type `lsdev -d sctl`.

About this task

The output should resemble the following example output, although the assigned number may differ from the values in this example:

Table 28 Determining the major number value

Character	Block	Driver	Class
HP-PB	75 -1	spt	spt
HSC or PCI	203 -1	sctl	ctl

The value for `majornum` is the number in the `Character` column.

Determining the minor number

To determine the value for `minornum`, use the `ioscan` command.

About this task

The relevant lines in the `ioscan` output are those:

- For the controller itself, which contains HP C6280-7000 in the `Description` column.
- For the adapter to which the controller is connected, which is the second line above the line for the controller and contains `ext_bus` in the `Class` column.

If the `schgr` driver is configured on the system, it appears associated with the library. The `ioscan` output line resembles:

Table 29 `ioscan` output when driver is configured

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
spt	0	10/4/4.6.0	schgr	CLAIMED	DEVICE	HP C6280-7000

If the `schgr` driver is not configured on the system, no driver appears to be associated with the library. The `ioscan` output line resembles:

Table 30 `ioscan` output when driver is not configured

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
unknown	-1	10/4/4.6.0	schgr	UNCLAIMED	DEVICE	HP C6280-7000

Testing the device driver and device file installation

After the device driver is installed and the device file is created, run the `inquire` command to list available SCSI devices.

About this task

NOTICE Use the `inquire` command with caution. Running `inquire` sends the SCSI inquiry command to all devices detected on the SCSI bus. Using the `inquire` command during normal operations may cause unforeseen errors and possible data loss may result.

An example of the output from this command (with the `-s` option) is as follows:

```
scsidev@0.1.0:HP C1194F 0.14 Autochanger (Jukebox), /dev/rac/
c0t1d0
scsidev@0.2.0:Quantum DLT4000 CC37 Tape, /dev/rmt/c0t2d0BESTnb
scsidev@0.3.0:Quantum DLT4000 CC37 Tape, /dev/rmt/c0t3d0BESTnb
scsidev@0.4.0:Quantum DLT4000 CC37 Tape, /dev/rmt/c0t4d0BESTnb
scsidev@0.5.0:Quantum DLT4000 CC37 Tape, /dev/rmt/c0t5d0BESTnb
```

As of HP-UX 11iv3, two different addressing modes are supported: `LEGACY` and `AGILE`. The `inquire` program lists devices using the B.T.L. notation for the `LEGACY` addressing mode, for example:

```
scsidev@B.T.L.
```

For the `AGILE` addressing mode, it lists devices using the DSF notation, for example:

```
/dev/rtape/tape106_BESTnb
```

Inquire command does not detect tape drive

When a tape drive is attached to the HP-UX 11i V2 64-bit host and the `inquire` command is run, the tape drive is not detected, even if the device is configured, labeled, and mounted and a save was successful.

To work around this issue, identify the drive path in the `/dev/rmt` folder, and configure the device with this path.

Whenever a new device is attached to the system, ensure that the cached file `/tmp/lgto_scsi_devlist` is updated. Remove this temp file and then run the `inquire` command, which rebuilds the file.

Errors from unsupported media in HP tape drives

Certain HP tape drives can only read 4-mm tapes of a specific length. Some, for example, read only 60-meter tapes. To determine the type of tape that is supported, refer to the drive's hardware manual.

If unsupported media is used, the following error message may appear when you use the `nsrmm` or `nsrjb` command to label the tape:

```
nsrmm: error, label write, No more processes (5)
```

The following error message may appear when you use the `scanner -i` command to label the tape when unsupported media is used:

```
scanner: error, tape label read, No more processes (11)
scanning for valid records ...
read: 0 bytes
read: 0 bytes
read: 0 bytes
```

Unloading tape drives on an HP-UX server or storage node

When the `nsrjb -u -s` command is used to unload a tape drive in an autochanger that is attached to an HP-UX server or storage node, the unload operation ejects all tape volumes inside the autochanger devices, and into their respective slots.

To unload a single drive to its corresponding slot, use the `nsrjb -u -f device_name` command instead.

SCSI pass-through driver required for HP-UX autochangers

Review the required procedures in the before you run the `jbconfig` program to configure an autochanger with a NetWorker server on HP-UX.

Follow the procedures to rebuild the kernel even if the SCSI pass-through driver is installed. Then run the `jbconfig` program to configure the *NetWorker Installation Guide* autochanger.

AIX device considerations

Review this section for information about using devices on AIX hosts.

STK-9840 drives attached to AIX

If you attach an STK-9840 drive to an AIX server, use SMIT to modify the IBM tape drive definition field to set the value of `Use Extended File Mark` to Yes.

LUS driver operation on AIX

When a library comes online, NetWorker obtains an exclusive lock on the library due to the operation of the LUS driver on AIX. This lock is maintained if the library is enabled. As a result, you cannot use diagnostic tools such as `inquire` and the `sji` utilities to access the library during this time. To access the library using these tools, you must first take the library offline.

SCSI and VTL libraries

SCSI libraries have automated robotic mechanisms to move tape media from a fixed number of library slots to devices for read or write operations. The number of slots can typically vary between 2 to 10,000 and the number of devices can be between 1 to 100 or more.

Traditionally, libraries are physical units with mechanical robotics, however the same functionality can also be provided by virtual tape libraries (VTLs) that emulate this functionality. VTLs can also be configured and used as Autochangers.

The robotic controller and associated tape devices are always all controlled through a SCSI interface which is available on one or more storage hosts.

Selecting a volume for the NetWorker server

When a backup takes place, the NetWorker server searches for a volume from the appropriate pool to accept the data for backup.

About this task

The available volumes are as follows:

- Mounted on stand-alone devices.
- Available for labeling and accessible to the NetWorker server through Auto Media Management or a library.
- Labeled for the appropriate pool and already mounted in a device, or are available for mounting, if a library is being used.

If two or more volumes from the appropriate pool are available, the server uses this hierarchy to select a volume.

- A volume in a jukebox device has priority over volume in a disk or tape device.
- A volume in a local disk device has priority over a volume in a local tape device.
- If two local disk are available, then the device less data sessions will have priority.
- If two local tapes devices have available volumes, then NetWorker will use the volume with the earliest label date.
- If two jukebox are available, then NetWorker will select the volume with the earliest label date..

Data recovery and volume selection

The NetWorker server determines which volumes are required for recovery. If the appropriate volume is currently mounted, the recovery begins. If the volume is not mounted and a library is used, the server attempts to locate and mount the volume in an eligible device for appropriate media pool. Preference is given to mount the volume in a read-only device, if one is available.

If a stand-alone device is used, or if the server cannot locate and mount the volume, the server sends a mount request notification.

If more than one volume is needed to recover the data, the NetWorker server displays all the volumes, in the order needed. During the recovery process, the server requests the volumes, one at a time.

NOTICE NetWorker will automatically unload volumes that have been placed in a jukebox device but have never been mounted (for example, `nsrjb -l -n <volume>`). Any command, such as the scanner command, that operates on volumes that have never been mounted will be affected by this behavior. To prevent NetWorker from unloading the volume, the device should be set to service mode while the command is being run.

Automatic volume relabeling

NetWorker has the ability to automatically relabel recyclable volumes when needed or when scheduled.

When you enable Auto Media Management, the NetWorker server will automatically relabel a volume when the mode is recyclable. A volume is automatically set to recyclable when all save sets on the volume, including partial save sets that span other volumes, are marked as recyclable. [Auto Media Management](#) on page 146 provides more information on Auto Media Management.

 **Note:** You can manually change the mode of a volume to recyclable. [Changing the volume mode](#) on page 537 provides information about changing the mode of a volume.

You can configure a media pool to automatically relabel recyclable volume at a user defined time and interval. [Automatically relabeling volumes in a media pool](#) on page 93 provides more information about configuring the automatic relabel process for recyclable volumes in a media pool.

Virtual tape library (VTL) configuration

During library configuration, the NetWorker software automatically attempts to detect if a library is a VTL, and updates the read-only Virtual Jukebox attribute to Yes, or if not, to No. VTLs that are mistakenly identified as autochangers can indicate what type of license should be used, either autochanger or VTL.

VTL licensing

The *NetWorker Licensing Guide* provides information about NetWorker licensing support for a Virtual Tape Library.

Multiplex backups to Data Domain VTL devices

You can configure multiplexed backups to Data Domain VTL devices on remote, non-dedicated NetWorker storage nodes. Multiplexing is the use of multiple parallel save streams or concurrent sessions to each device. Each additional save stream (max sessions value) to a VTL device reduces the number of devices needed by somewhat less than one because deduplication efficiency decreases slightly.

The following prerequisites, restrictions, and considerations apply:

- NetWorker dedicated storage nodes (DSNs) and NetWorker backup to local VTLs cannot use this configuration.
- Multiplexing decreases deduplication efficiency on the VTLs by 4% to 8% per additional save stream. For example, given a sufficiently large device block size, 4 parallel streams (max sessions=4) results in deduplication ratios that are 12%-24% below the non-multiplexed rate (max sessions=1).
- Deduplication ratios may be initially low when you increase max sessions due to extra processing, following which efficiency improves.
- Heavily used Data Domain systems, with 75% or more disk space already used, can suffer impaired performance when used with multiplexing.
- As a best practice, do not use client-side or server-side encryption during backup to the Data Domain system.

Multiplex to Data Domain VTL prerequisites and considerations

Ensure the following prerequisites and practices.

- If currently using DD OS 5.0.x, upgrade to DD OS 5.7 or later.
- The recommended settings for VTL are: max sessions=4; target sessions=4; and device block size=512 KB.
- Best max sessions and device block size values depend on the environment. For example, max sessions=2 might provide better stability and deduplication while still meeting the backup window.
- Deduplication efficiency on the VTLs is reduced by 4% to 8% per additional save stream. For example, given a sufficiently large device block size, 4 parallel streams

(max sessions=4) results in deduplication ratios that are 12%-24% below the non-multiplexed rate (max sessions=1).

- Typically, deduplication ratios are initially low when you increase max sessions and device block size due to re-priming and re-analysis overhead, following which efficiency improves.
- Heavily used Data Domain systems, with 75% or more disk space that is already used, can suffer impaired performance when used with multiple sessions.
- As a best practice, do not use client-side or server-side encryption during backup to the Data Domain system.

Configuring multiplex backup to Data Domain VTL devices

Configure Data Domain VTL devices for multiple session backups as follows.

Procedure

1. Shut down backup service on the NetWorker VTL storage node, or shut down the NetWorker server if that is possible, and verify that there is no backup activity on the storage node.
2. Use NMC or the `nsradmin` command to set the sessions values for each VTL device. The recommended values are as follows:
 - Max sessions=4 (32 maximum)
 - Target sessions=4
 - Device block size=512KB

Optimal max sessions and device block size values depend on the environment. For example, max sessions=2 might provide better stability and deduplication while still meeting the backup window.

Note: If you shut down the NetWorker server in step 1, you can run the `nsradmin` command with the `-d resdir` option. This option uses the NetWorker resource database, `resdir`, without opening a network connection.

For example, on UNIX/Linux or Microsoft Windows systems, run the following command:

```
nsradmin -i input_file.txt
```

where *input_file.txt* is a text file that contains the following lines that you can customize to the own environment:

```
option regexp: on
. type: nsr device; media type: LTO Ultrium-3; media family:
tape; name: /dev/rmt*
update max sessions: 4; target sessions: 4; device block size:
512KB
```

3. Create a no intra-block multiplexing (nibmp) tag file in the NetWorker `debug` folder on the NetWorker storage node.

For example, you can use the standard NetWorker installation paths for the tag file. You can limit the tag file path to a specific pool by adding the `_poolname` variable as a suffix to the tag file. The `_poolname` can include spaces, for example, `_My Pool`. On Microsoft Windows systems, ensure that the specified pathname is enclosed in quotes.

Unix/Linux system examples.

```
touch /nsr/debug/nibmp
touch /nsr/debug/nibmp_My Pool
```

Microsoft Windows system examples.

```
echo > "NetWorker_install_path\nsr\debug\nibmp"
echo > "NetWorker_install_path\nsr\debug\nibmp_My Pool"
```

4. Restart the NetWorker services to enable the multiplexing functionality.

The technical note that is named , available on the Online Support website, provides more details.

Non-rewinding tape device usage (UNIX/Linux only)

Tape drives used as storage devices must be accessed by non-rewinding device files. The NetWorker server assumes that a tape is in the same position in which it was the last time it was accessed. If the operating system's device driver rewinds the tape, then the position is lost, and previously written data will be overwritten by the next backup.

The NetWorker configuration software automatically chooses the correct device pathname for tape devices. If the user specifies the pathname, then it must be non-rewinding, and it must follow the Berkeley Software Distribution (BSD) semantic rules.

For example, `/dev/rmt/0mbn`, where:

- The `b` satisfies the BSD semantics requirement on Solaris and HP-UX.
- The `n` specifies non-rewinding behavior on Solaris, HP-UX, Linux, and HP-Tru64.

On AIX, the number following the decimal selects the BSD and non-rewinding behavior and must be either 1 or 5 for NetWorker software (for example `/dev/rmt2.1`).

Note: Never change a device pathname from non-rewinding (`/dev/rmt/0cbn`) to rewinding (`/dev/rmt/0cb`). When the pathname is changed to rewinding, the data could only be saved, but never recovered. All but the last save are overwritten by later saves.

Pools with libraries

If the backup strategy includes both full and nonfull backups, estimate the number of volumes needed for the full backups and assign them to the Full pool. This ensures that the full backups are located in a consecutive range of slots in the library. This allows all of the volumes to be removed at the same time.

Persistent binding and naming

Some operating systems provide the persistent binding option to permanently bind logical and physical addressing so that the associations are retained. This guarantees that the operating system always uses and creates the same symbolic path for a device is known as persistent naming.

Proper configuration of the operating system to use persistent binding and persistent naming resolves issues related to device ordering by forcing the operating system to always assign the same device filename regardless of external events.

Persistent binding

Persistent binding guarantees that the operating system always uses the same SCSI target ID for SAN devices, regardless of reboots or other events, by statically mapping a target's WWN address to a desired SCSI address. On some operating systems, this

is done by default, while on others it has to be set manually. The operating system documentation provides further information.

In most cases, persistent binding should also be set on the Host Bus Adapter (HBA) by using the configuration utility that comes with the Fibre Channel HBA. The HBA device driver documentation provides details.

Persistent binding is required for consistent library operations within NetWorker, because the NetWorker server communicates with the library controller over a SCSI address that is chosen during initial library configuration. If the SCSI address changes, the library will become unavailable. In this case, disable the library and change the “control port” address to reflect the new SCSI address of the library controller.

If devices have already been configured in NetWorker prior to enabling persistent binding on the host, delete existing devices from the library resource and perform a re-scan of devices followed by a reconfiguration of the tape library.

Persistent naming

Persistent naming is used to ensure that the operating system or device driver of a server always creates and uses the same symbolic path for a device (referred to as device file).

After you create persistently named device files and they are present on the host, enable the **Use persistent names** option when scanning for tape devices from the NetWorker Management Console.

If devices have already been configured in NetWorker prior to enabling persistent naming on the host, delete existing devices from the library resource and perform a re-scan of devices followed by a reconfiguration of the tape library.

Whether to add or recycle volumes

The NetWorker server saves files on volumes marked appen (appendable). If the volumes are marked full, they cannot receive backups. There are situations best suited to either adding a new volume, or recycling an existing volume.

If volumes are marked full, you can:

- Remove the full volumes and replace them with new media if the volumes are being kept for long-term storage.
- Change the volume mode to recyc (recyclable) if the data on the full volumes is not needed. The NetWorker server overwrites the data with new backups, but maintains the existing labels. [Changing the volume mode](#) on page 537 provides information about changing the volume mode.

When all of the save sets on the volume have passed the time period specified by the retention policy, the mode of the volume automatically changes to recyclable.

There are advantages both to recycling media and adding more media to a pool. With recycling, the same volumes are used repeatedly, and there is no need to add new volumes to the pool. The volumes can, however, wear out over time and exhibit a higher failure rate.

On the other hand, if backups are to be stored for some time, then it might be necessary to add more media to the pool instead of recycling. For example, a library might need new volumes every three months if the company policy is to maintain the backups for a year. In this case, new media must be added to the pool until the volumes that contain expired or old backups can be recycled.

Configure libraries

A library resource must be created on a storage node for each library, including silos, that you want to use with NetWorker. Because the NetWorker server is also a storage node, this procedure applies to a NetWorker server and all storage nodes. You can configure a library either automatically with the **Configure All Libraries** wizard or manually with the user interface.

Before you create devices, you must create the storage node that will manage the devices. [Storage nodes](#) on page 98 provides details. When you create the new devices, you can use NetWorker to perform a device scan, which searches for new devices across multiple storage nodes.

NetWorker can only automatically create tape devices that have serial numbers. Use the `inquire` or `sn` commands to determine if a device returns a serial number. UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about how to use the `inquire` and `sn` commands.

NetWorker can automatically configure the following library types:

- SCSI
- NDMP
- ACSLS Silo

Use the `jbconfig` command to configure a library that contains tape devices or a robotic arm that does not have serial numbers. Use the `jbconfig` command to configure IBM tape libraries that are controlled through the use of the IBMs tape driver. This is because the device autodetection code uses the internal `lus` driver to control libraries.

 **Note:** Before you create devices on a storage node, update the devices to the most recent firmware and driver versions.

Autodetection of libraries and tape devices

Autodetection is a scanning process that applies only to physical tape libraries and virtual tape libraries (VTLs). The NetWorker software automatically discovers libraries and devices that are being used for backups and recoveries.

The maximum number of configured devices for any NetWorker server and storage node combination is 750. The maximum number, including non-configured devices, can vary depending on the specific server that is being administered.

The following options are available from many of the menus throughout the Devices task:

- Configure all Libraries
- Scan for Devices

If you start these options from the server folder instead of from the storage node folder, then all storage nodes on the NetWorker server are automatically selected for configuration in the wizard, or for scanning, respectively.

As with other Console functions, you can view and work with only those NetWorker servers for which you have access permission.

 **NOTICE** Autodetection should not be used for devices on a Storage Area Network (SAN) while any of the devices are in use, because this may cause the device in use to become unresponsive. To avoid this situation, do not configure a device in multiple NetWorker datazones.

Adding a library resource

Procedure

1. In the server's **Administration** interface, click **Devices**.
2. Open the **Storage Nodes** folder in the navigation tree.
3. Right-click the storage node to which the device is to be configured, and select **Configure All Libraries** (which is available from many of the menus throughout the **Devices** task). This opens a wizard that can configure all detected libraries, except those explicitly excluded in the library exclusion list during configuration.

 **NOTICE** If Configure All Libraries is started from the server folder instead of from the Storage Node folder, then all storage nodes on the NetWorker server are automatically selected for configuration in the wizard.

The Configure All Libraries wizard appears. This lets you step through library configuration, including this input (some of which is filled in by default):

- Library type (select SCSI/NDMP).
 - An NDMP remote username and a password are required for an NDMP device that acts as a storage node.
 - Adjust the **Enable New Device** option, if necessary.
 - Current server sharing policy. Use maximal sharing with Dynamic Drive Sharing (DDS). By default, the sharing policy is displayed as “server default,” which is maximal sharing.
 - Storage nodes to which libraries can be configured (select a storage node to see its details). If the appropriate storage node is not listed, click **Create a New Storage Node**.
 - When creating a new storage node, replace the default value in the **Name** field with the fully-qualified domain name or short name of the new storage node.
 - Update storage node properties, if required.
4. After specifying the required information, click **Start Configuration**. The configuration window displays a message that the Configure All Libraries process has started. The status of the configuration activity can be viewed by the **Monitoring > Log** screen.
 5. When the configuration is complete, click **Finish** to close the configuration wizard. If problems occur during configuration, you can click the **Back** button on the configuration window to adjust the settings.

Scanning for libraries and devices

Devices already known to the NetWorker server can be seen in the enterprise hierarchy in the navigation tree. Use the Scan for Devices option described here to find devices that are not yet known to the NetWorker server. Be aware that:

About this task

- A storage node must be added to the hierarchy before its devices can be scanned.
- The Scan for Devices option does not detect file type or advanced file type devices.
- By default, the Linux kernel configures a maximum of 128 st devices by default. Refer to [The inquire command and the Scan for Devices operation do not detect more than 128 tape devices](#) on page 131 if the Scan for Devices option does not detect more than 128 tape devices on Linux operating systems.

- A specific network interface can be used between the NetWorker server and the storage node when scanning for devices. [Identifying a specific network interface for device scan operations](#) on page 145 provides more information.

Procedure

1. In the **Console** window, click **Enterprise**.
2. In the navigation tree, select a NetWorker server.
3. In the **Name** column of the **Host detail** table, double-click **NetWorker**. The **NetWorker Administration** window for the selected server opens. Note that while multiple **NetWorker Administration** windows can be open simultaneously, each one displays information about only one host or server.
4. In the **Administration** window, click **Devices**.
5. In the navigation tree:
 - a. Right-click the server name, and select **Scan for Devices**.
 - b. Click the storage node to be scanned.
 - c. If the appropriate storage node is not listed, click **Create a New Storage Node**.
 - d. When creating a new storage node, replace the default value in the **Name** field with the fully-qualified domain name or short name of the new storage node.
 - e. Fill in any required information, such as whether to scan for SCSI or NDMP devices and whether to search all LUNs.
 - f. Click **Start Scan**. To monitor the scan activity, click **Monitoring**, then select the **Log** tab. Any relevant status information is displayed there.
6. Return to the **Devices** navigation tree to view the refreshed device information (configured and unconfigured):
 - To display SCSI and NDMP libraries available to the NetWorker server, select **Libraries** in the navigation tree. Any available library or silo appears in the Libraries detail table.
 - To display stand-alone devices available to the NetWorker server, select **Devices** in the navigation tree. Any available stand-alone device appears in the **Devices** detail table, along with devices available in libraries.
 - To display the libraries and devices that are available to a storage node, select the storage node in the navigation tree. Available storage nodes appear in the table. Double-click a storage node to see its details, along with the devices that are available in the storage node.

Barcode labeling tips

The NetWorker server uses volume labels and barcode labels to identify volumes. Both label types are recorded in the media database. The volume label is also recorded internally on the media (internal volume label). The NetWorker server uses barcode labels to inventory volumes, and uses volume labels to identify the volumes needed for backup and recovery. A requirement to match the volume label with the barcode label can be set in the library's **Properties** window.

Follow these guidelines when using barcode labels with the NetWorker software:

- When NetWorker software relabels volumes automatically, it reuses the original volume label name. A label name can be changed only if the volume is relabeled

manually. The NetWorker software scans the barcode label during the labeling process and updates the media database with the new volume name and its associated barcode label.

- Do not use identical barcode labels for any of the NetWorker volumes. The use of identical labels defeats the purpose of using barcode labels, which is to facilitate the inventory process and ensure label accuracy.
- Volume names must be unique on the NetWorker server. Give each volume a unique volume label. If a second volume is labeled with an existing barcode label and the Match Barcode Labels attribute in the library's properties is enabled, the NetWorker server displays an error message and does not allow the second volume to be labeled. The error message identifies the library slots containing the two volumes with identical labels and the barcode label.
To correct this problem, either apply a different label to one of the volumes and restart the labeling process, or disable the Match Barcode Labels attribute in the library's properties while labeling the second volume.
- It is not necessary to label existing volumes with barcode labels if they are stored in a vault or offsite for long periods. These volumes are rarely, if ever, inventoried.
- Before using barcode labels on existing volumes, affix the barcode labels to them. Then, load and mount each volume individually, so that the NetWorker server can match the barcode label with the existing volume label.
- Record the volume label on the tape.
- A variety of barcode labels can be purchased from third-party vendors. Choose from among numeric labels, alphanumeric labels, or a special combination of numbers and characters. Furthermore, barcode labels can be ordered to match a current volume labeling scheme.
- Use a consistent labeling scheme. If volumes are labeled with the server name and an extension such as "001," order a range of labels starting with "server_name.001" and ending with "server_name.100", or as wide a range as necessary. Instructions for barcode labels should be provided with the library hardware documentation. Contact the hardware manufacturer with questions about barcode labels. A consistent labeling scheme helps better organize and track volumes. It also facilitates the inventory process if all of the volumes, use barcode labels.

Configuring a library to use volumes with barcodes

About this task

[Barcode labeling tips](#) on page 143 provides more information.

Procedure

1. In the **Administration** window, click **Devices**.
2. Open the **Libraries** folder. The **Libraries** detail table appears.
3. Right-click the appropriate library, and select **Properties**. The **Properties** window appears.
4. Select the **Configuration** tab.
5. In the **Media Management** area of the **Configuration** tab, select:
 - Bar Code Reader
 - Match Bar Code Labels
6. Click **OK**.

Using unmatched volume and barcode labels

About this task

 **Note:** If unmatched volume and barcode labels are to be used, ensure that labels are attached to the outside of the volumes.

Procedure

1. Apply barcode labels to the volumes.
2. Place the volumes with the barcode labels in the library.
3. In the **Administration** window, click **Devices**.
4. Open the **Libraries** folder. The **Libraries** detail table appears.
5. Right-click the appropriate library, and select **Properties**. The **Properties** window appears.
6. Select the **Configuration** tab.
7. In the **Media Management** area of the **Configuration** tab:
 - Select **Bar Code Reader**.
 - Ensure that **Match Bar Code Labels** is not selected.
8. Click **OK**. The NetWorker server uses the next available label from the label template for the volume name. It labels the volumes and records both labels in the media database.
9. Inventory the volumes to ensure that the NetWorker server has the most current volume information.
10. Use **Media > Volumes** to match the correct volume labels to the barcode labels. Consider making a list of the name correlations.

 **Note:** If the barcode function is enabled, but no barcode label is affixed to the volume, an error message indicates that a barcode label does not exist.

Identifying a specific network interface for device scan operations

If the NetWorker server has multiple network interfaces, you can specify that a specific network interface be used for scan operations. In this case, the dvdetect (device scan) program will use the specified network address or hostname to communicate with the NetWorker server.

Procedure

1. In the server's **Administration interface**, click the **Devices** button.
2. Select **View > Diagnostic Mode**.
3. In the left pane, click on the **Storage Nodes** folder.
4. In the right pane, select a storage node.
5. Right-click the storage node and select **Properties**.
6. Select the **Configuration** tab.
7. In the **Server network interface** field, type the network address or the unique hostname of the network interface on the NetWorker server that is to be used.
8. Click **OK**.

Media Library parallelism

To define the media library parallelism, use the **Max parallelism** attribute on the **Configuration** tab of the Library resource .

Media library parallelism allows you to define the maximum number of available devices for inventory and label operations.

It is recommended that you set the **Max parallelism** attribute of the Library resource to one less than the number of devices within the library, which allows you to reserve one device for recovery operations.

To improve the efficiency of library operations that operate on multiple volumes, use multiple devices in parallel for these operations. However, you may want to restrict the number of devices that NetWorker uses for inventorying and labeling operations, to ensure that some devices are available for other library operations.

Managing the library configuration

This section provides detailed information about managing a tape library in the NetWorker environment.

Auto Media Management

Auto Media Management gives the NetWorker server automatic control over media that are loaded in the storage device.

When you enable the Auto Media Management feature during device configuration, the NetWorker server automatically:

- Labels the volume (recognizes EDM labels and does not overwrite them).
-  **NOTICE** If the Auto Media Management feature is enabled, the NetWorker server considers volumes that were labeled by a different application to be valid re-label candidates. Once the NetWorker server re-labels the volume, the previously stored data is lost.
- Mounts the volume.
- Overwrites volumes that are considered to be unlabeled. The NetWorker server considers a volume to be unlabeled under the following conditions:
 - Has no internal label.
 - Is labeled with information other than a NetWorker label.
 - Is labeled with a NetWorker label, but the density that is indicated on the internal label differs from that of the device where the volume is mounted.
- Recycles volumes eligible for reuse that are loaded into the device.

When you do not enable the Auto Media Management feature, the NetWorker server ignores unlabeled volumes and does not use the volume for backup.

The Auto Media Management feature can re-label a volume that has a different density, it is possible, inadvertently, to overwrite data that still has value. For this reason, be careful if NetWorker volumes are shared among devices with different densities.

Existing tapes with NetWorker labels

When Auto Media Management is used with tapes that have NetWorker labels that have not been recycled, the volumes must be removed from the media database before a utility such as tar is used to overwrite the labels. Also ensure that the tapes have been fully rewound before overwriting the labels. Auto Media Management can then properly relabel the tapes.

Auto Media Management for stand-alone devices

The Auto Media Management feature can be enabled for stand-alone devices during manual device configuration, or from the **Properties** window after configuration.

When Auto Media Management is enabled for a stand-alone device, the following processes occur when a volume becomes full during a backup:

- A notification is sent that indicates that the server or storage node is waiting for a writable volume. Simultaneously, the NetWorker server waits for the full, verified volume to be unmounted.
- The device is monitored and the software waits for another volume to be inserted into the device.
- After a volume is detected, a check is performed to determine whether the volume is labeled. If so:
 - The volume is mounted into the device.
 - The NetWorker server checks to see whether the newly mounted volume is a candidate to receive data:
 1. If yes, the write operation continues.
 2. If no, the NetWorker server continues to wait for a writable volume to continue the backup.
- If the volume is recyclable and is a member of the required pool, it is recycled the next time a writable volume is needed.
- If the volume is unlabeled, it is labeled when the next writable volume is needed for a save. Note that Auto media management does not label disk type devices such as AFTD and Data Domain.

 **NOTICE** If a partially full volume is unmounted, the NetWorker server automatically ejects the volume after a few seconds. If a stand-alone device is shared between storage nodes, then Auto Media Management should not be enabled for more than one instance of the device. Enabling Auto Media Management for more than one instance of the stand-alone device will tie up the device indefinitely. No data is sent to the device and no pending message is sent.

Enabling Auto Media Management for libraries

Auto Media Management is not enabled for libraries during autoconfiguration. Auto Media Management for a library can be set by changing the library's properties after configuration.

Procedure

1. In the server's **Administration** window, click **Devices**.
2. Select the **Libraries** folder in the navigation tree. The Libraries detail table appears.
3. Right-click the library, and select **Properties**. The **Properties** window appears.
4. Select the **Configuration** tab.
5. In the **Media Management** area, select **Auto Media Management**.
6. Click **OK**.

Labeling volumes

The NetWorker software applies a label template to create a unique internal label for each volume. The label corresponds to a pool and identifies the pool for the volume during backup and other operations.

Several preconfigured label templates are supplied with the NetWorker software. You cannot delete these preconfigured label templates. [Naming label templates](#) on page 79 provides more information.

When you label a volume, the labeling process:

- Writes a label on the volume.
- Adds the volume label to the media database.
- Prepares tape media to have data written to it.

When you re-label tape, the data on the tape is effectively gone.

During data recovery, the server requests the volume that contains the required data, identifying the required volume by the name with which it was labeled.

Label templates

Several preconfigured label templates are supplied with the NetWorker software. These preconfigured label templates cannot be deleted. [Naming label templates](#) on page 79 provides more information about label templates and preconfigured label template.

Labeling or re-labeling library volumes

Labeling volumes in a library is time-consuming, so consider labeling volumes before it is time to back up or recover files. When a volume is re-labeled, that volume is initialized and becomes available for writing again.

Procedure

1. In the **Administration** window, click **Devices**.
2. In the left pane, select **Libraries**.

A list of libraries appears in the right pane.

3. Right-click the library and select **Label**.

Details for the selected library appear, including divided tables for devices and slots. The **Label Library Media** dialog box also appears.

4. From the **Target Media Pool** list, select the pool for the volume.

The pool determines the label template that is used to label the volume.

5. To require manual recycling of the volume, select **Allow** > **Manual Recycle**.

With manual recycling, the volume is not automatically marked as recyclable when all save sets expire. You must manually mark the volume as recyclable.

 **NOTICE** A volume that has been set to manual recycle retains that setting, even after the volume is re-labeled. You must explicitly reset the volume to automatic recycle by right-clicking the volume in the **Media** window, selecting **Recycle**, and then selecting the **Auto** option.

6. To be prompted before the existing label is overwritten, select **Prompt to overwrite label**.
7. Click **OK**.

The **Library Operation** dialog box appears, stating that the library operation has started.

8. To track the status of the label operation, click **Monitoring** in the **Administration** window.
9. If you selected **Prompt to overwrite label**, confirm the overwrite of the existing volume label with a new label:
 - a. Right-click the label operation in the **Monitoring** window and select **Supply Input**.
A confirmation message appears.
 - b. Click **Yes**.

Verifying the label when a volume is unloaded

If a SCSI reset is issued during a backup, the volume rewinds and NetWorker may overwrite the volume label.

To detect if the label is overwritten in this circumstance, select the **Verify label on eject** checkbox in the Device resource, or set the **Verify label on unload** setting in the Jukebox resource to **Yes**. With these settings, NetWorker verifies that a volume label exists before ejecting the volume. If the volume label cannot be read, all save sets on the volume are marked as suspect and the volume is marked as full.

Empty slots in label operations

Slots that have been intentionally left empty (such as bad slots) are skipped during labeling operations. The NetWorker software logs a message similar to: "Slot 5 empty, skipping."

Barcode labels

The option to label a library volume with a barcode is available during automatic device configuration. This option can be set in the library's Properties tab after configuration.

Barcode labels make volume inventory fast and efficient. They eliminate the need to mount the volumes in a device. The library scans the external barcode labels with an infrared light while the volumes remain in their slots. Inventorying with barcode labels greatly reduces the time needed to locate a volume or determine the contents of a library.

Barcode labels also provide greater labeling accuracy. The labels are placed on the volumes before the volumes are loaded and scanned in the library. Once the library has scanned the barcode, the NetWorker server records and tracks the label in the media database. The NetWorker server uses barcode labels only to inventory volumes. A volume must have a label, but it need not have a barcode label.

 **Note:** Libraries include hardware that reads barcode labels. The barcode information is then forwarded to the NetWorker server. Problems reading barcode labels indicate hardware problems. In the event of a barcode-related problem, consult the library's documentation or the hardware vendor.

Requirements for performing an inventory with barcodes

To perform an inventory by using barcodes, the following requirements must be met:

- The library must have a barcode reader.
- A barcode label must be present on the tape.
- The location field within the NetWorker media database must be correct or null. To view the location field, use the mmlocate command.

Device Service mode

Use the service mode setting to take a device offline temporarily. Service mode differs from the disabled state in that the `nsrmmmd` process is not stopped.

While a device is in service mode, save or recover sessions that are either in process or pending are completed. No new sessions are assigned to the device while it is in service mode.

Although a drive in service mode is taken out of the collection of drives that the NetWorker software can select for automated operations, the drive is available for some manual operations that use the `nsrjrb` or `nsrmm` command with the `-f` option. For more information, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

The device might also go into service mode, rather than become disabled, if consecutive errors occur in excess of the maximum consecutive error count specified for the device. This means that if there are no hardware issues, the tape can be ejected and used in other drives. [Media handling errors](#) on page 172 provides more information about how to set the maximum consecutive error count.

 **Note:** The drive must be manually reset to Enabled for the NetWorker software to use the device again.

Setting the Service mode for a device

Procedure

1. Open the device's **Properties** window.
2. On the **General** tab, set **Status Enabled** to **Service**.

Reconfiguring a library

Use this procedure to reconfigure a tape library.

Before you begin

To reconfigure a library or to add or remove access paths to the devices in a library, use an account with the Configure NetWorker privilege. This includes access paths that allow libraries to be shared.

About this task

 **Note:** The following procedure does not support adding NDMP devices to a non-NDMP library if both the NDMP server and the NetWorker storage node are on the same host. Instead, use the `jbedit` command.

Procedure

1. Run **Scan for Devices**, in case a device path has been added to, or removed from, the library since the latest scan.
2. In the server's **Administration** window, click **Devices**.
3. Select **Libraries** in the navigation tree. The **Libraries detail** table appears.
4. In the navigation tree, right-click the entry for the library to be reconfigured, or open the **Storage Nodes** folder, open the library folder, and then right-click the library entry there.
5. Select **Reconfigure Library**. The **Reconfigure Library** window appears. Note that the storage node name and library name cannot be changed in this window.
6. Make appropriate changes in the **Configure devices on various storage nodes using existing drive connectivity** area, selecting or clearing checkboxes as

necessary, or using the buttons at the right side of the area (**Check All**, **Clear All**, **Reset**).

Drives that are already configured to be used by the library display check marks in the boxes that are adjacent to their names:

- Selecting a box adds the drive to the library.
 - Clearing a box removes the drive from the library.
 - The **Reset** button returns the checkboxes to the condition they had when the Reconfigure Library window was opened.
7. Click **Start Configuration** to reconfigure, or **Cancel** to leave the window.
 8. Run **Scan for Devices** to refresh the navigation tree and show the reconfiguration results.

Specifying library slots

The available slots feature controls which volumes the NetWorker server uses for backup. The server uses all volumes in a library to perform recoveries, but the volumes that are automatically selected for backups can be controlled by designating a range of available slots in the library.

About this task

Perform the following steps to define the available slots in a tape library.

Procedure

1. Ensure that volumes have been placed in all the available slots of the library so that the NetWorker server can continue uninterrupted with an automatic backup.

With two-sided media, the number of available slots is effectively doubled. For example, with 32 optical disks labeled “jupiter.001.a” to “jupiter.032.b,” there are a total of 64 sides, and therefore, there are 64 slots from which to choose.
2. In the server’s **NetWorker Administration** interface, select **View > Diagnostic Mode** from the menu bar.
3. Click **Devices**.
4. Open the **Libraries** folder in the navigation tree. The **Libraries** detail table appears.
5. In either the navigation tree or in the **Libraries** detail table, right-click the library on which the slots are to be designated, and select **Properties**.
6. Select the **Advanced** tab of the **Properties** window.
7. In the **Media Management Area**, in the **Available slots** field, type a range of contiguous slots, then click **+** to add the range of slots.

For example (assuming that no slots have already been configured), to designate slots 1 through 3 as available, then skip a defective slot 4, and designate slots 5 through 7 as available, type this information in the Available Slots field:

- a. Type **1-3**, then click **+** to add these slots.
- b. Type **5-7**, then click **+** to add these slots.
- c. Click **OK**. Slot 4 will be skipped when tapes are loaded.

Reset a library

A library must be reset each time the library and the NetWorker software become out of sync. A library reset can be done using either the Administration interface or the command prompt.

Resetting a library in the Administration interface

To reset a library in the Administration interface:

Procedure

1. In the **Administration** window, click **Devices**.
2. Open the **Libraries** folder in the navigation tree. The **Libraries** detail table appears.
3. Select a library in the navigation tree or double-click a library in the **Libraries** detail table to open the double-paned **Library Operations** view.

The library's drives are listed in the pane on the left in the **Device** column. The library's slots are listed in the pane on the right.
4. Right-click a library in the **Device** column, and select **Reset**. You are prompted to reset the library.
5. Click **Yes**. The **Library Operation** window appears and displays this message:

```
The library operation has started.
Please see the Monitoring->Operations screen for its status.
```

6. Click **OK**.

Resetting a library from the command prompt

Use the `nsrjb -HE` command to reset a library from the command prompt. For example, the library inventory must be correct after adding drives to an SJI-compliant library, such as adding DLT7000 drives to an ETL 7/3500 device.

About this task

To make the NetWorker software aware of these new drives, run `nsrjb -HE` to reset the library. The `-E` option reinitializes the library's element status. Some libraries can track whether there is media in a component in the library. This feature is known as an *element status* capability.

A series of commands exists that allow direct interaction with libraries (sji commands) and tape drives (cdi commands). These commands should only be used by the most knowledgeable of NetWorker users, as the consequences of using them can be unknown. For information about these commands, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

Deleting libraries

The library's devices remain, and can still respond to NetWorker operations (such as monitoring, labeling, deletion, and so on) after the library definition is deleted. A deletion of a library deletes the library, not its devices.

About this task

Procedure

1. In the server's **Administration interface**, click **Devices**.
2. Select **Libraries** in the navigation tree. The **Libraries detail** table appears.

3. In either the navigation tree or in the **Libraries detail** table, right-click the entry for the library to be deleted, and select **Delete**.
4. When prompted, click **Yes**.

This message appears:

```
"Are you sure you want to delete this jukebox? If so, please
re-attempt
deletion within a minute."
```

5. Click **OK** to confirm the deletion.

Library notifications

The NetWorker server uses notifications to send messages about NetWorker events. Several preconfigured notifications, such as the following, provide information about various situations:

- Volumes in the library are 90% full
- Library needs more volumes to continue
- Library has a mechanical problem
- Library device needs cleaning
- Cleaning cartridge needs attention.

The NetWorker software automatically mounts a required volume as long as the volume is loaded in the library. If a recovery operation requires a volume that is not loaded in the library, the Tape mount request 1 notification sends an alert to Monitoring > Alerts, with a request to do something with a specific volume.

After a library problem is corrected, it might be necessary to mount a volume so the NetWorker server can continue to back up or recover files.

Refreshing enterprise library views on request

Procedure

1. From the **Console** window, click **Libraries**.
2. In the navigation pane, select a server to update, or select the top item in the hierarchy to update library information for all NetWorker servers.
3. Right-click the server, and select **Refresh**.

Changing the polling interval for enterprise library views

Enterprise library views are updated periodically without user intervention.

About this task

Procedure

1. From the **Console** window, click **Setup**.
2. From the **Setup** menu, select **System Options**.
3. In the **Polling Interval for NetWorker Libraries** field, type the appropriate time, in hours.
4. Click **OK**.

Adding and removing media by using the library front panel

Certain media libraries allow for media to be added and removed by using the front panel display. This operation circumvents the NetWorker server's normal procedures

for adding and removing volumes and may cause the server information to become out of sync with the library. Normally, you should use the NetWorker server procedures for adding and removing media, rather than the library's front panel display. This is more efficient and guarantees that the server and the library will be in sync.

About this task

If it is necessary to use the library's front panel display to add and remove volumes.

Note: When a library is partitioned, the NetWorker software does not become aware of the partitioning. This means that the entire physical library will be disabled, not just one partition.

Procedure

1. In the **Properties** window for the Library, on the **General** tab, set **Status Enabled** to **Service**.

Note: Putting the library in service mode will cancel all operations or wait for operations to complete that cannot be canceled, and then put the library into disabled mode.
2. Once the library is in disabled mode, use the library's front panel to add and remove tapes.
3. In the **Properties** window for the Library, on the **General** tab, set **Status Enabled** to **Enabled**.
4. Inventory the library. [Inventorying library volumes](#) on page 160 has information about inventorying libraries.

Volume mounting and unmounting

A volume must be mounted before files can be backed up. If no volume is mounted at the start of a backup, an error message appears and requests that a volume be mounted.

Mounting or unmounting a volume in a library

Procedure

1. In the **Administration** window, click **Devices**.
2. Open the **Libraries** folder in the navigation tree. The **Libraries** detail table appears.
3. Select a library in the navigation tree or double-click a library in the **Libraries** detail table to open the double-paned library operations view. The library's drives are listed in the **Devices** column, and its slots are listed in the **Slot** column.
4. To mount a volume:
 - a. In the **Devices** column, select the appropriate drive.
 - b. In the **Volume** column, right-click a volume to mount, and select **Mount**.
 - The Library Operation window displays this message:

The library operation has started.

- The **Monitoring > Operations** screen displays its status.

c. Click **OK**.

5. To unmount the volume:

a. Right-click the device or the volume in the double-paned table view of the library and select **Unmount**.

- The Library Operation window displays this message:

The library operation has started.

- The **Monitoring > Operations** screen displays its status.

6. Click **OK**.

Unmounting volumes automatically (idle device timeout)

At times, a volume that is mounted in one device might be needed by another device in the same library. For example, data being recovered by one device could span more than one volume, and the required volume could be mounted on another device. To address this need, a value can be defined in the Idle Device Timeout attribute for that particular library.

About this task

The Idle Device Timeout attribute specifies the number of minutes a mounted volume can remain idle before it is automatically unmounted from the device and returned to its slot, where it can then be accessed by another device. For libraries, this attribute appears on the Timers tab of a library's Properties. The default value for a library is 10 minutes.

Procedure

1. In the server's **NetWorker Administration** interface, click **Devices**.
2. Open the **Libraries** folder in the navigation tree.
3. Right-click the appropriate library in the detail table, and select **Properties**. The **Properties** window appears.
4. Select the **Timers** tab.
5. Specify a value in the **Idle Device Timeout** attribute.
1. You can also override the library's Idle Device Timeout attribute for a specific device in the library.

To specify the Idle Device Timeout value for a specific device:

6. In the server's **Administration interface**, click **Devices**.
7. Select **View > Diagnostic Mode**.
8. Select **Devices** in the navigation tree. The Devices detail table appears.
9. Right-click the device and select **Properties**.
10. Select the **Advanced** tab.
11. Specify a value in the **Idle Device Timeout** attribute.

The default value is 0 (zero) minutes, which means that the device never times out and the tape must be ejected manually. However, when the value of this attribute is set to 0, the value specified in the device library's Idle Device Timeout attribute will take precedence.

Mounting or unmounting a volume in a stand-alone tape drive

Procedure

1. Manually insert a volume in the stand-alone drive, or ensure that a volume is already loaded.

In a stand-alone device, a volume that has been loaded into the drive is not considered to be mounted until it has been explicitly mounted in the user interface or from the command prompt.
2. In the **Administration** window, click **Devices**.
3. Select **Devices** in the navigation tree. The **Devices detail** table appears.
4. Select the device. To mount the volume, in the **Devices detail** table, right-click the device, and select **Mount**.
5. To unmount the volume, in the **Devices > detail** table, right-click the device, and select **Unmount**.

- The **Library Operation** window displays this message:

```
The library operation has started.
```

- The **Monitoring > Operations** screen displays its status.
6. Click **OK**.

Labeling and mounting a volume in one operation (stand-alone tape drive)

When multiple storage devices are connected to the NetWorker server, the device for labeling must first be selected from the list of available devices. Remember that labeling a volume makes it impossible for the NetWorker server to recover original data from that volume.

Procedure

1. In the **Administration** window, click **Devices**.
2. Manually insert an unlabeled or recyclable volume in the NetWorker server storage device, or ensure that a volume of this type is already present for the NetWorker server to access.
3. Select **Devices** in the navigation tree. The **Devices detail** table appears.
4. Right-click the stand-alone device in the detail table, and select **Label**. The **Label** window appears:
 - a. Type a unique label name, or accept the default name that is associated with the selected pool.

If the volume is unlabeled, the NetWorker server assigns the next sequential label from the label template that is associated with the selected pool. If a recyclable volume from the same pool is being re-labeled, then the volume label name and sequence number remain the same. Access to the original data on the volume is destroyed, and the volume becomes available.
 - b. Select a pool on the **Pools** menu. The NetWorker server automatically applies the label template that is associated with the **Default** pool unless a different pool is selected.

- c. Select the **Manual Recycle** attribute if the volume should be manually recycled.

If the Manual Recycle attribute is enabled when the volume is labeled, the volume cannot automatically be marked as recyclable according to the retention policy. When a volume is marked as manual recycle, the NetWorker server disregards the assigned browse and retention policies. Therefore, only an administrator can mark the volume recyclable.

A volume that has been set to manual recycle retains that setting, even after re-labeling. A Manual Recycle policy cannot be changed back to Auto Recycle by clearing the Manual Recycle checkbox. The volume must be explicitly reset to use auto recycle.

- d. The **Mount After Labeling** attribute is selected by default. The NetWorker server automatically labels the volume, and then mounts the volume into the device.

5. Click **OK**.
6. If the volume is recyclable, a message warns that the named volume is about to be recycled, and asks whether to continue. Click **Yes** to re-label and recycle the volume.
7. After a volume is labeled and mounted in a device, the volume is available to receive data. Since the NetWorker label is internal and machine-readable, place an adhesive label on each volume that matches that internal volume label.

[Configuring a library to use volumes with barcodes](#) on page 144 provides information on using barcode labels.

 **Note:** If you are in the process of re-labeling a mounted volume and you choose not to overwrite the existing label, the volume is left in an unmounted state. To use this volume, mount it again.

Labeling volumes without mounting

Volumes can be prelabeled without being mounted.

To label a volume without mounting, follow the same procedures as for labeling and mounting in one operation, but clear the **Mount After Labeling** attribute in the **Label** window.

Mounting uninventoried volumes

You can mount volumes that are not included in the library inventory, but are valid (properly labelled) NetWorker volumes.

Procedure

1. In the **Administration** window, click **Devices**.
2. Select **View > Diagnostic Mode** on the toolbar.
3. Manually insert the volume in an empty library slot.
4. Open the **Libraries** folder in the navigation tree. The **Libraries** detail table appears.
5. Select the library in the navigation tree in which the volume was manually inserted, or double-click the same library in the **Libraries** detail table. The **Libraries** detail table changes to the double-paned library operations view. The library's drives are listed in the **Devices** column, and its slots are listed in the **Slot** column.

6. In the **Devices** column, right-click the library in which the volume was manually inserted, and select **Inventory**. The **Inventory Library** window appears.
7. Type the slot number of the volume in both the **First** and **Last** field of the **Slot Range**.
8. Select **Operation Type**: either **Slow/Verbose** (the default) or **Fast/Silent**.
 - When **Slow/Verbose** is selected, the **Supply Input** option and icon on the **Operations** screen of the **Monitoring** window can be used to confirm the choice to relabel a volume. The device path appears in the **Device** field.
 - When **Fast/Silent** is selected, the **Supply Input** option and icon are not available, and relabeling proceeds automatically, without user input. The device path does not appear in the **Device** field. [Entering user input](#) on page 61 provides details.
9. Click **OK**.
 - The **Library Operation** window displays this message:
The library operation has started.
 - The **Monitoring > Operations** screen displays its status.
The NetWorker software then inventories the specified slot.
10. Mount the inventoried volume.

 **NOTICE** Unlabeled tapes may not be mounted for inventoring. Unlabeled tapes can only be mounted to be labeled. An attempt to mount an uninventoried volume by using unlabeled media results in an I/O error. The volume will also be ejected.

Libraries with volume import and export capability

The NetWorker software supports the use of the SCSI-II import/export feature found in many brands of library. Depending on the library model, this feature is also known as cartridge access port (CAP), mail slot, and loading port. The import/export feature deposits and withdraws (ejects) volumes from slots in the library. This feature enables the operator to deposit and withdraw cartridges without invalidating the device inventory list. Normally, if the operator opens the door to load or unload media, the element status of the autoloader is invalidated, which requires the reinitialization the library. The NetWorker server does not, however, automatically inventory the volume after a deposit and withdrawal.

The reinitialization usually consists of the following:

- An inventory of all slots
- A reset of the robotic arm
- A check to see whether each drive is working

The Deposit attribute causes a library to take the first available volume from the CAP and place it in the first empty library slot. The Eject/Withdraw attribute moves a volume from a slot (never from a drive) to the CAP.

Depositing a volume by using the import/export feature

Use these general instructions when working with a CAP. Specific instructions for working with a CAP can vary, depending on the library manufacturer. For specific instructions, refer to the library's documentation.

Procedure

1. Ensure that volumes are available in the CAP for deposit.

2. In the **Administration** window, click **Devices**.
3. Select **Libraries** in the navigation tree.
The **Libraries detail** table appears.
4. Double-click the library in which to deposit the volume.
The **Libraries detail** table changes to the double-paned library operations view.
5. Right-click either the device or the slot, and select **Deposit**.
You are prompted to deposit the volume.
6. Click **Yes**. The **Library Operation** window displays this message:

The library operation has started.

The **Monitoring > Operations** screen displays its status.

7. Click **OK**.
8. Click **Monitoring** to go to the **Monitoring** window, and then select the **Operations** tab.
9. Right-click the **User Input** icon for the deposit job and select **Supply Input**.
You are prompted to load the cartridges into the ports and type **Yes** to continue.
10. Click **Yes**.
11. Right-click the **User Input** icon for the deposit job and select **Supply Input** again.
You are prompted to continue depositing volumes.
12. Click **Yes** to continue depositing volumes, or **No** when done.

Withdrawing a volume by using the import/export feature

About this task

Note: If the library is partitioned into logical libraries and the import/export slots are shared between the partitions, you must withdraw volumes by using the `nsrjb -P` command to specify the ports from which to withdraw the volumes. The `nsrjb` man page or *NetWorker Command Reference Guide* for more information.

Procedure

1. Ensure that the volume to be withdrawn is in a known slot, and that the CAP has an empty port to hold the withdrawn volume.
2. In the **Administration** window, click **Devices**.
3. Select **Libraries** in the navigation tree. The **Libraries detail** table appears.
4. Double-click the library from which the volume is to be *NetWorker Command Reference Guide* withdrawn. The **Libraries detail** table changes to the double-paned library operations view.
5. Right-click the slot that contains the volume, and select **Eject/Withdraw**.
You are prompted to withdraw the volume.
6. Click **Yes**.
 - The **Library Operation** window displays this message:

The library operation has started.

- The **Monitoring > Operations** screen displays the status.
7. Click **OK**.
 8. To review the result, select **Monitoring > Log**. A successful **Eject/Withdraw** operation ends with a **Succeeded** comment in the log.

Inventorying library volumes

When the NetWorker software labels the contents of a library, the software registers the location of the volumes in the library slots when it assigns the volume label. This process is called taking inventory. When the volumes in the library are inventoried, the NetWorker software reads the label of each volume and records its slot number. If the volumes are not moved in the library after they have been labeled, then the NetWorker server can access the volumes because each volume label is assigned to a specific slot.

About this task

If, however, the contents of the library are changed without being labeled, or if volumes are moved into new slots, the NetWorker software must be notified that the library now holds a different set of labeled volumes or that the volumes are in a different order. For example, if the library has more than one magazine, the volumes must be inventoried each time that a magazine is removed, and another one is loaded into the library.

When the volumes in a new magazine are labeled, there is no need to inventory them. The NetWorker software automatically records the slot number in which each newly labeled volume is located.

The NetWorker software can use barcode labels to speed up the inventory process. If the library supports the use of barcode labels, consider using them if large numbers of volumes, and/or if the library contents change often. [Barcode labels](#) on page 149 provides more information on using barcode labels.

Procedure

1. In the **Administration** window, click **Devices**.
2. Open the **Libraries** folder in the navigation tree. The **Libraries** detail table appears.
3. Select a library in the navigation tree or double-click a library in the Libraries detail table. The **Libraries** detail table changes to the double-paned library operations view.
4. Right-click anywhere within the **Devices** pane, and select **Inventory**. The **Inventory > Library** window appears.
5. Type the numbers of the first and last slots to be inventoried in the **Slot Range** area.
6. Select **Operation Type**: either **Slow/Verbose** (the default) or **Fast/Silent**.
7. Click **OK**.
 - The **Library Operation** window displays this message:

```
The library operation has started.
```

- The **Monitoring > Operations** screen displays its status.
8. Click **OK**. If the volumes do not have barcode labels, the NetWorker software must mount each volume, read its label, and unmount it. In this case, the inventory process can take some time to complete.

Library maintenance

Periodically clean a storage library to keep it working correctly. The NetWorker server provides automatic cleaning of devices located in libraries. The server does not support automatic cleaning for stand-alone devices. Cleaning is an option set during configuration.

The service mode feature allows a library to be taken offline temporarily for cleaning or other maintenance.

Automatic tape device cleaning

Tape device cleaning is an automated, self-contained operation. It is no longer part of a media-loading operation. Tape device cleaning is automatically triggered if one of these conditions exist:

- The last time the device was cleaned was a full cleaning interval ago.
- The Cleaning Required attribute for the device is set to Yes in one of the following ways:
 - Manually by the user.
 - Automatically by the NetWorker server, after it receives a “device needs cleaning” notification.

When one of these conditions is met for a device, cleaning begins as soon as the device becomes available. Loaded devices are unloaded before a cleaning operation begins. Loading a cleaning cartridge (with the `nsrjb -l` cleaning cartridge command) to force a cleaning operation is no longer supported.

Selecting a tape device manually for cleaning

About this task

 **NOTICE** Do not enable automated cleaning for silos in the NetWorker software. The automated device cleaning feature cannot be used in a silo, because it depends on fixed slot numbers. For information about how to clean devices in a silo, refer to the silo manufacturer’s software documentation.

Procedure

1. In the server’s **NetWorker Administration** interface, click **Devices**.
2. Open the **Libraries** folder in the navigation tree and select the drive that contains the mounted volume with the block size being checked. The drive’s detail table appears.
3. Right-click the drive in the detail table, and select **Properties**. The **Properties** window appears.
4. Select the **General** tab.
5. Set the **Cleaning Required** attribute to **Yes**.

Delaying tape device cleaning

Occasionally it is necessary to set the **Cleaning Delay** attribute in order to allow a tape device to sleep before attempting to unload a cleaning cartridge.

Procedure

1. In the server’s **NetWorker Administration** interface, click **Devices**.
2. Select **View > Diagnostic Mode**.

3. Open the **Libraries** folder in the navigation tree.
4. Right-click the appropriate library in the detail table, and select **Properties**. The **Properties** window appears.
5. Select the **Timers** tab.
6. Select a value in seconds for the **Cleaning Delay** attribute.

Tape alert

The TapeAlert feature provides, among other things, diagnostic information for devices for which hardware cleaning is enabled.

NetWorker provides the following attributes for tape device cleaning:

- Cleaning required
- Cleaning interval
- Date last cleaned

When the Common Device Interface (CDI) is enabled, TapeAlert attributes provide tape drive status. SCSI Commands must be selected for the CDI attribute on the Configuration tab of the relevant device's Properties. If CDI cannot be enabled, TapeAlert is not supported.

Devices that are capable of TapeAlert perform constant self-diagnostics and communicate the diagnostic information via the nsrmmd program to logs that can be viewed in the Monitoring task.

The following TapeAlert attributes are found in the device's Properties, on the Volume tab.

- TapeAlert Critical: Displays critical diagnostic information, such as for media or drive failure, when user intervention is urgent and data is at risk.
- TapeAlert Warning: Displays a message when the media or device needs servicing.
- TapeAlert Information: Displays status information.

The following table describes the nature of the tape alert levels.

Table 31 Tape alert severity

Severity	Urgently requires user intervention	Risks data loss	Explanatory
Critical	X	X	
Warning		X	X
Informative			X

The messages indicate tape and drive states related to tape drive read/write management, cleaning management, or drive hardware errors.

Informative messages

Informative messages indicate status information:

- A data or cleaning tape is nearing its end of life.
- A tape format that is not supported.

 **Note:** When automatic cleaning is enabled, a diagnostic message to indicate that a drive needs cleaning initiates NetWorker drive cleaning.

Warning messages

Warning messages indicate the following types of drive errors:

- Recoverable read or write errors occurred.
- Media is at end of life.
- Read-only tape format is in the drive.
- Periodic cleaning is required.

Critical messages

Critical messages are warnings that a drive might be disabled and requires immediate attention to avoid data loss:

- Unrecoverable read or write errors occurred.
- Tape is marked read-only.
- Drive require immediate cleaning.
- Drive is predicting hardware failure.

Informative and warning messages should clear automatically by nsrmmd once the reported issue is handled.

Critical messages about hardware errors are not cleared by nsrmmd because they might indicate intermittent hardware problems.

Troubleshooting libraries and devices

This section provides detailed information about how to troubleshoot issues with libraries and devices, including how to correct drive ordering issues and block size issues between UNIX and Windows devices.

Troubleshooting autoconfiguration failure

Common symptoms of library autoconfiguration failure include the following:

About this task

- The library is not listed in the Libraries folder in the Administration interface.
- The library is listed, but is listed as being unconfigured.

Common causes include:

- Device drivers are not properly installed.
- Autodetection fails to match a detected library with its devices due to:
 - Out-of-date device firmware.
 - Failure of the library to return its devices' serial numbers.
- Autodetection failed to start on the storage nodes.

Procedure

1. Check **Monitoring > Log** for relevant messages.
2. From the command prompt, type the following command to verify that the library returns the serial numbers of its devices:

```
sn -a b.t.l.
```

where b.t.l. refers to the bus target LUN of the library. If the bus target LUN is not known, run the inquire command first, to obtain this information.

Library configuration using the jbedit command

If the autoconfiguration program cannot be used, the `jbedit` (jukebox edit) program can be used as a fallback means of editing library configurations. This command can be run on a NetWorker server, storage node, or client (if the client is a storage node). It operates without disrupting any backup or recovery operations on the library.

Running the `jbedit` program requires Configure NetWorker user privileges.

The `jbedit` program supports all direct-attached SCSI/SJI, SAN, and NDMP libraries.

The `jbedit` program is not intended to be a full-fledged editor of the Library resource. The editing of Library resource attributes should be done as described in [Reconfiguring a library](#) on page 150. The `jbedit` options provide selection lists that make it easy to find drives or devices to be added or deleted.

The following table lists the most commonly used `jbedit` program options.

Table 32 Common `jbedit` options

Option	Description
-a	Add a drive or device.
-d	Deletes a drive or device.
-j	Name of the autochanger to be edited.
-f	Name of the device to be added or deleted.
-E	Element address of the device to be added or deleted.

The *NetWorker Command Reference Guide* or the UNIX man page provides a detailed description of the `jbedit` command, its options, and associated diagnostic messages.

Device ordering

The NetWorker server uses logical device names assigned by the operating system when communicating with devices. It is possible for the operating system to re-associate logical device names with the physical addresses of the devices, generally after rebooting the host or after plug-and-play events. This may cause device reordering, where the physical device will have a different device filename. As a result, tape devices configured in the NetWorker software no longer match the names of the devices as recognized by the operating system.

If device reordering occurs, the NetWorker software is unable to use any affected drives until the configuration is manually corrected.

The NetWorker server detects device reordering events by comparing the current serial number of the device to the serial number of the device at configuration. If the serial numbers do not match, the NetWorker server stops all operations on that device and an error message will be posted, similar to the alert identified for device serial number mismatch in the table [Preconfigured notifications](#) on page 717. CDI must be enabled for this functionality. [Setting the common device interface](#) on page 172 provides more information about enabling CDI.

Detecting device ordering issues

To determine if there is a problem with device ordering in your environment, you first determine if the device order that appears in `nsrjb` output matches the device order

from the `inquire` and `sjisn` commands, then verify that the device configuration within your NetWorker configuration conforms to this.

Procedure

1. Execute the `inquire` command with the `-cl` option to determine the device path, scsi address, and serial number of the device.
2. Execute the `sjisn` command to determine the current order of the devices:

```
sjisn scsidev@bus.target.lun
```

where `bus.target.lun` is the SCSI address of the robotic arm returned by the `inquire` command in step 1, for example, 1.2.0.

3. Match the serial numbers of the devices in the `sjisn` output to the device names that correspond to these serial numbers in the `inquire -cl` output. This will give you the current device order by device filename.
4. Execute the `nsrjb` command to determine the order of devices as configured in NetWorker. Drive entries towards the end of the `nsrjb` output list the device order as configured in NetWorker.
5. Compare the device ordering as determined in step 3 and step 4. If the device ordering in these two steps do not match, the device ordering has changed and the library will need to be reconfigured.

Drive ordering change corrections

After a drive ordering change has taken place and the NetWorker software is no longer correctly communicating with devices, you can correct the problem within your NetWorker configuration by using the NetWorker Console or the `jbedit` command line program.

Using NetWorker Console to correct drive ordering changes

You can correct drive ordering changes by using the NetWorker Console.

Procedure

1. Ensure that you have a current backup of the resource database.
2. Delete the library resource in the NetWorker Console. [Deleting libraries](#) on page 152 provides details.
3. Rescan the library. [Scanning for libraries and devices](#) on page 142 provides more information.

Using the `jbedit` command to correct drive ordering changes

You can correct drive ordering changes by using the `jbedit` command.

Procedure

1. Use the `jbedit` command with the `-d` option to delete devices from the NetWorker configuration.
2. Use the `jbedit` command with the `-a` option to add the devices again.

[Library configuration using the `jbedit` command](#) on page 164, or the UNIX man page for `jbedit` or the *NetWorker Command Reference Guide* provides more information about the `jbedit` command.

Clearing device ordering/serial mismatch errors from the NetWorker Console

After a device ordering error has been detected, a message is displayed in the Alerts and Notifications windows of the NetWorker Management Console, as well as the log files. The error message is similar to the following:

About this task

```
"Check system device ordering. Moving device on %s to . To correct, scan for devices in NMC and re-enable the device."
```

An Event ID for the error is also created, which will be removed along with the alert when the problem is resolved. You can resolve the problem and clear the error message.

Procedure

1. Disable the drive.
2. Perform one of the above procedures to correct the problem.
3. Re-enable the drive, and retry the operation that was being performed prior to receiving the error.

Results

The Alert will be removed and the event dismissed.

Tape drive number reordering (Microsoft Windows only)

If more than one tape drive is attached to the NetWorker server when both the server and drives are shut down, restart all of the tape drives, either before or immediately after the NetWorker server is restarted. If Windows does not locate all of its previously configured tape drives at the time of startup, it automatically reassigns the tape registry name.

For example, assume that these three tape drives are attached to the server:

- The first one, \\.\Tape0, is a 4 mm tape drive.
- The second, \\.\Tape1, is an 8 mm tape drive.
- The third, \\.\Tape2, is also an 8 mm tape drive.

If only the second and third tape drives are restarted, Windows reassigns the tape registry numbers so that the second storage device becomes \\.\Tape0 and the third storage device becomes \\.\Tape1. The tape registry numbers no longer match the defined storage devices within the NetWorker software. As a result, the server mishandles the drives and their volumes.

It might be easier to leave a nonoperational drive (device) attached to the server until a replacement is available. If the drive is removed, the name must be deleted, and then the new drive must be added.

To disable the drive, select No for the Enabled attribute in the device's Properties.

Device calibration

For information about the frequency and method for calibrating the loading mechanism for the device, refer to the library manufacturer's documentation.

SCSI data block size issues between UNIX and Windows

Different SCSI hardware limitations exist between UNIX and Microsoft Windows operating systems. This can lead to data block size compatibility problems (although they are less likely to occur now than in the past, given larger Fibre-Channel capacities). For example, with a device defined in UNIX that is physically attached to a Windows HBA, it is possible to define a block size greater than that allowed by the

Windows hardware. This could lead to I/O errors in both write and read states on the device. In order to use both operating systems, it is necessary to determine a block size that is acceptable to both.

NOTICE In NetWorker 8.0.1 and later, the default block size for an LTO device increases from 128 KB to 256 KB. When NetWorker labels a new or used volume in an LTO device and the Device block size attribute of the device is handler default, the label operation uses a 256 KB block size.

Determining the allowable block size

You can determine the allowable block size by checking the **Properties** window of a mounted volume while in Diagnostic Mode.

Procedure

1. In the server's **NetWorker Administration interface**, click **Devices**.
2. Select **View > Diagnostic Mode**.
3. Open the **Libraries** folder in the navigation tree and select the drive that contains the mounted volume with the block size being checked. The drive's detail table appears.
4. Right-click the drive in the detail table, and select **Properties**. The **Properties** window appears.
5. Select the **Volume** tab. In the **Loaded Volume** area, one of the displayed volume attributes is the **Volume Block Size**.
6. Click **OK**.

Solving block-size compatibility problems

About this task

Note: It is also possible to solve problems with block-size compatibility by changing the block size for an entire device type. The change, however, must be made on each storage node where it is to be available. Once the block size is changed, it affects only those volumes that are labeled after the change. Volumes can be relabeled to use the new block size, but if they contain data that should be saved, be sure to clone the data beforehand to a volume that already uses the new block size.

Procedure

1. In the server's **NetWorker Administration interface**, click **Devices**.
2. Select **View > Diagnostic Mode** on the menu bar.
3. Open the **Libraries** folder in the navigation tree and select the drive that contains the mounted volume with the block size being checked. The drive's detail table appears.
4. Right-click the drive in the detail table, and select **Properties**. The **Properties** window appears.
5. Select the **Advanced** tab. In the **Device Configuration** area, the currently configured **Device Block Size** value is displayed.
6. Select the appropriate **Device Block Size** value.
7. Click **OK**.

Setting the block size for a device type

Procedure

1. Change the block size:

- On UNIX, change the block size by setting this environment variable to the greatest common value for both systems. For example:

```
setenv NSR_DEV_BLOCK_SIZE_MEDIA_TYPE value
```

where:

- *MEDIA_TYPE* is the backup device type available to the NetWorker server (also found in the Media Type attribute on the General tab of the device's properties). The media type syntax must be all uppercase, with underscores (_) replacing blank spaces and hyphens. Therefore, a device displayed in the NetWorker software as "8mm Mammoth-2" would be listed as:
8MM_MAMMOTH_2
 - *value* must be a multiple of 32 KB, with a minimum value of 32 KB.
 - On Microsoft Windows only, install a later model HBA, or upgrade to drivers that can support up to 128 KB blocks. Windows also accepts the same environment variable format as UNIX to set block size.
2. Restart the NetWorker server in order for changed environment variables to take effect.

Device block size for read and write operations

The block size for a volume is defined during the label operation. The label operation uses the value defined in the Device block size attribute for the Device or the value defined by the appropriate block size environment variable.

The block size for both read and write operations uses the block size defined in the volume header during the label operation rather than the device block size.

Block-size mode (UNIX/Linux only)

Ensure that the block size mode for tape devices that are used with NetWorker software is set to variable. Otherwise, data recovery might fail. The procedure for setting the device block size varies depending on the operating system.

The operating system's documentation provides information about setting the tape device block size in the operating system.

Device parameter settings

Device parameter settings can be modified for the devices the NetWorker software uses in two ways:

- Individually, through the NetWorker Administration interface.
- Globally, for all devices through operating system environment variables. The adjustment of environment variables should only be done by users who know the server environment and performance tuning requirements. For example, an administrator who wants to fine-tune performance by changing a certain setting for all LTO devices on a particular NetWorker server.

The variables (and their equivalent names in the Administration interface) are described in the following sections.

Device setting environment variables

There are several device-related environment variables available to configure devices for the NetWorker software.

Device-related environment variables include the following:

- NSR_DEV_BLOCK_SIZE_MEDIA_TYPE

- NSR_DEV_TAPE_FILE_SIZE_MEDIA_TYPE
- NSR_DEV_LOAD_TIME_MEDIA_TYPE
- NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE
- NSR_DEV_LOAD_TRY_LIMIT_MEDIA_TYPE
- NSR_DEV_DEFAULT_CAPACITY_MEDIA_TYPE

where:

MEDIA_TYPE is the backup device type available to the NetWorker server.

Note: The media type syntax must be all uppercase, with underscores (_) replacing blank spaces and hyphens. For example, a device displayed in the NetWorker software as “8mm Mammoth-2” would be listed as: 8MM_MAMMOTH_2
To determine the media type, right-click the device and select the **General** tab. The **Media Type** attribute contains the media type that should be used in these environment variables.

NSR_DEV_BLOCK_SIZE_MEDIA_TYPE

NSR_DEV_BLOCK_SIZE_MEDIA_TYPE is organized in units of kilobytes. This environment variable will cause NetWorker to override the default block-size setting defined for the tape drive in the operating system. The value set must be a multiple of 32, with a minimum value of 32. Maximums are determined by platform, SCSI driver, and device.

For example:

NSR_DEV_BLOCK_SIZE_4MM_20GB=64

For information about using this environment variable to set block-size compatibility between UNIX and Microsoft Windows, [SCSI data block size issues between UNIX and Windows](#) on page 166 provides more information.

NSR_DEV_TAPE_FILE_SIZE_MEDIA_TYPE

NSR_DEV_TAPE_FILE_SIZE_MEDIA_TYPE is organized in units of NSR_DEV_BLOCK_SIZE_MEDIA_TYPE and is the number of blocks written between filemarks. These filemarks are used to locate a particular spot on the tape during recovery, and more filemarks generally lead to faster positioning. For example:

NSR_DEV_TAPE_FILE_SIZE_TZ89=512

On UNIX and Linux platforms, the NetWorker software writes a filemark by closing and reopening the tape device, which takes one or two seconds. If this value is too small, throughput could be slowed and recoveries may take longer to complete.

On Microsoft Windows platforms, the NetWorker software writes asynchronous filemarks. This setting has a minimal effect on performance.

NSR_DEV_LOAD_TIME_MEDIA_TYPE

NSR_DEV_LOAD_TIME_MEDIA_TYPE is the number of seconds that nsrmmd polls and waits for a drive to become ready after the library inserts a tape into the device. NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE is used to set the number of seconds nsrmmd waits between polls during load time.

If the value of NSR_DEV_LOAD_TIME_MEDIA_TYPE is too short, there could be unnecessary load failures. If it is too long, then labeling new tapes takes longer than necessary. The minimum allowable value is 10 seconds. The maximum value is 600 seconds. For example:

NSR_DEV_LOAD_TIME_DTL8000=300

NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE

NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE is the number of seconds that nsrmmd waits between each attempt to read a newly inserted tape. The minimum allowable value is 1 second, the maximum value is 30 seconds. For example:

NSR_DEV_LOAD_POLL_INTERVAL_DLT=10

NSR_DEV_LOAD_TRY_LIMIT_MEDIA_TYPE

NSR_DEV_LOAD_TRY_LIMIT_MEDIA_TYPE is the number of times that nsrmmd will attempt to open a drive. The nsrmmd program will poll the drive until the limit set in NSR_DEV_LOAD_TIME_MEDIA_TYPE is reached. After the limit is reached, it will retry until the NSR_DEV_LOAD_TRY_LIMIT_MEDIA_TYPE is reached. The default value and minimum allowable value is 2, the maximum value is 120.

NSR_DEV_LOAD_TRY_LIMIT_DLT=4

NSR_DEV_DEFAULT_CAPACITY_MEDIA_TYPE

NSR_DEV_DEFAULT_CAPACITY_MEDIA_TYPE is the size of the particular tape used to base the percent full calculation. This variable value has no effect on the actual tape capacity. Any integer value is allowed, with a KB, MB or GB designation to indicate a range of values. Any value less than 200 MB will be overridden by the normal default capacity. There is no obvious maximum, with the only practical limitation being the actual storage size. For example:

NSR_DEV_DEFAULT_CAPACITY_DTL7000=12GB

Setting device parameters in the NetWorker Administration interface

You can locate and change the device parameters in the Administration interface.

About this task**Procedure**

1. In the server's **Administration interface**, click **Devices**.
2. Select **View > Diagnostic Mode**.
3. Select **Devices** in the navigation tree. The Devices detail table appears.
4. Double-click the device in the devices table or right-click the device and select **Properties**. The **Properties** window appears, with the **General** tab selected.
5. Select the **Advanced** tab. In the **Device Configuration** area, the device settings are the first fields shown. The following table lists the fields and their corresponding environment variables:

Results

Table 33 Device settings and environment variables

Device setting	Corresponding environment variable
Device Block Size	NSR_DEV_BLOCK_SIZE_MEDIA_TYPE
Device File Size	NSR_DEV_TAPE_FILE_SIZE_MEDIA_TYPE
Device Load Time	NSR_DEV_LOAD_TIME_MEDIA_TYPE
Device Eject Time	None
Device Poll Interval	NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE
Device Min Load Tries	NSR_DEV_LOAD_TRY_LIMIT_MEDIA_TYPE

Table 33 Device settings and environment variables (continued)

Device setting	Corresponding environment variable
Device Default Capacity	NSR_DEV_DEFAULT_CAPACITY_MEDIA_TYPE
Device Tape Flags	None

When device parameters are set in this interface, it is not necessary to stop and restart the NetWorker server in order for the settings to take effect.

Setting device environment variables on Windows

Setting environment variables for the NetWorker software differs on Windows and UNIX operating systems.

About this task

Environment variables on Microsoft Windows are set using the Control Panel System applet on the NetWorker server.

Procedure

1. Browse to **Control Panel > System and Security > System > Advanced System Settings**.
2. In the **General** tab click **Environment Variables...**
3. Click **New**.
4. Specify the environment variable name and value.
5. Stop and start the NetWorker Backup and Recover Server service in order for the environment variables to take effect.

Setting device environment variables on UNIX

Setting environment variables for the NetWorker software differs on Windows and UNIX operating systems.

About this task

On UNIX and Linux NetWorker sources the `/nsr/nsrrc` file before starting the NetWorker processes.

Procedure

1. On the NetWorker server, modify the `/nsr/nsrrc` file. If this file does not exist, create this file as a Bourne shell script file.
2. Add the environment variables in the following format:

```
ENV_VAR_NAME = value
export ENV_VAR_NAME
```

3. Stop and start the NetWorker server processes in order for the environment variables to take effect.

Setting the common device interface

Common device interface (CDI) allows the NetWorker server to send commands to tape devices. The CDI feature is not supported within an NDMP environment. You can set CDI support in the NetWorker Administration interface.

Procedure

1. In the server's **NetWorker Administration interface**, click **Devices**.
2. Select **View > Diagnostic Mode**.
3. Select **Devices** in the navigation tree. The Devices detail table appears.
4. Double-click a device in the **Devices** table (or right-click the device and select **Properties**). The **Properties** window appears, with the **General** tab selected.
5. Select the **Advanced** tab. In the **Device Configuration** area, locate the CDI settings:

- **Not Used:** Disables the CDI feature and uses standard tape driver calls for tape operations.
- **SCSI Commands:** Sends explicit SCSI commands to tape devices.

When enabled, the CDI feature:

- Provides clearer tape status messages.
- Informs when a tape is write protected.
- Enables Tape Alert, which provides diagnostic information for devices.

Although you can disable the CDI feature by selecting the **Not Used** option, it can be time-consuming to disable CDI on a large number of devices.

In this situation, create an empty file named `/nsr/debug` directory and create an empty file that is named `cdidisable`. Then restart the NetWorker server. The presence of this file disables the use of CDI for that server and all the storage nodes that are controlled by that server.

Note: Use of CDI does not change what is written to tape. A tape that is written with CDI enabled can be read with CDI disabled. Conversely, a tape that is written with CDI disabled can be read with CDI enabled. The CDI feature enables NetWorker software to collect better diagnostic information and facilitates tape usage when enabled. Only set or disable the CDI feature on the advice of an Customer Service representative. If tape or SCSI issues occur while the CDI feature is enabled, contact Customer Service.

Media handling errors

The architecture of device drivers can produce media handling errors. The NetWorker software automatically retries a failed operation such as a mount or read of a volume. The number of times the NetWorker software retries the failed operation depends on the value of the Max Consecutive Errors attribute, which is set in the Advanced tab of the device's Properties window. The default value is 20. When the device's Max Consecutive Errors value is reached, the device stops retrying the operation and becomes disabled.

A mount or read operation might fail for several reasons, for example:

- Attempts to mount and read a damaged tape in a library can result in a loop of failed actions: the device might repeatedly try to mount the tape, replace it in the slot, and then retry the action with the same result. In this example, to bring the drive back into use, remove the damaged tape, then reenable the device.

- A drive that always reports a fixed number of failures before correctly mounting and reading a tape, even if the tape is not damaged, can cause a failure loop. In this example, ensure that the Max Consecutive Errors value is higher than the number of times that particular drive fails before working correctly.

Re-enabling a device

Once the number of retries equals the Max Consecutive Errors value, the device becomes disabled. After the problem that disabled the device has been fixed, the device (drive) must be reenabled before it can be used again.

Procedure

1. When the NetWorker computer is idle, remove any volume from the disabled drive and ensure that the drive is in good working order.
2. In the **Administration** window, click **Devices**. The **Devices** detail table appears.
3. Right-click the drive to be reenabled, and select **Properties**. The **Properties** window appears.
4. In the **Status** area of the **General** tab, set **Enabled** to **Yes**.
5. Click **OK**.
1. If the disabled drive is part of a library, it might be necessary to reset the device. To do this:
6. From the command prompt, change the path to the directory that contains the NetWorker binaries.
7. Type this command:

```
nsrjb -HE
```

 **NOTICE** A device retains its enabled or disabled status in the Properties window and in the Devices detail table regardless of whether its storage node is enabled or disabled. Therefore, it is possible that the storage node Properties window is set to disabled while its devices appear to be enabled in the GUI.

Silo libraries

This section describes silos and silo devices. Silos and libraries are managed similarly by NetWorker software.

A silo tape library (STL) is a peripheral that usually contains many storage devices.

Silo libraries have a robotic controller that moves tape media between slots and devices. Silos do not use a SCSI interface to access and control the media movements. Media movements are controlled by a separate host that is called the silo server. The silo server uses silo management software to manage media movement requests over the network. The silo vendor provides the silo management software. The silo server cannot be the same computer as the NetWorker server.

The silo can be shared among many applications, systems, and platforms. As with libraries, silos make data and media operations more automatic. Silos can load, change, and manage volumes, and clean the devices automatically.

NetWorker only supports silos that use the Automated Cartridge System Library Software (ACSL) Manager software.

NetWorker software interactions with a silo

A NetWorker server acts as a client of the silo management software, which resides on the silo server. The NetWorker server communicates with the silo through the Silo

Tape Library Interface (STLI), which must be installed on the NetWorker server that uses the silo.

To access the volumes and devices in a silo, the NetWorker server sends a request to the silo management software, in the form of an STLI call. For example, to mount a volume in a silo device, the NetWorker media service sends a request to the silo management software to mount the volume into a particular device in the silo. The silo server responds to the request and mounts the volume in the requested device.

The silo management software controls many of the operations that NetWorker software controls with a library. For example, the silo management software keeps track of the slot where each silo volume resides, and might control the deposit and withdrawal of volumes, as well as automated cleaning of silo devices.

Naming conventions for silo devices

The silo name of the storage devices is supplied during the configuration process. The silo name is the name that the silo management software uses to refer to the storage device. Depending on the type of silo, the device name can take several forms. This section describes the naming conventions of the currently supported silos.

StorageTek device naming conventions

The StorageTek (STK) silo management software uses either a program that is called ACSLS that runs on a UNIX system, or a program that is called Library Attach that runs on a Multiple Virtual Storage (MVS) system. These programs name devices according to a coordinate system based on the physical location of the devices in the silo. When you configure the silo in NetWorker, you supply the name of the silo that the silo management software uses to refer to the storage device.

For tape drives, the name consists of four digits that are separated by commas:

- The first digit refers to the automated cartridge system (ACS) with which the drive is associated.
- The second digit refers to the library storage module (LSM) in which the drive is located.
- The third and fourth digits refer to the panel and slot location in which the drive is located.

A typical name for an STK drive is similar to: 1,0,1,0.

You cannot determine the drive names from the NetWorker software. Contact the silo administrator for the drive names of the devices that the NetWorker server can use. To connect to more than one drive, determine the SCSI IDs for each drive and correctly match the IDs to the silo names. If the operating system device names and silo names are accidentally swapped, NetWorker can only mount and unmount volumes. NetWorker cannot read or write to the volumes after they are mounted. To reconfigure the device names correctly, modify the Library resource in the **Administration** window and change the order of the device names in the **STL Device Names** attribute.

Installing a silo

Procedure

1. Install the silo management software on the silo server.
2. If required, install the **STLI** library on the NetWorker server. For more information, refer to the documentation from the silo vendor.

For example, for a NetWorker server or storage node running Windows to control an STK silo, the libattach program must be installed.

On UNIX systems, do not install the STLI library because all the necessary software is installed when the NetWorker software is installed.

3. Ensure that the NetWorker server is properly connected to the media devices in the silo.
4. Add the silo. [Configuring silo libraries](#) on page 175 provides further details.

Configuring silo libraries

Procedure

1. In the server's **Administration interface**, click **Devices**.
2. Open the **Storage Nodes** folder in the navigation tree.
3. Right-click the storage node to which the device is to be configured, and select **Configure All Libraries** (which is available from many of the menus throughout the Devices task). This action opens a wizard that can configure all detected libraries, except those libraries that are explicitly excluded in the library exclusion list during configuration.

 **Note:** If **Configure All Libraries** is started from the server folder instead of from the **Storage Node** folder, then all storage nodes on the NetWorker server are automatically selected for configuration in the wizard.

The **Configure All Libraries** wizard appears, and allows the user to step through library configuration, including the following input (some of which is filled in by default):

- Library type (select **STL Silo**).
 - Adjust the **Enable New Device** option, if required.
 - Current server sharing policy (use maximal sharing with Dynamic Drive Sharing [DDS]).
 - Storage nodes on which the libraries should configure. You can select a storage node to see its details that are displayed. If the appropriate storage node is not listed, click **Create a New Storage Node**. When creating a storage node, replace the default value in the **Name** field with the name of the new storage node:
 - a. Update storage node properties, if required.
 - b. Type the **Silo Controller** count, which sets the number of silos to be configured for the selected storage node. The default is 1. If a silo count of greater than one is selected, then a library name and hostname must be typed for each one.
 - c. Type the **Hostname** of the silo controller.
 - d. (Optional) Use the **Test Silo Controller Connectivity** button to see whether the connection to a silo controller works. Use it once for each silo. If the connection to a given silo fails, an error message appears.
4. Click **Start Configuration** after filling in the requested information. The **Configuration** window displays a message that the Configure All Libraries process has started, and that the configuration activity can be viewed by checking the **Monitoring > Log** screen for status.
 5. Click **Finish** on the **Configuration** window to close the configuration wizard. If problems occur during configuration, then the **Back** button on the

Configuration window becomes active, which allows the user to return to the input screen to adjust input.

NetWorker software with ACSLS silos

In this section, the term “ACSLS server” refers to the name of the system that is running any one of StorageTek's library manager programs.

The `ssi` program is used indirectly by the `nsrjb` program to communicate with an ACSLS server. The `nsrjb` program loads `libstlstk`, which handles the TCP calls to and from the `ssi` program. The `ssi` program then handles all of communication to and from the ACSLS server. Starting with ACSLS version 5.3, it is possible to run either a NetWorker server or storage node on the same host that is running ACSLS.

To configure a library, the `ssi` and `mini_el` programs must be running on the system on which library configuration is performed. The `ssi` and `mini_el` programs are generally run as background processes, and are usually started automatically by the system.

In addition to the `ssi` and `mini_el` programs, a shared library file (usually called `libstlstk.xxx` where `xxx` is an operating system-dependent extension) is also required. An appropriate version of this library is installed as part of NetWorker installation.

ACSLS silos and firewalls

With `ssi` version 2.0, communication with the ACSLS server on a specified port number is supported, using the `-a` command line option. This is part of the STK firewall enhancement. The ACSLS version 7 or later must be running on the ACSLS server to use this functionality.

The UNIX man pages for these commands, or see the *NetWorker Command Reference Guide*, which provides information on the `ssi` and `mini_el` programs.

Releasing a silo device

When a silo device is configured for use with a NetWorker server, it is possible to restrict silo access only to the NetWorker server. These restrictions allow increased availability to the silo for those with full access. These restrictions can be lifted by using the Release Device feature.

Procedure

1. In the **Administration** window, click **Devices**.
2. Open the **Libraries** folder in the navigation tree. The **Libraries** detail table appears.
3. Select a silo in the navigation tree or double-click a silo in the **Libraries** detail table to open the double-paned **Library Operations** view. The silo's drives are listed in the **Device** column. The slots are listed in the **Slot** column.
4. Right-click a silo in the **Slot** column, and select **Release Device**. A window appears and asks whether to release devices.
5. Click **Yes**. The **Library Operation** window appears and displays this message:

```
The library operation has started.
Please see the Monitoring->Operations screen for its status.
```

6. Click **OK**.
7. Repeat all steps for each device to be released.

Silo device cleaning

Do not enable automated cleaning for silos in the NetWorker software. The automated device cleaning feature depends on fixed slot numbers, so it cannot be used in a silo, which does not have fixed slot numbers. For information about how to clean devices in a silo, refer to the ACSLS silo manufacturer's software documentation.

Environment variables for StorageTek silos

Environment variables must be set for StorageTek silos. The following table lists the environment variables to set.

Table 34 StorageTek environment variables

Silo model	Environment variables
StorageTek	For UNIX systems: - CSI_HOSTNAME = <i>name_of_ACSLS_system</i> The following commands should also be running on the system and can be in the NetWorker startup script: <binaries_path>/mini_el & <binaries_path>/ssi & For Windows systems: The LibAttach Configurator program is available from StorageTek. It creates a ssi process, and a link is available to start the mini_el process from Start > Programs > LibAttach menu tree. Once installed and configured, it starts on restart.

Setting environment variables for UNIX systems

Procedure

1. Create a Bourne shell script file named `/nsr/nsrrc` on the NetWorker server if it does not already exist.
2. Add the variables in this format:


```
ENV_VAR_NAME = value
export ENV_VAR_NAME
```
3. Stop and start the NetWorker server daemons in order for the environment variables to take effect.

Media management in a silo

More than one software application can use a single silo. Therefore, media management in a silo requires extra operations to prevent the NetWorker software from overwriting volumes used by other programs.

Silo slot numbering

In a library, the NetWorker software specifies many functions by slot number. A library has a fixed number of slots, and NetWorker software uses the slot number to refer to a volume's physical location.

A silo works similarly, but a silo has a variable number of slots, starting at zero when it is first configured, and limited by the silo license purchased. The fundamental identifier of a silo volume is its barcode, or volser (volume serial number). The volser never changes over the life of a particular volume.

When the `nsrjb` command lists the contents of a silo, it also lists a slot number. Use the slot number to specify which volumes to mount, unmount, label, and inventory. Volumes are not always assigned the same slot number in the silo. The slot numbers in the silo are assigned dynamically, based on the sorted order of the barcodes that have been allocated. If additional barcodes that fall earlier in the sort sequence are allocated later, then the slot numbers change for all volumes that are later in the sequence.

The `nsrjb` UNIX man page or the *NetWorker Command Reference Guide* provide more information.

Silo volume mounting and unmounting

The mount and unmount operations for silos are the same as for library volumes.

Consider the following when mounting and unmounting library volumes:

- A volume must be mounted before it can be labeled, read, or had data written on it. The robotic mechanism mounts volumes in the devices of a silo.
- Volumes must be unmounted before they can be inventoried in a silo or removed from a NetWorker pool.

[Volume mounting and unmounting](#) on page 154 provides more information.

Silo volume labeling

The NetWorker labels for volumes in a silo include both a regular NetWorker volume label (written on the media of the volume) and a silo barcode identifier. The volume label is usually based on the volume pool's label template. The barcode identifier is written on a physical label on the outside of the volume, which the barcode reader in the silo can scan during inventory. [Labeling volumes](#) on page 148 and [Barcode labels](#) on page 149 provide instructions on how to label silo volumes.

The use of barcodes with matching barcode labels and NetWorker volume labels, are both available for a silo. The Barcode Reader attribute must be selected, however the Match Barcode Labels attribute is optional. When both attributes are selected, the internal volume label that NetWorker software writes on the media of each volume will match the barcode label on the outside of the volume. When the labels match, it is easier to track volumes. But the NetWorker software does not require the internal and external labels to match.

With most silo management software, unlabeled volumes can be used. The silo management software assigns a "virtual" barcode label to those volumes. Although volumes can be used without barcodes, it is difficult to maintain integrity, since once the volume has been removed from the silo, the information about the virtual barcode is lost. Any volume without an actual barcode can be reinserted into the silo under a virtual barcode that NetWorker software (or another application) associates with some of the data.

Using silos with volume import and export capability

NetWorker software supports the use of the import/export feature that is found in many brands of silos. Depending on the silo model, this feature is also known as CAP, mail slot, and loading port. The import/export feature deposits and withdraws volumes from slots in the silo.

About this task

The import/export feature enables the operator to deposit and withdraw cartridges without invalidating the device inventory list. If the operator opens the door to load or unload volumes, the element status of the autoloader is invalidated, requiring the time-consuming operation of reinitializing the silo. Note, however, that NetWorker software does not automatically inventory the volume after a deposit.

Either the NetWorker software or the silo management software can be used to control the import/export feature on the supported silos to deposit and withdraw volumes in a silo. But it is often more efficient to use the silo management software, especially to deposit or withdraw many volumes.

If the import/export feature is set to automatic mode, the silo management software inserts volumes automatically and the NetWorker software cannot be used to insert volumes.

To issue deposit and withdraw commands:

- To add and deposit volumes, type: `nsrjb -a -T tags -d`
- To remove and eject/withdraw volumes, type: `nsrjb -x -T tags -w` where tags specifies the tags or barcodes of volumes in a remote silo.

NOTICE You cannot deposit a volume from the CAP (I/O Port) using the `nsrjb -d` command. A silo volume deposit requires the `-T` and `-a` options in sequence to add a volume in the media database.

The sequence of operations is:

- `nsrjb -d -T Barcode`
- Ignore the error message that appears.
- `nsrjb -a -T Barcode`

Barcode IDs

A list of available barcode-labeled volumes is available from the silo management software. Refer to the silo manufacturer's documentation for how to generate the list of barcode IDs.

To specify a barcode identifier or template for the volumes from a command prompt, use the `-T` option with the `nsrjb` command. The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrjb` command.

Silo volume allocation

When volumes are added, the NetWorker server is directed to the volumes it can use.

NOTICE Because silos can be used by more than one software application, it is possible that a different application could read or write to volumes that belong to the NetWorker software. To prevent this from happening, most silo management software includes methods to limit access to volumes based on the hostname of the computer on which various programs run. The NetWorker software does not

provide a method for setting up this sort of protection. The silo management software must configure it.

The addition of a volume causes the NetWorker software to query the silo management software to verify that the requested volume exists.

If the volume exists, the volume is allocated to the NetWorker software.

Adding a silo volume

Procedure

1. In the **Administration** window, click **Devices**.
2. Open the **Libraries** folder in the navigation tree. The **Libraries** detail table appears.
3. Double-click a silo in the **Libraries** detail table to open the double-paned library operations view. The silo's drives are listed in the **Device** column, and its slots are listed in the **Slot** column.
4. Right-click a silo in the **Device** column, and select **Add**. The **Add Library Volumes** window appears, with the option to select either **Template** or **List** for barcode selection.
5. Select either **Template** or **List** to enter barcode volume identifiers.
 - The **Template** option allows the use of wildcards in creating a list of barcode IDs. Each entry should be on a separate line. For example, to name four tapes A01B, A02B, A03B, and A04B, type:

```
A0
1-4
B
```

- The **List** option allows the entry of barcode IDs, separately. Each entry should be on a separate line. For example, type the name for each tape:

```
A01B
A02B
A03B
A04B
```

6. Type the appropriate volume identifiers in the **Barcodes** field.
7. Click **OK** (or **Cancel**, to continue adding to the list).
 - Click "+" to add an entry.
 - Click "<" to insert above a highlighted selection.
 - Click "-" to delete an entry.

The Library Operation window displays this message:

The library operation has started.

The **Monitoring > Operations** > screen displays the status.

8. Click **OK**.

The **Library** detail table displays the added volumes.

Inventory silos

Taking inventory of the volumes in a silo ensures that the mapping between slot number and volume name is correct, or reconciles the actual volumes in a silo with the volumes listed in the NetWorker media database.

The slot number of a silo volume is not a numbered slot inside the silo, as it is in a library. The slot number of a silo volume is the number of the volume's position in the list of volumes in a silo.

The tasks for inventorying volumes in a silo are the same as those for a library. [Inventorying library volumes](#) on page 160 provides information about inventorying a library.

The NetWorker software examines all of the volumes in the silo and compares the new list of volumes to the NetWorker media database. Then the NetWorker software produces a message listing any volumes located in the silo that are not in the media database.

When the NetWorker software inventories a silo, the silo's barcode label reader reads the barcode labels on the outside of each volume. When a barcode matches an entry in the NetWorker media database, the volume does not need to be loaded. The inventory proceeds rapidly. If, however, the NetWorker software reads a barcode that does not match any of the entries in the media database, the volume must be mounted and read in order for a proper inventory to be taken.

Troubleshooting a silo

If the particular silo model does not automatically deposit the volume, then place the volumes in the insert area, right-click the volume, and select Deposit.

To perform the Deposit and Add operations from a command prompt:

- On silos that require manual depositing, type `nsrjb -a -T tags -d`
- On silos where the silo management software deposits volumes automatically, such as StorageTek silos, type

```
nsrjb -a -T tags
```

where:

- `tags` specifies the tags or barcodes of volumes in a remote silo.
- `-d` performs the manual deposit.

[NetWorker software interactions with a silo](#) on page 173 provides more information on STLs.

Deallocating (removing) silo volumes

When an STL volume in a silo is no longer needed, the volume can be deallocated from the silo. Deallocation is basically the same operation as removing a volume from a library. Although the volume cannot be loaded by the robotic mechanism, the entries in the NetWorker media database remain intact. If the volume is allocated again, NetWorker software can retrieve the data from it later.

About this task

Use deallocation when the silo license limits the number of usable slots, or when data is moved offsite for safer storage. When the license limits the number of slots, it might be possible to leave the volumes in the silo, if it is certain that the volumes will not be used by another application. That way, the volumes can easily be added again when the data on them must be accessible.

The allocation operation is not automatic. The volumes must be manually allocated again and reinventoried to let the NetWorker server access the data. If the volume is to be removed from the silo for offsite storage, it must be removed with NetWorker software and then ejected from the silo by using the silo management software.

Procedure

1. Unmount the volume from the device. [Volume mounting and unmounting](#) on page 154 provides instructions on unmounting volumes.
2. In the **Administration** window, click **Devices**.

3. Open the **Libraries** folder in the navigation tree. The Libraries detail table appears.
4. Double-click a silo in the **Libraries** detail table to open the double-paned library operations view. The silo's drives are listed in the **Device** column.
5. Right-click a silo in the **Device** column, and select **Remove**.
The Remove Library Volumes window appears, with the option to select either Template or List for barcode selection.
6. Select either **Template** or **List** to enter barcode volume identifiers.
 - The **Template** option allows the use of wildcards in creating a list of barcode IDs. For example, to name four tapes A01B, A02B, A03B, and A04B, type A0, 1-4, and B.
 - The **List** option allows the entry of barcode IDs, separately. For example, type the name for each tape: A01B, A02B, A03B, and A04B.
7. Type the appropriate volume identifiers in the **Barcodes** field.
8. Click **OK**.
 - The **Library Operation** window displays this message:

```
The library operation has started.
```

- The **Monitoring > Operations** screen displays the silo's status.
9. Click **OK**. Notice that on return to the **Libraries** detail table, the removed volumes are no longer listed.

Results

[NetWorker software interactions with a silo](#) on page 173 provides information on STLs.

NDMP libraries

NDMP libraries or devices are accessed by using the NDMP protocol and are typically used by network attached storage (NAS) systems. These devices do not allow direct access to control from the host operating system. Control and data movement is performed over the network by using the NDMP protocol.

The NDMP guide provides more information.


NetWorker hosts with shared libraries

The NetWorker software permits different NetWorker hosts (a NetWorker server or storage node) within a datazone to control individual devices within a library. This is known as library sharing.

The presence of a SAN within the datazone is not required for library sharing. Dynamic Drive Sharing (DDS) does not support sharing libraries across datazones.

How library sharing works

Library sharing enables one NetWorker host to control the library's robotic arm, while other NetWorker hosts (as well as the host controlling the robotic arm) can each control and use specific library devices. A specific device can be controlled only by a single NetWorker host. The following figure shows how multiple NetWorker hosts can share library devices.

Figure 14 How library sharing works

Library task inactivity periods

Library resources include attributes that are used by older, slower libraries that specify the number of seconds a library is inactive after certain operations (such as loading, unloading, or ejecting a volume). For example, once a tape is loaded, the library must read and, possibly, reposition the tape before the next operation can begin. This period of delay is known as *sleeping*.

While sleeping, the library cannot receive or perform other operations. Without the sleep period, the loading or unloading of volumes might fail.

The NetWorker software automatically configures default sleep periods. Change these values only when troubleshooting a library's performance, or if a NetWorker technical support specialist requests it. Typically, the higher the sleep values specified in the attributes, the longer it takes the library to perform the task. Be cautious when changing these values.

The sleep attributes and their default values are shown in this table.

Table 35 Library resource sleep attributes

Attribute	Description	Default value
Load Sleep	Number of seconds that the NetWorker software waits for a library to complete loading a cartridge.	15 seconds
Unload Sleep	Number of seconds that the NetWorker software waits for a library to complete unloading a cartridge.	60 seconds
Eject Sleep	Number of seconds that the NetWorker software waits for an eject operation to complete.	60 seconds

Table 35 Library resource sleep attributes (continued)

Attribute	Description	Default value
Deposit Timeout	Number of seconds for a library to wait for a tape to be deposited in the mail slot before it times out.	15 seconds
Withdraw Timeout	Number of seconds for a library to wait for a tape to be withdrawn from the mail slot before it times out.	15 seconds
Cleaning Delay	Number of seconds that the NetWorker software waits between the completion of a drive cleaning operation and the ejection of the cleaning cartridge from the drive.	60 seconds
Idle Device Timeout	The number of minutes NetWorker allows a device with a volume to be idle before automatically unmounting it. For specific devices, this value can be overridden. Unmounting volumes automatically (idle device timeout) on page 155 provides more information.	10 minutes
Port Polling Period	Number of seconds for a library to wait before polling a mail slot to check for the updated status.	3 seconds

Server Network Interface attribute

The Server Network Interface attributes in the Device resource are used to determine the network address or the hostname used by the `nsrmmd` program to communicate with the NetWorker server. Similarly, the Server Network Interface attribute in the Library resource is used to determine the network address or the hostname used by the `nsrlcpd` program to communicate with the NetWorker server. These attributes are displayed in the NetWorker Console in diagnostic mode only. The Server Network Interface attributes are only relevant if the device or library is connected to a storage node.

Note: For devices, the `nsrmmd` program will read the Server Network Interface value for the first enabled device from the list of storage node devices, and each subsequent `nsrmmd` started by the NetWorker server will use the same value. Therefore, the NetWorker server will always use the same Server Network Interface value for every `nsrmmd` it starts or restarts, regardless of whether or not the Server Network Interface attribute is different for each device.

Dynamic drive sharing

Dynamic Drive Sharing (DDS) is a feature that provides NetWorker software with the ability to recognize shared physical tape drives. DDS enables NetWorker software to perform the following operations:

- Skip the shared tape drives that are in use.
- Route the backups or recoveries to other available shared tape drives.

Introduction to DDS

DDS controls application requests for tape media and allows the NetWorker server and all storage nodes to access and share all attached devices.

A system administrator can configure DDS by setting a sharing policy for devices that are accessible from multiple storage nodes.

There are two terms that are central to the use of DDS are drive and device. Within the context of DDS, these terms are defined as follows:

- Drive—The physical backup object, such as a tape drive, disk, or file.
- Device—The access path to the physical drive.

 **Note:** NetWorker only supports DDS in a storage area network (SAN) Fibre Channel environment and not in a direct-connect SCSI environment.

Benefits of DDS


Enabling DDS on a NetWorker system provides these benefits:

- Reduces storage costs—You can share a single tape drive among several storage nodes. In fact, since NetWorker software uses the same open tape format for UNIX, Windows, NetWare and Linux, you can share the same tape between different platforms (assuming that respective save sets belong to the same pool).
- Reduces LAN traffic—You can configure clients as SAN storage nodes that can send save sets over the SAN to shared drives.
- Provides fault tolerance—Within a SAN environment, you can configure hardware to eliminate a single point of failure.
- Provides configuration over a greater distance—You can configure a system over a greater distance than with SCSI connections.

DDS configuration overview

The following figure illustrates the DDS process and potential device sharing configurations. This basic configuration consists of a server, two storage nodes, and a library with two tape drives.

Figure 15 Dynamic Drive Sharing



In this figure:

- Storage nodes sn_1 and sn_2 are attached to the library.
- Each storage node, on its own, has access to drive_1 and drive_2.
- With DDS enabled, both storage nodes have access to both drives and can recognize when a shared drive is in use.

This configuration requires two DDS licenses, one for each drive.

Note: Ensure that all applicable devices can be seen from each storage node by running the `inquire -l` command locally on each storage node.

DDS block-size compatibility between UNIX and Windows

With DDS enabled, drives can be shared between storage nodes on different platforms, such as UNIX and Microsoft Windows. For NetWorker software operations (such as backups and recoveries) to take place successfully, ensure that the block size is compatible between different platforms or hardware.

To ensure compatibility, make sure one of the following conditions is met:

- The various storage nodes sharing a drive support the same block sizes.
- When a tape is labeled on a drive, it is labeled with the block size defined on the storage nodes.

Block-size incompatibility between UNIX and Windows

Incompatible block-size settings between UNIX and Microsoft Windows storage nodes could result in any of these error scenarios:

- A backup taken on a UNIX node might not be recoverable on a Microsoft Windows node if the Windows node does not support large block sizes.
- A UNIX process labels and saves data to a tape and leaves the tape mounted. A Microsoft Windows process subsequently attempts to verify the label on this tape and fails because the label verification is done by reading a header from the data portion.
- A tape on a UNIX node is labeled with a large block size. The backup is started on a Microsoft Windows node and the Windows node attempts to write the backup by using the default block size. Internally, the backup on Windows is written by breaking down the big buffer of data into smaller segments of writable block sizes. Attempting to recover a specific file on Windows in this situation fails due to positioning errors on the tape. The data is still recoverable from the Windows side, since the NetWorker software will switch from using file and block positioning to reading the tape from the beginning to reach the correct position. The data might not, however, be recoverable from the UNIX side.

Unintended Access to DDS device prevention

The Reserve/Release attribute has been added to the Device resource for tape devices to support Reserve/Release, including the Persistent Reserve commands.

Reserve/Release is a mechanism that uses SCSI commands to attempt to prevent unintended access to tape drives that are connected by using a shared-access technology, such as Fibre Channel, iSCSI, or SCSI multiplexers. It is a “cooperative” and host-based mechanism, which means that all applications should respect the reservations and not purposely break them. Access is granted based on the host system that reserved the device. Other applications that run on that host cannot be prevented from accessing a reserved device.

Reserve/Release cannot prevent a malicious or badly behaved application from accessing a reserved device. It also cannot prevent all problems caused by hardware issues (such as SCSI resets or FC LIPs) from interrupting data access.

The basic sequence requires that a host reserve a tape drive (using specific SCSI commands) before attempting to access the tape drive. If this “reservation” succeeds, then the host can use the drive. If the reservation fails (usually because the device is reserved by someone else), then the host attempting the reservation should not attempt to use the drive. When a host has finished using a reserved drive, that host must release the drive by using the appropriate SCSI commands.

The reservation is maintained by the drive itself. With older (called “Simple” in NetWorker software) Reserve/Release, the reservation is based on the SCSI ID of the system that issued the reserve command. For tape drives connected to Fibre Channel (FC) using FC-SCSI bridges, the mapping between FC host and reservation is done inside the bridge, since the initiator on the SCSI side is always the bridge itself, regardless which host actually issued the reserve command.

For Persistent Reserve, the reservation is associated with a 64-bit “key” that is registered by the host. Several keys can be registered with a given drive at any given time, but only one may hold the active reservation. NetWorker software uses the “exclusive” reservation method for Persistent Reserve. Only the host that holds the active reservation is allowed to access the drive.

The Reserve/Release attribute does not support file type or advanced file type devices.

The settings that relate to Reserve/Release and Persistent Reserve are found in a device's **Properties** window, on the **Advanced** tab. They are visible only when diagnostic mode is turned on.

The default setting for Reserve/Release is None. Once any other Reserve/Release setting is selected, it works automatically, without further user intervention. The Reserve/Release attribute is supported only on Common Device Interface (CDI) platforms, so if the CDI attribute in a device's **Properties** is set to Not Used, then Reserve/Release settings are ignored.

For newer hardware, once a Reserve/Release setting (other than None) has been selected, the appropriate Persistent Reserve commands are automatically issued before a device is opened for reading or writing, and before the device is closed. With older hardware, a SCSI-2 Reserve command is issued before opening the device, and a SCSI-2 Release command is issued after the device is closed.

Reserve/Release has these possible settings:

- None (the default)
- Simple
- Persistent Reserve
- Persistent Reserve + APTPL (Activate Persist Through Power Loss)

The Persistent Reserve Key attribute has also been added. It is used with Persistent Reservation calls.

Restrictions for use of the SCSI Reserve/Release setting

There are restrictions for using the SCSI Reserve or Release setting.

Consider the following:

- It is available on CDI platforms only. Consequently, since CDI is not supported within an NDMP environment, Reserve/Release is not supported with NDMP.
- Not all drives support persistent Reserve/Release. (All drives support at least simple reserve release. The code automatically drops back from Persistent +APTPL or Persistent to Simple on drives that do not support Persistent.)
- SCSI resets can clear Simple reservations at the device.
- Even with Reserve/Release, there is no guarantee against data loss.
- If the operating system has its own Reserve/Release feature, that feature must be disabled in order for the NetWorker Reserve/Release feature to work.
- Even if all of the enterprise's NetWorker storage nodes have this feature enabled, then it is possible that, on the storage node where a backup operation is run, data loss can be caused by the operating system's utilities or by third-party programs.

DDS attributes in the device properties

Configure the attributes that DDS uses, in the **Properties** window for a device.

The attributes include:

- Hardware ID
- Shared Devices

Hardware ID attribute

The Hardware ID attribute tracks the drives that are shared between multiple hosts. Device instances that share the same physical drive across multiple hosts have the same hardware ID. The device autoconfiguration process automatically assigns the

Hardware ID to a device, or it is added when manually configuring a device. Users cannot edit the Hardware ID.

You can view the Hardware ID in the **Properties** window for a device, on the **General** tab, in the **Device Sharing** area.

NetWorker generates the Hardware ID when a device is scanned or configured. The Hardware ID consists of the following components:

- Hardware serial number
- Device type
- Worldwide part number (WWPN)
- Worldwide name (WWN)

Shared Devices attribute

The Shared Devices attribute appears on the **Operations** tab of a device's **Properties** window when in diagnostic mode. It features values that can be used to manipulate all shared instances of a drive simultaneously. This attribute enables or disables all devices that share the same Hardware ID with a single action. The following table lists allowed values and descriptions for the attribute.

Table 36 Shared Devices attributes

Value	Description
Enable All	When selected, enables all devices with the same Hardware ID.
Disable All	When selected, disables all the devices with the same Hardware ID.
Done	This value is the default setting. After the server has enabled or disabled all devices with the same Hardware ID, the attribute value is reset to Done.

You cannot configure the Shared Devices attribute with the `jbconfig` program.

Idle Device Timeout attribute and DDS

A tape might remain mounted in a drive after a backup completes. Other requests for the drive from another device path must wait during this timeout period. Use the Idle Device Timeout attribute to adjust the timeout value.

The Idle Device Timeout attribute is not specifically a DDS attribute, but is useful in configuring shared drives. This attribute appears on the device **Properties** window on the **Advanced** tab when displayed in Diagnostic Mode. The default value is 0 (zero) minutes, which means that the device never times out and you must manually eject the tape.

If the device belongs to a library, you can also specify the Idle Device Timeout value for all devices in the library. However, the library value will take effect only on those devices whose **Idle Device Timeout** value is 0. The Idle Device Timeout value for a library is located on the **Timer** tab of the library **Properties** window.

Max active devices

In a DDS environment, use the Max active devices attribute, on the **General** tab of the Storage Node resource to define the maximum number of active devices for a storage node.

This attribute sets the maximum number of devices that NetWorker may use from the storage node in a DDS configuration. In large environments with media libraries that have a large number of devices, storage nodes might not have the ability to optimize all the drives in the library. The Max active devices attribute allows you to limit the number of devices that the storage node uses at a specified time, which allows the storage node to have access to all the devices in the library, but does not limit the storage node to the number of devices it can fully optimize.

File type devices

File type devices (FTDs) are legacy devices and their use is limited. Continued support for legacy and test purposes is maintained, however you are encouraged to use AFTD or DD Boost devices in preference to FTD. An FTD can be configured on the NetWorker server by creating a new Device resource in the same manner as for other storage devices.

The following conditions and restrictions apply to FTDs:

- The upper limit of save set size on an FTD may be either:
 - The upper limits supported by the operating system
 - The file size specified by the disk device vendor
- If multiple FTDs are configured on a system, each device must have a unique name.
- To use multiple FTDs on the same disk, partition the disk and create only one FTD per partition.
- Dynamic Drive Sharing is *not* supported.
- For FTDs created on a UNIX or Linux network file system (NFS):
 - The file system used for the FTD must not be used for any other data.
 - There must be one FTD per NFS system.
 - The Volume Default Capacity attribute for the FTD must be set to a size that is less than 100 percent of the total capacity of the file system.

 **NOTICE** Data loss will result if a full FTD is made appendable while a backup is pending completion and a save set is partially written to the full FTD. In this case, the partial save set (currently in “incomplete” state) will be overwritten.

FTD capacity issues

For FTDs, the Volume Default Capacity is a hard limit on the amount of data that can be written to the device. The Volume Default Capacity value is an estimate of what the volume capacity is likely to be. If the value is not set correctly, the NetWorker percent-used calculation will be incorrect.

 **Note:** By contrast, AFTDs ignore the Volume Default Capacity value to allow dynamic expansion of disk space.

The Volume Default Capacity attribute displays on the Configuration tab of the Device properties when Diagnostic Mode (View > Diagnostic Mode) is enabled:

- To avoid accidentally filling an FTD, set the Volume Default Capacity attribute to restrict the size of the device. For example, if a capacity of 100 MB is set, then the device will be marked full when 100 MB is reached.
 - Volume Default Capacity attribute must not be set to a value of more than 4 TB.
 - If the Volume Default Capacity of a volume changes, the changes do not take effect until the FTD is re-created, the directory contents are deleted, and the volume is relabeled.
-  **NOTICE** If the FTD is used before the Volume Default Capacity attribute is set, then the legacy data on that FTD must be staged or cloned to another device. Otherwise, this data will be overwritten.

Full FTD prevention

To prevent the file system from becoming full when backing up data to FTDs, policies can be used to move the data off the disk as soon as necessary. Save sets from FTDs can be staged or cloned to an AFTD to take advantage of advanced file type device features.

To make space for additional backups:

- Configure a save set staging policy. [Staging save sets](#) on page 509 provides details.
- Review and, if required, modify the retention policy of the save sets.

Stand-alone devices

A Device resource must be created for each stand-alone tape device on a storage node. Stand-alone drives must be configured individually.

Storage nodes must have been created before devices can be configured to be used by them. [Storage nodes](#) on page 98 provides information about storage nodes and how to create them. Note that all scanning for devices is done at the storage node level, and can be done across multiple storage nodes. Only devices that have serial numbers can be autoconfigured. Use the `jbconfig` command to configure devices that do not have serial numbers.

 **Note:** Devices must be updated to the most recent firmware and drivers.

Autodetecting and configuring a stand-alone tape drive

You can configure a new stand-alone tape drive, automatically by using Scan for Devices.

Procedure

1. In the server's **NetWorker Administration interface**, click **Devices**.
2. Right-click **Devices** in the navigation tree, and select **Scan for Devices** to detect available devices. The **Scan for Devices** window appears.
3. Click **Start Scan**.
4. Check the scan status by clicking the **Monitoring** button and selecting the **Log** tab. Then return to the Devices navigation tree.
5. Select either the **Devices** folder or the **Storage Nodes** folder in the navigation tree. All detected drives are listed. Any still-unconfigured drives are preceded by a circular icon that displays a wrench.

6. Right-click the stand-alone drive to be configured, and select **Configure Drive**. A **Configuration** dialog box appears.
7. Click **Yes** to confirm that the drive should be configured. The new drive is automatically configured.

Adding a stand-alone device manually

Procedure

1. In the server's **NetWorker Administration** interface, click **Devices**.
2. Right-click **Devices** in the navigation tree, and select **New**. The **Create Device** window appears, with the **General** tab selected, and a default device path in the **Name** field of the **Identity** area of the window.
3. Replace the default name with the path and name of the device:
 - a. If the device is configured on the server's storage node, the name is the simple device path, such as `/tmp/d0` for a file type device. A tape device on Windows would have a format similar to `\\.\Tape0`.
 - b. If the device is configured on a remote storage node, then the name must indicate that the storage node is remote by including `rd=` and the name of the remote storage node in the device path. For example, if the remote storage node is `neptune`, then the device path might be `rd=neptune:/tmp/d0` or `rd=neptune:\\.\Tape0`.

[File type devices](#) on page 190 provides instructions and restrictions on backing up to a file type device.
4. In the **Identity** area, configure the following:
 - a. In the **Comment** field, add an optional, descriptive comment.
 - b. In the **Media Type** field, select a media type.
5. In the **Status** area, configure the applicable checkboxes:
 - **Read Only**
 - **Auto Media Management**
6. In the **Cleaning** area, configure the applicable fields:
 - **Cleaning Required**
 - **Cleaning Interval**

The Date Last Cleaned is filled in automatically once a drive has been cleaned.
7. Select the **Configuration** tab to set attributes, such as:
 - Target Sessions
 - Max Sessions
 - Local Backup to a dedicated storage node

NDMP settings (NDMP remote username and password are required for an NDMP device that acts as a storage node.)
8. Click **OK** when the configuration is complete.

Auto Media Management for stand-alone devices

The Auto Media Management feature can be enabled for stand-alone devices during manual device configuration, or from the **Properties** window after configuration.

When Auto Media Management is enabled for a stand-alone device, the following processes occur when a volume becomes full during a backup:

- A notification is sent that indicates that the server or storage node is waiting for a writable volume. Simultaneously, the NetWorker server waits for the full, verified volume to be unmounted.
- The device is monitored and the software waits for another volume to be inserted into the device.
- After a volume is detected, a check is performed to determine whether the volume is labeled. If so:
 - The volume is mounted into the device.
 - The NetWorker server checks to see whether the newly mounted volume is a candidate to receive data:
 1. If yes, the write operation continues.
 2. If no, the NetWorker server continues to wait for a writable volume to continue the backup.
- If the volume is recyclable and is a member of the required pool, it is recycled the next time a writable volume is needed.
- If the volume is unlabeled, it is labeled when the next writable volume is needed for a save. Note that Auto media management does not label disk type devices such as AFTD and Data Domain.

 **NOTICE** If a partially full volume is unmounted, the NetWorker server automatically ejects the volume after a few seconds. If a stand-alone device is shared between storage nodes, then Auto Media Management should not be enabled for more than one instance of the device. Enabling Auto Media Management for more than one instance of the stand-alone device will tie up the device indefinitely. No data is sent to the device and no pending message is sent.

Mounting or unmounting a volume in a stand-alone tape drive

Procedure

1. Manually insert a volume in the stand-alone drive, or ensure that a volume is already loaded.

In a stand-alone device, a volume that has been loaded into the drive is not considered to be mounted until it has been explicitly mounted in the user interface or from the command prompt.
2. In the **Administration** window, click **Devices**.
3. Select **Devices** in the navigation tree. The **Devices detail** table appears.
4. Select the device. To mount the volume, in the **Devices detail** table, right-click the device, and select **Mount**.
5. To unmount the volume, in the **Devices > detail** table, right-click the device, and select **Unmount**.

- The **Library Operation** window displays this message:

```
The library operation has started.
```

- The **Monitoring > Operations** screen displays its status.

6. Click **OK**.

Labeling and mounting a volume in one operation (stand-alone tape drive)

When multiple storage devices are connected to the NetWorker server, the device for labeling must first be selected from the list of available devices. Remember that labeling a volume makes it impossible for the NetWorker server to recover original data from that volume.

Procedure

1. In the **Administration** window, click **Devices**.
2. Manually insert an unlabeled or recyclable volume in the NetWorker server storage device, or ensure that a volume of this type is already present for the NetWorker server to access.
3. Select **Devices** in the navigation tree. The **Devices detail** table appears.
4. Right-click the stand-alone device in the detail table, and select **Label**. The **Label** window appears:
 - a. Type a unique label name, or accept the default name that is associated with the selected pool.

 If the volume is unlabeled, the NetWorker server assigns the next sequential label from the label template that is associated with the selected pool. If a recyclable volume from the same pool is being re-labeled, then the volume label name and sequence number remain the same. Access to the original data on the volume is destroyed, and the volume becomes available.
 - b. Select a pool on the **Pools** menu. The NetWorker server automatically applies the label template that is associated with the **Default** pool unless a different pool is selected.
 - c. Select the **Manual Recycle** attribute if the volume should be manually recycled.

 If the Manual Recycle attribute is enabled when the volume is labeled, the volume cannot automatically be marked as recyclable according to the retention policy. When a volume is marked as manual recycle, the NetWorker server disregards the assigned browse and retention policies. Therefore, only an administrator can mark the volume recyclable.

 A volume that has been set to manual recycle retains that setting, even after re-labeling. A Manual Recycle policy cannot be changed back to Auto Recycle by clearing the Manual Recycle checkbox. The volume must be explicitly reset to use auto recycle.
 - d. The **Mount After Labeling** attribute is selected by default. The NetWorker server automatically labels the volume, and then mounts the volume into the device.
5. Click **OK**.

6. If the volume is recyclable, a message warns that the named volume is about to be recycled, and asks whether to continue. Click **Yes** to re-label and recycle the volume.
7. After a volume is labeled and mounted in a device, the volume is available to receive data. Since the NetWorker label is internal and machine-readable, place an adhesive label on each volume that matches that internal volume label.

[Configuring a library to use volumes with barcodes](#) on page 144 provides information on using barcode labels.

 **Note:** If you are in the process of re-labeling a mounted volume and you choose not to overwrite the existing label, the volume is left in an unmounted state. To use this volume, mount it again.

Labeling volumes without mounting

Volumes can be pre-labeled without being mounted.

To label a volume without mounting, follow the same procedures as for labeling and mounting in one operation, but clear the **Mount After Labeling** attribute in the **Label** window.

Mounting uninventoried volumes

You can mount volumes that are not included in the library inventory, but are valid (properly labelled) NetWorker volumes.

Procedure

1. In the **Administration** window, click **Devices**.
2. Select **View > Diagnostic Mode** on the toolbar.
3. Manually insert the volume in an empty library slot.
4. Open the **Libraries** folder in the navigation tree. The **Libraries** detail table appears.
5. Select the library in the navigation tree in which the volume was manually inserted, or double-click the same library in the **Libraries** detail table. The **Libraries** detail table changes to the double-paned library operations view. The library's drives are listed in the **Devices** column, and its slots are listed in the **Slot** column.
6. In the **Devices** column, right-click the library in which the volume was manually inserted, and select **Inventory**. The **Inventory Library** window appears.
7. Type the slot number of the volume in both the **First** and **Last** field of the **Slot Range**.
8. Select **Operation Type**: either **Slow/Verbose** (the default) or **Fast/Silent**.
 - When **Slow/Verbose** is selected, the **Supply Input** option and icon on the **Operations** screen of the **Monitoring** window can be used to confirm the choice to relabel a volume. The device path appears in the **Device** field.
 - When **Fast/Silent** is selected, the **Supply Input** option and icon are not available, and relabeling proceeds automatically, without user input. The device path does not appear in the **Device** field. [Entering user input](#) on page 61 provides details.
9. Click **OK**.

- The **Library Operation** window displays this message:
The library operation has started.
 - The **Monitoring > Operations** screen displays its status.
The NetWorker software then inventories the specified slot.
10. Mount the inventoried volume.

 **NOTICE** Unlabeled tapes may not be mounted for inventorying. Unlabeled tapes can only be mounted to be labeled. An attempt to mount an uninventoried volume by using unlabeled media results in an I/O error. The volume will also be ejected.

Labeling volumes

The NetWorker software applies a label template to create a unique internal label for each volume. The label corresponds to a pool and identifies the pool for the volume during backup and other operations.

Several preconfigured label templates are supplied with the NetWorker software. You cannot delete these preconfigured label templates. [Naming label templates](#) on page 79 provides more information.

When you label a volume, the labeling process:

- Writes a label on the volume.
- Adds the volume label to the media database.
- Prepares tape media to have data written to it.

When you re-label tape, the data on the tape is effectively gone.

During data recovery, the server requests the volume that contains the required data, identifying the required volume by the name with which it was labeled.

Labeling or re-labeling library volumes

Labeling volumes in a library is time-consuming, so consider labeling volumes before it is time to back up or recover files. When a volume is re-labeled, that volume is initialized and becomes available for writing again.

Procedure

1. In the **Administration** window, click **Devices**.
2. In the left pane, select **Libraries**.
A list of libraries appears in the right pane.
3. Right-click the library and select **Label**.
Details for the selected library appear, including divided tables for devices and slots. The **Label Library Media** dialog box also appears.
4. From the **Target Media Pool** list, select the pool for the volume.
The pool determines the label template that is used to label the volume.
5. To require manual recycling of the volume, select **Allow > Manual Recycle**.
With manual recycling, the volume is not automatically marked as recyclable when all save sets expire. You must manually mark the volume as recyclable.

NOTICE A volume that has been set to manual recycle retains that setting, even after the volume is re-labeled. You must explicitly reset the volume to automatic recycle by right-clicking the volume in the **Media** window, selecting **Recycle**, and then selecting the **Auto** option.

6. To be prompted before the existing label is overwritten, select **Prompt to overwrite label**.
7. Click **OK**.

The **Library Operation** dialog box appears, stating that the library operation has started.

8. To track the status of the label operation, click **Monitoring** in the **Administration** window.
9. If you selected **Prompt to overwrite label**, confirm the overwrite of the existing volume label with a new label:
 - a. Right-click the label operation in the **Monitoring** window and select **Supply Input**.
A confirmation message appears.
 - b. Click **Yes**.

Verifying the label when a volume is unloaded

If a SCSI reset is issued during a backup, the volume rewinds and NetWorker may overwrite the volume label.

To detect if the label is overwritten in this circumstance, select the **Verify label on eject** checkbox in the Device resource, or set the **Verify label on unload** setting in the Jukebox resource to **Yes**. With these settings, NetWorker verifies that a volume label exists before ejecting the volume. If the volume label cannot be read, all save sets on the volume are marked as suspect and the volume is marked as full.

Troubleshooting devices and autochangers

This section explains how to resolve problems with devices and autochangers.

NOTICE Do not edit device files and directories, this can result in unpredictable behavior and make it impossible to recover data.

Additional attributes in the Autochanger resource

The Autochanger resource contains attributes that provide a detailed view of the hidden options that the `nsrjb` program uses. [Displaying diagnostic mode attributes](#) on page 887 provides information about how to display hidden attributes.

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about these attributes.

NOTICE Do not change time related attributes unless advised to do so by a Technical Support representative.

Maintenance commands

NetWorker device driver software provides maintenance commands, such as `lusbinfo` and `lusdebug`, that you can use to diagnose problems on tape devices and autochangers.

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about how to use these commands.

Autodetected SCSI jukebox option causes server to stop responding

If you use the `jbconfig` command to create an autodetected SCSI jukebox and the server stops responding, perform the following steps.

1. Start the `jbconfig` program
2. Select the option that installs an SJI jukebox.
3. Type the number that corresponds to the type of jukebox you are installing.
4. Continue with `jbconfig` until this message appears:

```
Jukebox has been added successfully.
```

Autochanger inventory problems

This section provides an overview of the situations that can result in an outdated autochanger inventory of volumes and how to update the inventory. When the jukebox inventory becomes outdated, the NetWorker software cannot use the autochanger.

The autochanger inventory can become out of date when:

- You manually eject the media from the autochanger drive.
- You manually remove the media is from the autochanger.
- You open the autochanger door.

To update the inventory and enable the NetWorker software to use the autochanger again, perform the following steps.

1. Verify that the volume is correctly installed in the autochanger and that the autochanger door is closed.
2. Log in as root or administrator on the NetWorker server.
3. Reset the autochanger:

```
nsrjb -Hv
```

4. Inventory: the autochanger:

```
nsrjb -Iv
```

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrjb` command.

Destination component full messages

When you perform a manual operation on an autochanger, for example when you use the buttons on the autochanger to unload the tape drive instead of unloading the tape drive by using NetWorker operations, a message similar to the following may appear:

```
Destination component full
```

To resolve the problem, use the `nsrjb -H` command to reset the autochanger.

Tapes do not fill to capacity

The data stored on a tapes may not always fill the tape to capacity. For example, the NetWorker server can mark a tape with an advertised capacity of 4,000 MB full, after writing only 3,000 MB of data.

To enable NetWorker to use the maximum tape capacity, select the highest density device driver for the device. Additional reasons that the server appears to fill tapes prematurely include:

- Write errors occur during a backup. With any tape error, the NetWorker server marks the tape as full. To prevent tape write errors, clean the tape drive regularly and use only data-quality tapes. If cleaning the drive does not help, ensure that you perform the following actions:
 - Confirm the configuration of the device driver.
 - Set any necessary switch settings on the tape drive, based on the manufacturer specifications.
 - Confirm that all cables are secure.
 - Address other potential SCSI problems.
- Space requirements for NetWorker to create file marks. The NetWorker server periodically writes file marks to facilitate rapid recovery of data. These file marks consume varying amounts of tape space, depending on the type of tape drive. The number of file marks the server writes to the tape depends on how many save sets are on the tape. Many small save sets require more file marks than a few larger ones.
- Tape capacity differences. Two apparently identical tapes from the same vendor can vary significantly in capacity. This can cause problems when you copy one full tape to another, especially if the destination tape holds less data than the source tape.
- Data compression affects the tape capacity. If you use compression on the tape drive, you cannot predict the effect on tape capacity. A compressing drive can provide twice the capacity of a non-compressing drive. Tape capacity can vary depending on the type of backup data. For example, if a non-compressing drive writes 2 GB of data to a specific tape, the compressing drive could write 10 GB, 2 GB, 5 GB, or some other unpredictable amount of data.
- Tape length. Verify the tape lengths, for example, a 120-meter DAT tape holds more data than a 90-meter DAT tape.

Tapes get stuck in drive when labeling tapes on Linux Red Hat platform

When you label a tape in a DDS configuration on an RHEL NetWorker server, the tape may become stuck in the drive and display the following error message:

```
unload failure-retrying 30 seconds
```

To resolve this issue, set the `auto_lock` setting attribute to “0” (Off) in the `/etc/stinit.def` file for the following drive types:

- Sony AIT-2 and AIT-3
- IBM LTO Gen1
- HP LTO Gen1

- IBM LTO GEN2
- IBM 3580 drive LTO-1
- IBM 3592 J1A
- Quantum DLT 7000

By default the `auto_lock` setting is set to 1(On).

Increasing the value of Save Mount Time-out for label operations

A label operation initiated by a backup operation may take more than 30 minutes before it fails when the Auto media management option is enabled and the label operation encounters a corrupted tape.

The NetWorker software keeps a record of the location of the corrupted tape only for the current backup operation, and NetWorker can attempt to use a corrupted tape for the other backup operation, unless an operator removes the volume.

To modify the time it takes the label operation timeout, modify the Save Mount Time-out attribute for the storage node. [Configuring timeouts for storage node remote devices](#) describes how to modify the attribute.

Server cannot access autochanger control port

The control port controls the autochanger loading mechanism. The autochanger hardware installation manual contains information about how to verify that the control port is correctly connected.

If you cannot determine that the control port is working, contact the autochanger vendor for assistance.

Modifying the control port

When a change in the control port of the robotic arm of a library occurs, NetWorker may not be able to perform library operations, such as labeling, mounting, and unmounting, and inventorying. You may see the error `no such file or directory` when NetWorker tries to perform library operations.

About this task

To update the NetWorker server or storage node to use the new control port, perform the following steps.

Procedure

1. Run the `inquire` command to determine the SCSI device address of the library arm and to confirm that a serial number is reported.

NOTICE Use the `inquire` command with caution. The `inquire` command sends the SCSI inquiry command to all devices detected on the SCSI bus. If you use `inquire` during normal operations, unforeseen errors and possible data loss may result.

- If `inquire` reports the serial number of the arm, follow the procedure at [Scanning for libraries and devices](#) on page 142 to scan the library for devices, then enable the library in NMC:
 - a. In the **Administration** window, click **Devices**.
 - b. Expand the **Libraries** folder, then right-click the library and select **Enabled/Disable**.

- If inquire does not report the serial number or if the scan for devices operation does not detect the control port change, use the `nsradmin` command to change the control port:
 - a. Log in as root or as Windows administrator on the NetWorker host that manages the control port.
 - b. At the command prompt, type `nsradmin`. The `nsradmin` prompt appears.
 - c. To disable the library, type the following commands:


```
type: NSR jukebox
update enabled: no
```
 - d. When `nsradmin` prompts you to update the resource, type `yes`.
 - e. To update the control port, type:


```
update control port: scsidev@b.t.l
```

 where `b.t.l` is the bus.target.lun of the library's robotic arm (as reported by the `inquire` command).
 - f. When `nsradmin` prompts you to update the resource, type `yes`.
 - g. To re-enable the library, type:


```
update enabled: yes
```
 - h. When `nsradmin` prompts you to update the resource, type `yes`.
 - i. To verify that the control port was changed and the library is now enabled, type `print` at the `nsradmin` prompt.

Changing the sleep times required for TZ89 drive types

When you unload a volume from a TZ89 tape device you may receive an error message similar to the following and NetWorker will repeatedly try to unload the tape:

```
nsrd: media info: unload retry for jukebox `COMPAQTL895' failed
- will retry again.
```

To resolve this issue, changes the sleep attributes in the Autochanger resource.

1. Shut down NetWorker services.
2. Shut down and restart the autochanger that contains the TZ89 drives.
3. When the autochanger is back online, restart NetWorker services. NetWorker will not try to unload the drive again.
4. Use NMC to edit the following autochanger sleep time attributes, and use the following values:
 - Eject Sleep: **18** secs
 - Unload Sleep: **40** secs
 - Load Sleep: **40** secs[Additional attributes in the Autochanger resource](#) on page 197 provides information about how to set the sleep attributes.
5. Try to unload the drive again. If the drive fails to unload, repeat this procedure and increase the sleep times.

Message displayed when CDI enabled on NDMP or file type device

If you enable the CDI feature for an NDMP tape device or file type device (FTD), a message similar to the following appears:

```
nsrd: media notice: The CDI attribute for device "/dev/rmt/3cbn" has been changed to "Not used".
```

To avoid this message, do not enable the CDI attribute for these device types.

Verify firmware for switches and routers

Ensure that the switches or routers firmware that you use on the network was manufactured after August 1995. Most of the switch and router vendors have significantly improved their handling of RPC traffic since August 1995.

Commands issued with nsrjb on a multi-NIC host fail

When you run `nsrjb` commands to manage a jukebox on a NetWorker server or storage node that has multiple network interface cards (NIC), the commands may fail.

To prevent this failure, add the domain name of each additional NIC to the Aliases attribute in the Client resource for the NetWorker server or storage node. [Editing a Client resource](#) on page 485 describes how to edit a Client resource.

SCSI reserve/release with dynamic drive sharing

When the NetWorker software uses Dynamic Drive Sharing (DDS) the operating system tape driver might use the SCSI reserve/release feature in a manner that interferes with the proper operations of the NetWorker software. To resolve this issue, disable the reserve/release feature.

Solaris

The `st.conf` file contains a setting for each device type in use that enables or disables the SCSI reserve/release feature. The Tape Configuration section of the `st` man page provides more information. Use the most up-to-date `st` driver that is available for the version of Solaris.

Edit the `st.conf` file only if one of the following conditions apply:

- The NetWorker configuration includes DDS.
- Solaris `st` does not support a tape drive that is configured on a Solaris host.

To determine if the Solaris `st` tape driver supports a tape drive, perform the following steps:

1. Use the `mt` command to load a tape in the drive. For example, with the tape device file `0cbn`, the type: `mt -f /dev/rmt/0cbn status`
 - If the output of the `mt` command includes the line SCSI tape drive or appears similar to the following, the `st` tape driver uses generic settings, which do not support the tape drive:

```
mt -f /dev/rmt/4cbn status
Vendor 'IBM ' Product 'ULT3580-TD2 ' tape drive:
sense key(0x6)= Unit Attention residual= 0
retries= 0 file no= 0 block no= 0
```

Tape operations may appear to work in NetWorker but you may run into problems when you try to recover saved data.

- If the output of the `mt` command appears similar to the following, the `st` tape driver recognizes the drive and uses the correct internal settings to manage the drive:

```
mt -f /dev/rmt/0cbn status
HP Ultrium LTO tape drive:
sense key(0x0)= No Additional Sense residual= 0
retries= 0 file no= 0 block no= 0
```

In this configuration, you must only edit the `st.conf` file when you use the drive in a DDS configuration.

AIX

To reset the reserve/release setting on an AIX operating system, use the SMIT interface.

1. From the **Devices** menu, select **Tapes**.
2. Change the value for the **RESERVE/RELEASE** support attribute from **No** to **Yes**.

HP-UX

To reset the reserve/release setting on an HP-UX 11 operating system, perform the following steps.

1. Change the `st_ats_enable` kernel variable to a value other than zero.
2. (Optional) Restart the computer to ensure that the operating system implements the change.

 **Note:** The reserve/release is a fixed setting in HP-UX 10.

Recovering save sets from a VTL on a different NetWorker server

The following procedure describes the steps that you need to perform before you can load a tape that was in a VTL managed by one NetWorker server into a different NetWorker server.

Before you begin

Ensure the destination VTL is the same model, has the same drive names and the same number of drives as the original VTL.

Procedure

1. Confirm the inventory of the VTL in the destination NetWorker storage node
2. Run the **inquire** command to determine the Control port of the VTL on the destination NetWorker storage node.
3. Run the **sjimm** command to load the tape into a drive on the destination NetWorker server.
4. Use the `mt` command to ensure that the tape status is online. For example: `mt -f device_name status`

When the `mt` command reports that the tape drive is online, you can use the `scanner` command to scan the save set information into the media database and client file index of the destination NetWorker server.

CHAPTER 4

Data Protection Policies

This chapter contains the following topics:

• Overview of protection policies.....	206
• Designing data protection policies.....	207
• Policy Notifications.....	263
• Monitoring policy activity.....	263
• Policy log files.....	265
• Starting, stopping, and restarting policies.....	267
• Starting actions in a workflow for an individual client.....	269
• Modifying data protection Policy resources.....	270
• Configuring NSR Protection Policies from nsradmin	283
• Managing policies from the command prompt.....	290
• Identifying clients that missed the workflow schedule.....	298
• Managing policies using the NetWorker Management Web UI.....	299
• Troubleshooting policies.....	352

Overview of protection policies

A protection policy allows you to design a protection solution for your environment at the data level instead of at the host level. With a data protection policy, each client in the environment is a backup object and not simply a host.

Data protection policies enable you to back up and manage data in a variety of environments, as well as to perform system maintenance tasks on the NetWorker server. You can use either the **NetWorker Management Web UI** or the NMC **NetWorker Administration** window to create your data protection policy solution.

A data protection policy solution encompasses the configuration of the following key NetWorker resources:

Policies

Policies provide you with a service-catalog approach to the configuration of a NetWorker datazone. Policies enable you to manage all data protection tasks and the data protection lifecycle from a central location.

Policies provide an organizational container for the workflows, actions, and groups that support and define the backup, clone, management, and system maintenance actions that you want to perform.

Workflows

The policy workflow defines a list of actions to perform sequentially or concurrently, a schedule window during which the workflow can run, and the protection group to which the workflow applies. You can create a workflow when you create a new policy, or you can create a workflow for an existing policy.

A workflow can be as simple as a single action that applies to a finite list of Client resources, or a complex chain of actions that apply to a dynamically changing list of resources. In a workflow, some actions can be set to occur sequentially, and others can occur concurrently.

You can create multiple workflows in a single policy. However, each workflow can belong to only one policy. When you add multiple workflows to the same policy, you can logically group data protection activities with similar service level provisions together, to provide easier configuration, access, and task execution.

Protection groups

Protection groups define a set of static or dynamic Client resources or save sets to which a workflow applies. There are also dedicated protection groups for backups in a VMware environment or for snapshot backups on a NAS device. Review the following information about protection groups:

- Create one protection group for each workflow. Each group can be assigned to only one workflow.
- You can add the same Client resources and save sets to more than one group at a time.
- You can create the group before you create the workflow, or you can create the group after you create the workflow and then assign the group to the workflow later.

Actions

Actions are the key resources in a workflow for a data protection policy and define a specific task (for example, a backup or clone) that occurs on the client resources in the group assigned to the workflow. NetWorker uses a work list to define the task. A

work list is composed of one or several work items. Work items include client resources, virtual machines, save sets, or tags. You can chain multiple actions together to occur sequentially or concurrently in a workflow. All chained actions use the same work list.

When you configure an action, you define the days on which to perform the action, as well as other settings specific to the action. For example, you can specify a destination pool, a retention period, and a target storage node for the backup action, which can differ from the subsequent action that clones the data.


When you create an action for a policy that is associated with the virtual machine backup, you can select one of the following data protection action types:

- **Backup** — Performs a backup of virtual machines in vCenter to a Data Domain system. You can only perform one VMware backup action per workflow. The VMware backup action must occur before clone actions.
- **Clone** — Performs a clone of the VMware backup on a Data Domain system to any clone device that NetWorker supports (including Data Domain system or tape targets). You can specify multiple clone actions. Clone actions must occur after the Backup action.

You can create multiple actions for a single workflow. However, each action applies to a single workflow and policy.

The following figure provides a high level overview of the components that make up a data protection policy in a datazone.

Figure 16 Data Protection Policy



Designing data protection policies

Designing and developing effective data protection policies requires thoughtful analysis of the client resources from which to back up data, the actions to perform on the data, and the order and timing of the actions.

Data protection policies can be grouped into six main strategies:

- **Traditional backups**—Includes file system backups, NDMP backups, NMDA backups, NMM backups, and Block Based Backups. The *NetWorker Network Data Management Protocol (NDMP) User Guide* provides detailed information about how to backup, clone, and recover NDMP data. The NMM and NMDA documentation provides information about how to backup, clone, and recover application data.
- **NetWorker and NMC Server database backups and maintenance activities**—Performs NetWorker server bootstrap and NMC database backups.
- **Snapshot backups**—Includes snapshot backups of supported storage arrays or appliances. You can clone snapshot data currently with the backup operation, or

after the snapshot backup completes. The *NetWorker Snapshot Management Integration Guide* describes how to configure data protection policies for EMC storage arrays and appliances with the NetWorker Snapshot Management feature.

- NAS device backups—Includes file system snapshots, and NAS snapshots. You can clone data after a snapshot backup job completes or concurrently. The *NetWorker Snapshot Management Integration Guide* describes how to configure data protection policies for snapshot backups.
- VMware backups—Includes NetWorker VMware Protection with the vProxy appliance (NVP), VMware Backup Appliance (VBA) backups, VBA checkpoint backups for disaster recovery, and virtual machine backups. The *NetWorker VMware Integration Guide* describes how to configure data protection policies for NVP, VBA, VBA checkpoint, and virtual machine backups and clones.
- Cloning- You can configure data protection policies that clone backup data by querying the media database for a list of save sets that are based on user defined criteria.

Note:

- You can also clone traditional, snapshot, bootstrap, and VMware backup data concurrently with the backup operation, or after the backup operation completes. The Integration Guides provide detailed information about how to clone Snapshot and VMware backup data.
- The NetWorker data protection policy applies to scheduled backups only, and it does not apply to manual backups. Some NetWorker module backups might appear to be scheduled backups that are initiated by a policy backup action, but they are manual backups because they are initiated or converted by a database or application. The *NetWorker Module for Databases and Applications Administration Guide* and the *NetWorker Module for SAP Administration Guide* provides additional details.

Default data protection policies in NMC's NetWorker Administration window

The NMC **NetWorker Administration** window provides you with pre-configured data protection policies that you can use immediately to protect the environment, modify to suit the environment, or use as an example to create resources and configurations. To use these pre-configured data protection policies, you must add clients to the appropriate group resource.

Note: NMC also includes a pre-configured Server Protection policy to protect the NetWorker and NMC server databases.

Platinum policy

The Platinum policy provides an example of a data protection policy for an environment that contains supported storage arrays or storage appliances and requires backup data redundancy. The policy contains one workflow with two actions, a snapshot backup action, followed by a clone action.

Figure 17 Platinum policy configuration



Gold policy

The Gold policy provides an example of a data protection policy for an environment that contains virtual machines and requires backup data redundancy.

Silver policy

The Silver policy provides an example of a data protection policy for an environment that contains machines where file systems or applications are running and requires backup data redundancy.

Bronze policy

The Bronze policy provides an example of a data protection policy for an environment that contains machines where file systems or applications are running.

Overview of configuring a new data protection policy

About this task

The following steps are an overview of the tasks to complete, to create and configure a data protection policy.

Procedure

1. Create a policy resource.

When you create a policy, you specify the name and notification settings for the policy.

2. Within the policy, create a workflow resource for each data type.

For example, create one workflow to protect file system data and one workflow to protect application data. When you create a workflow, you specify the name of the workflow, the time to start the workflow, notification settings for the workflow, and the protection group to which the workflow applies.

3. Create a protection group resource.

The type of group that you create depends on the types of clients and data that you want to protect. The actions that appear for a group depend on the group type.


4. Create one or more action resources for the workflow resource.
5. Configure client resources, to define the backup data that you want to protect, and then assign the client resources to a protection group.

Example 6 Example of a data protection policy with 2 workflows

The following figure illustrates a policy with two different workflows. Workflow 1 performs a probe action, then a backup of the client resources in Client group 1, and then a clone of the save sets from the backups. Workflow 2 performs a backup of the client resources in Dynamic client group 1, and then a clone of the save sets from the backup.

Example 6 Example of a data protection policy with 2 workflows (continued)

Figure 18 Data protection policy example



NetWorker resource considerations

When you create NetWorker workflow and action resources, consider the following recommendation:

- The total number of clients in a single workflow should not exceed 100.

Note: Action parallelism defines the maximum number of simultaneous data streams that can occur on all clients in a group that is associated with the workflow that contains the action. Data streams include back up data streams, `savefs` processes, and probe jobs. For a backup action, the default parallelism value is 100 and maximum value is 1000. For all other action types, the default value is 0, or unlimited.

Strategies for traditional backups

The primary considerations for a traditional backup strategy are the groups of Client resources, the workflows that define the series of actions that are associated with the backup, and the schedule for the backup.

Creating a policy

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Policies**, and then select **New**.

The **Create Policy** dialog box appears.

3. On the **General** tab, in the **Name** field, type a name for the policy.

The maximum number of characters for the policy name is 64.

- Legal Characters: `_ : - + = # , . % @`
- Illegal Characters: `/ \ * ? [ ] ( ) $ ! ^ ; ' " ` ~ > < & | { }`

 **Note:** After you create a policy, the **Name** attribute is read-only.

4. In the **Comment** field, type a description for the policy.
5. From the **Send Notifications** list, select whether to send notifications for the policy:
 - To avoid sending notifications, select **Never**.
 - To send notifications with information about each successful and failed workflow and action, after the policy completes all the actions, select **On Completion**.
 - To send a notification with information about each failed workflow and action, after the policy completes all the actions, select **On Failure**.
6. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

To send email messages or the `smtpmail` application on Windows, use the default mailer program on Linux:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
- On Linux, to send an email notification, type the following command:
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

```
nsrlog -f policy_notifications.log
```

```
mail -s subject recipient
```

```
/usr/sbin/sendmail -v recipient_email "subject_text"
```

- On Windows, to send a notification email, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.

- *recipient1@mailserver*—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.
7. To specify the Restricted Data Zone (RDZ) for the policy, select the **Restricted Data Zones** tab, and then select the RDZ from the list.
 8. Click **OK**.

After you finish

Create the workflows and actions for the policy.

Create a workflow for a new policy in NetWorker Administration

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the left pane, expand **Policies**, and then select the policy that you created.
3. In the right pane, select **Create a new workflow**.
4. In the **Name** field, type the name of the workflow.

The maximum number of allowed characters for the **Name** field is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[]()\$!^;'""`~><&|{ }

5. In the **Comment** box, type a description for the workflow.

The maximum number of allowed characters for the **Comment** field is 128.

6. From the **Send Notifications** list, select how to send notifications for the workflow:
 - To use the notification configuration that is defined in the policy resource to specify when to send a notification, select **Set at policy level**.
 - To send notifications with information about each successful and failed workflow and action, after the workflow completes all the actions, select **On Completion**.
 - To send notifications with information about each failed workflow and action, after the workflow completes all the actions, select **On Failure**.
7. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages, or use the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:


```
nsrlog -f policy_notifications.log
```
- On Linux, to send an email notification, type the following command:


```
mail -s subject recipient
```

- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

```
/usr/sbin/sendmail -v recipient_email "subject_text"
```

- On Windows, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- **-s *subject***—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- **-h *mailserver***—Specifies the hostname of the mail server to use to relay the SMTP email message.
- ***recipient1@mailserver***—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

8. In the **Running** section, perform the following steps to specify when and how often the workflow runs:

- a. To ensure that the actions that are contained in the workflow run when the policy or workflow starts, in the **Enabled** box, leave the option selected. To prevent the actions in the workflow from running when the policy or workflow that contains the action starts, clear this option.
- b. To start the workflow at the time that is specified in the **Start time** attribute, on the days that are defined in the action resource, in the **AutoStart Enabled** box, leave the option selected. To prevent the workflow from starting at the time that is specified in the **Start time** attribute, clear this option.
- c. To specify the time to start the actions in the workflow, in the **Start Time** attribute, use the spin boxes.

The default value is 9:00 PM.

- d. To specify how frequently to run the actions that are defined in the workflow over a 24-hour period, use the **Interval** attribute spin boxes. If you are performing transaction log backup as part of application-consistent protection, you must specify a value for this attribute in order for incremental transaction log backup of SQL databases to occur.

The default **Interval** attribute value is 24 hours, or once a day. When you select a value that is less than 24 hours, the **Interval End** attribute appears. To specify the last start time in a defined interval period, use the spin boxes.

- e. To specify the duration of time in which NetWorker can manually or automatically restart a failed or canceled workflow, in the **Restart Window** attribute, use the spin boxes.

If the restart window has elapsed, NetWorker considers the restart as a new run of the workflow. NetWorker calculates the restart window from the start of the last incomplete workflow. The default value is 24 hours.

For example, if the **Start Time** is 7:00 PM, the **Interval** is 1 hour, and the **Interval End** is 11:00 PM., then the workflow automatically starts every hour beginning at 7:00 PM. and the last start time is 11:00 PM.

9. To create the workflow, click **OK**.

After you finish

Create the actions that will occur in the workflow, and then assign a group to the workflow. If a workflow does not contain a group, a policy does not perform any actions.

Protection groups for traditional backups

A protection groups for traditional backups identifies the client resources to back up.

Traditional backups support the following types of protection groups:

- Basic client group—A static list of client resources to back up.
- Dynamic client group—A dynamic list of client resources to back up. A dynamic client group automatically generates a list of the client resources that use a client tag which matches the client tag that is specified for the group.

Create multiple groups to perform different types of backups for different Client resources, or to perform backups on different schedules. For example:

- Create one group for backups of clients in the Accounting department, and another group for backups of clients in the Marketing department.
- Create one group for file system backups and one group for backups of Microsoft Exchange data with the NetWorker Module for Microsoft.
- Create one group for a workflow with backups actions that start at 11 p.m., and another group for a workflow with backup actions that start at 2 a.m.

 **Note:** A Client resource can belong to more than one group.

Creating a basic client group

Use basic client groups to specify a static list of client resources for a traditional backup, a check connectivity action, or a probe action.

Before you begin

Create the policy and workflow resources in which to add the protection group to.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Groups** and select **New** from the drop-down, or right-click an existing group and select **Edit** from the drop-down.

The **Create Group** or **Edit Group** dialog box appears, with the **General** tab selected.

3. In the **Name** attribute, type a name for the group.

The maximum number of characters for the group name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[]()\$!^;"'~><&|{ }

 **Note:** After you create a group, the **Name** attribute is read-only.

4. From the **Group Type** list, leave the default selection of **Clients**.
5. In the **Comment** field, type a description of the group.
6. From the **Policy-Workflow** list, select the workflow that you want to assign the group to.

 **Note:** You can also assign the group to a workflow when you create or edit a workflow.

7. (Optional) To specify the Restricted Datazone (RDZ) for the group, on the **Restricted Datazones** tab, select the RDZ from the list.
8. Click **OK**.

After you finish

Create Client resources. Assign clients to a protection group, by using the Client Configuration wizard or the **General** tab on the **Client Properties** page.

Creating a dynamic client group

Dynamic client groups automatically include group settings when you add client resources to the NetWorker datazone. You can configure a dynamic group to include all the clients on the NetWorker server or you can configure the dynamic client group to perform a query that generates a list of clients that is based on a matching tag value.

About this task

A tag is a string attribute that you define in a Client resource. When an action starts in a workflow that is a member of a tagged dynamic protection group, the policy engine dynamically generates a list of client resources that match the tag value.

Use dynamic client groups to specify a dynamic list of Client resources for a traditional backup, a probe action, a check connectivity action, or a server backup action.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Groups** and select **New** from the drop-down, or right-click an existing group and select **Edit** from the drop-down.

The **Create Group** or **Edit Group** dialog box appears, with the **General** tab selected.

3. In the **Name** attribute, type a name for the group.

The maximum number of characters for the group name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: / \ * ? [] () \$! ^ ' " ` ~ > < & | { }

 **Note:** After you create a group, the **Name** attribute is read-only.

4. From the **Group Type** list, select **Dynamic Clients**. For steps 5 to 8, follow the instructions given in [Creating a client group](#).

Supported actions in traditional backup workflows

Traditional backup workflows can optionally include a probe or check connectivity action before the backup, and a clone action either concurrently with or after the backup.

Probe

A probe action runs a user-defined script on a NetWorker client before the start of a backup. A user-defined script is any program that passes a return code. If the return code is 0 (zero), then a client backup is required. If the return code is 1, then a client backup is not required.

Only a backup action can follow a probe action.

Check connectivity

A check connectivity action tests the connectivity between the clients and the NetWorker server before the start of a probe or backup action occurs. If the connectivity test fails, then the probe action and backup action does not start for the client.

Traditional backup

A traditional backup is a scheduled backup of the save sets defined for the Client resources in the assigned group. You must specify the destination storage node, destination pool, the schedule (period and activity), and the retention period for the backup.

Clone

A clone action creates a copy of one or more save sets. Cloning enables secure offsite storage, the transfer of data from one location to another, and the verification of backups.

You can configure a clone action to occur after a backup in a single workflow, or concurrently with a backup action in a single workflow. You can use save set and query groups to define a specific list of save sets to clone, in a separate workflow.

Note: The clone action clones the scheduled backup save sets only, and it does not clone the manual backup save sets. Some NetWorker module backups might appear to be scheduled backups that are initiated by a policy backup action, but they are manual backups because they are initiated or converted by a database or application. The *NetWorker Module for Databases and Applications Administration Guide* and the *NetWorker Module for SAP Administration Guide* provides more details.

Actions sequences in traditional backup workflows


Workflows enable you to chain together multiple actions and run them sequentially or concurrently.

A workflow for a traditional backup can optionally include a probe or check connectivity action before the backup, and a clone action either concurrently with or after the backup.

The following supported actions can follow the lead action and other actions in a workflow.

Workflow path from a traditional backup action

The only action that can follow a traditional backup is a clone action.

Figure 19 Workflow path from a traditional backup action

Creating a check connectivity action

A check connectivity action tests the connectivity between the clients and the NetWorker server, usually before another action such as a backup occurs.

Before you begin

Create the policy and the workflow that contain the action. The check connectivity action should be the first action in the workflow.

Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[]()\$!^;'""`~><&|{ }

3. In the **Comment** field, type a description for the action.
4. To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.

Note: When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

5. From the **Action Type** list, select **Check Connectivity**.
6. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
7. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
8. Specify a weekly, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.

- To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
9. Specify the days to check connectivity with the client:
- To check connectivity on a specific day, click the **Execute** icon on the day.
 - To skip a connectivity check on a specific day, click the **Skip** icon on the day.
 - To check connectivity every day, select **Execute** from the list, and then click **Make All**.

The following table provides details about the icons.

Table 37 Schedule icons

Icon	Label	Description
	Execute	Check connectivity on this day.
	Skip	Do not check connectivity on this day.

10. Click **Next**.

The **Specify the Connectivity Options** page appears.

11. Select the success criteria for the action:

- To specify that the connectivity check is successful only if the connectivity test is successful for all clients in the assigned group, select the **Succeed only after all clients succeed** checkbox.
- To specify that the connectivity check is successful if the connectivity test is successful for one or more clients in the assigned group, clear the checkbox.

12. Click **Next**.

The **Specify the Advanced Options** page appears.

13. (Optional) Configure advanced options and schedule overrides.

Note: Although the **Retries**, **Retry Delay**, **Inactivity Timeout**, or the **Send Notification** options appear, the Check Connectivity action does not support these options and ignores the values.

14. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

15. From the **Failure Impact** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.

 **Note:** The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

 **Note:** If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.
16. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
17. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
18. (Optional) In **Start Time** specify the time to start the action.
 Use the spin boxes to set the hour and minute values, and select one of the following options from the drop-down list:
 - **Disabled**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
19. (Optional) Configure overrides for the task that is scheduled on a specific day.
 To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:
 - Select the day in the calendar, which changes the action task for the specific day.
 - Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.
 - To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

 **Note:**

 - You can edit or add the rules in the **Override** field.
 - To remove an override, delete the entry from the **Override** field.
 - If a schedule is associated to an action, then override option is disabled.

20. Click **Next**.

The **Action Configuration Summary** page appears.

21. Review the settings that you specified for the action, and then click **Configure**.

After you finish

(Optional) Create one of the following actions to automatically occur after the check connectivity action:

- Probe
- Traditional backup
 -  **Note:** This option is not available for NAS snapshot backups.
- Snapshot backup

Creating a probe action

A probe action runs a user-defined script on a NetWorker client before the start of a backup. A user-defined script is any program that passes a return code. If the return code is 0 (zero), then a client backup is required. If the return code is 1, then a client backup is not required.

Before you begin

- Create the probe resource script on the NetWorker clients that use the probe. Create a client probe resource on the NetWorker server. Associate the client probe resource with the client resource on the NetWorker server.
- Create the policy and workflow that contain the action.
- Optional. Create a check connectivity action to precede the probe action in the workflow. A check connectivity action is the only supported action that can precede a probe action in a workflow.

Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[]()\$!^;'""`~><&|{ }

3. In the **Comment** field, type a description for the action.
4. To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.

 **Note:** When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

5. From the **Action Type** list, select **Probe**.
6. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
7. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
8. Specify a weekly, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
 - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
9. Specify the days to probe the client:
 - To perform a probe action on a specific day, click the **Execute** icon on the day.
 - To skip a probe action, click the **Skip** icon on the day.
 - To perform a probe action every day, select **Execute** from the list, and then click **Make All**.

The following table provides details on the icons.

Table 38 Schedule icons

Icon	Label	Description
	Execute	Perform the probe on this day.
	Skip	Do not perform a probe on this day.

10. Click **Next**.
The **Specify the Probe Options** page appears.
11. Specify when to start the subsequent backup action:
 - To start the backup only if all the probes associated with client resources in the assigned group succeed, select the **Start backup only after all probes succeed** checkbox.
 - To start the backup if any of the probes are associated with a client resource in the assigned group succeed, clear the **Start backup only after all probes succeed** checkbox.
12. Click **Next**.
The **Specify the Advanced Options** page appears.
13. In the **Retries** field, specify the number of times that NetWorker should retry a failed probe or backup action, before NetWorker considers the action as failed.

When the **Retries** value is 0, NetWorker does not retry a failed probe or backup action.

Note: The **Retries** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option for other actions, NetWorker ignores the values.

14. In the **Retry Delay** field, specify a delay in seconds to wait before retrying a failed probe or backup action. When the **Retry Delay** value is 0, NetWorker retries the failed probe or backup action immediately.

Note: The **Retry Delay** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. When you specify a value for this option in other actions, NetWorker ignores the values.

15. In the **Inactivity Timeout** field, specify the maximum number of minutes that a job run by an action can try to respond to the server.

If the job does not respond within the specified time, the server considers the job a failure and NetWorker retries the job immediately to ensure that no time is lost due to failures.

Increase the timeout value if a backup consistently stops due to inactivity. Inactivity might occur for backups of large save sets, backups of save sets with large sparse files, and incremental backups of many small static files.

Note: The **Inactivity Timeout** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option in other actions, NetWorker ignores the value.

16. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

17. From the **Failure Impact** list, specify what to do when a job fails:

- To continue the workflow when there are job failures, select **Continue**.
- To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.

Note: The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.

- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

Note: If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

18. Do not change the default selections for the Notification group box. NetWorker does not support notifications for probe actions and ignores and specified values.
19. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
20. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
21. (Optional) In **Start Time** specify the time to start the action.

Use the spin boxes to set the hour and minute values, and select one of the following options from the drop-down list:

- **Disabled**—Do not enforce an action start time. The action will start at the time defined by the workflow.
- **Absolute**—Start the action at the time specified by the values in the spin boxes.
- **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.

22. (Optional) Configure overrides for the task that is scheduled on a specific day.

To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:

- Select the day in the calendar, which changes the action task for the specific day.
- Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.
 - To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

Note:

- You can edit or add the rules in the **Override** field.
- To remove an override, delete the entry from the **Override** field.
- If a schedule is associated to an action, then override option is disabled.

23. Click **Next**.

The **Action Configuration Summary** page appears.

24. Review the settings that you specified for the action, and then click **Configure**.

Creating a traditional backup action

A traditional backup is a scheduled backup of the save sets defined for the Client resources in the assigned group for the workflow.

Before you begin

- Create the policy and workflow that contain the action.
- (Optional) Create actions to precede the backup action in the workflow.
Supported actions that can precede a backup include:
 - Probe
 - Check connectivity

Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: / \ * ? [] () \$! ^ ; ' " ` ~ > < & | { }

3. In the **Comment** field, type a description for the action.
4. To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.

 **Note:** When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

5. From the **Action Type** list, select **Backup**.
6. From the secondary action list, select the backup type, for example, **Traditional**.
7. (Optional) From the **Force Backup Level** list select a backup level.

For workflows that have more than one scheduled backup within a 24-hour period, use the **Force Backup Level** attribute to allow more than one backup to occur at two different backup levels in a 24-hour period. When you select a backup level in the **Force Backup Level** attribute, the first backup is performed at the scheduled backup level. Each subsequent occurrence of the backup action in the next 24 hours occurs at the level defined in the **Force Backup Level** attribute. For example, if the level defined by the schedule is Full and the **Force Backup Level** attribute is set to Incr, the first backup started by the action occurs at a level full and subsequent backups, within 24 hours of the start of the full backup are incremental. By default this option is cleared, which means that if the action runs multiple backup operations in a 24 period, all the backups occur at the scheduled backup level.

8. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
9. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
10. Specify a weekly, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
 - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
11. To specify the backup level to perform when **Define** option under **Select Schedule** is selected, click the icon on each day. If it is **Select** option under **Select Schedule**, choose the customized NSR schedule using the drop down, after specifying the backup level within the NSR schedule resource.

The following table provides details about the backup level that each icon represents.

Table 39 Schedule icons

Icon	Label	Description
	Full	Perform a full backup on this day. Full backups include all files, regardless of whether the files changed.
	Incr	Perform an incremental backup on this day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
	Cumulative Incr	Perform a cumulative incremental backup. Cumulative incremental backups include files that have changed since the last full backup.
	Logs Only	Perform a backup of only database transaction logs.
	Incremental Synthetic Full  Note: Not supported for NDMP.	Perform an incremental synthetic backup on this day. An incremental synthetic full backup includes all data that changed since the last full

Table 39 Schedule icons (continued)

Icon	Label	Description
		backup and subsequent incremental backups to create a synthetic full backup.
	Skip	Do not perform a backup on this day.

To perform the same type of backup on each day, select the backup type from the list and click **Make All**.

NetWorker does not support the use of synthetic full backup levels for NDMP data.

Celerra, Isilon, VNX, Unity, and NetApp filers with NDMP version 4 or later support token-based backups (TBB) to perform NDMP full and incremental backups. NetWorker supports the same number of incremental levels that the NAS vendor supports. Celerra, Isilon, and NetApp documentation provide the maximum number of incremental levels that the TBB incremental backup can support.

When you configure TBB after you update the NetWorker server from 7.6 SP1 or earlier, the first incremental backup does not occur until after one complete full backup.

Filers that do not support TBB, do not support incremental backups. If you select the level Incr, the NetWorker server performs a full backup.

Verify that the NAS storage vendor supports NDMP incremental backups before you use this feature.

12. Click **Next**.

The **Specify the Backup Options** page appears.

13. From the **Destination Storage Node** box, select the storage node with the devices on which to store the backup data.
14. From the **Destination Pool** box, select the media pool in which to store the backup data.
15. From the **Retention** boxes, specify the amount of time to retain the backup data.

After the retention period expires, the save set is removed from the client file index and marked as recyclable in the media database during an expiration server maintenance task.

When you define the retention policy an NDMP client, consider the amount of disk space that is required for the client file index. NDMP clients with several thousands of small files have significantly larger client file indexes on the NetWorker server than a non-NDMP client. A long retention policy for an NDMP client increases disk space requirements on the file system that contains the client file indexes.

16. From the **Client Override Behavior** box, specify how NetWorker uses certain client configuration attributes that perform the same function as attributes in the Action resource:

- **Client Can Override**—The values in the Client resource for **Schedule**, **Pool**, **Retention policy**, and the **Storage Node** attributes take precedence over the values that are defined in the equivalent Action resource attributes.

Note: If the NetWorker policy action schedule is set to the **Skip** backup level, the **Client can Override** option is not honored. For NetWorker to consider the **Client can Override** option, change the action schedule to a different level.

- **Client Can Not Override**—The values in the Action resource for the **Schedule**, **Destination Pool**, **Destination Storage Node**, and the **Retention** attributes take precedence over the values that are defined in the equivalent Client resource attributes.
- **Legacy Backup Rules**—This value only appears in actions that are created by the migration process. The updating process sets the **Client Override Behavior** for the migrated backup actions to **Legacy Backup Rules**.

17. Click **Next**.

The **Specify the Advanced Options** page appears.

18. In the **Retries** field, specify the number of times that NetWorker should retry a failed probe or backup action, before NetWorker considers the action as failed. When the **Retries** value is 0, NetWorker does not retry a failed probe or backup action.

Note: The **Retries** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option for other actions, NetWorker ignores the values.

19. In the **Retry Delay** field, specify a delay in seconds to wait before retrying a failed probe or backup action. When the **Retry Delay** value is 0, NetWorker retries the failed probe or backup action immediately.

Note: The **Retry Delay** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. When you specify a value for this option in other actions, NetWorker ignores the values.

20. In the **Inactivity Timeout** field, specify the maximum number of minutes that a job run by an action can try to respond to the server.

If the job does not respond within the specified time, the server considers the job a failure and NetWorker retries the job immediately to ensure that no time is lost due to failures.

Increase the timeout value if a backup consistently stops due to inactivity. Inactivity might occur for backups of large save sets, backups of save sets with large sparse files, and incremental backups of many small static files.

Note: The **Inactivity Timeout** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option in other actions, NetWorker ignores the value.

21. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

22. From the **Failure Impact** list, specify what to do when a job fails:

- To continue the workflow when there are job failures, select **Continue**.
- To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.
 **Note:** The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.
- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

 **Note:** If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

23. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
24. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
25. (Optional) In **Start Time** specify the time to start the action.

Use the spin boxes to set the hour and minute values, and select one of the following options from the drop-down list:

- **Disabled**—Do not enforce an action start time. The action will start at the time defined by the workflow.
- **Absolute**—Start the action at the time specified by the values in the spin boxes.
- **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.

26. (Optional) Configure overrides for the task that is scheduled on a specific day.

To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:

- Select the day in the calendar, which changes the action task for the specific day.
- Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.

- To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

 Note:

- You can edit or add the rules in the **Override** field.
- To remove an override, delete the entry from the **Override** field.
- If a schedule is associated to an action, then override option is disabled.

27. From the **Send Notifications** list box, select whether to send notifications for the action:
 - To use the notification configuration that is defined in the Policy resource to send the notification, select **Set at policy level**.
 - To send a notification on completion of the action, select **On Completion**.
 - To send a notification only if the action fails to complete, select **On Failure**.
28. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:


```
nsrlog -f policy_notifications.log
```
- On Linux, to send an email notification, type the following command:


```
mail -s subject recipient
```
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

```
/usr/sbin/sendmail -v recipient_email "subject_text"
```

- On Window, to send a notification email, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.

- *recipient1@mailserver*—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

29. Click **Next**.

The **Action Configuration Summary** page appears.

30. Review the settings that you specified for the action, and then click **Configure**.

After you finish

(Optional) Create a clone action to automatically clone the save sets after the backup. A clone action is the only supported action after a backup action in a workflow.

Creating a clone action

A clone action creates a copy of one or more save sets. Cloning allows for secure offsite storage, the transfer of data from one location to another, and the verification of backups.

Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[]()\$!^;'""`~><&|{ }

3. In the **Comment** field, type a description for the action.
4. To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.

 **Note:** When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

5. From the **Action Type** list, select **Clone**.
6. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
7. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
8. Specify a weekly, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.

- To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
 - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
9. Specify the days to perform cloning:
- To clone on a specific day, click the **Execute** icon on the day.
 - To skip a clone on a specific day, click the **Skip** icon on the day.
 - To check connectivity every day, select **Execute** from the list, and then click **Make All**.

The following table provides details on the icons.

Table 40 Schedule icons

Icon	Label	Description
	Execute	Perform cloning on this day.
	Skip	Do not perform cloning on this day.

10. Click **Next**.

The **Specify the Clone Options** page appears.

11. In the **Data Movement** section, define the volumes and devices to which NetWorker sends the cloned data:
- From the **Destination Storage Node** list, select the storage node with the devices on which to store the cloned save sets.
 - In the **Delete source save sets after clone completes** box, select the option to instruct NetWorker to move the data from the source volume to the destination volume after clone operation completes. This is equivalent to staging the save sets.
 - From the **Destination Pool** list, select the target media pool for the cloned save sets.
 - From the **Retention** list, specify the amount of time to retain the cloned save sets.
- After the retention period expires, the save sets are marked as recyclable during an expiration server maintenance task.
12. In the **Filters** section, define the criteria that NetWorker uses to create the list of eligible save sets to clone. The eligible save sets must match the requirements that are defined in each filter. NetWorker provides the following filter options:
- Time filter—In the **Time** section, specify the time range in which NetWorker searches for eligible save sets to clone in the media database. Use the spin boxes to specify the start time and the end time. The **Time** filter list includes the following options to define how NetWorker determines save set eligibility, based on the time criteria:
 - **Do Not Filter**—NetWorker inspects the save sets in the media database to create a clone save set list that meets the time filter criteria.

- **Accept**—The clone save set list includes save sets that are saved within the time range and meet all the other defined filter criteria.
 - **Reject**—The clone save set list does not include save sets that are saved within the time range and meet all the other defined filter criteria.
- b. **Save Set filter**—In the **Save Set** section, specify whether to include or exclude ProtectPoint and Snapshot save sets, when NetWorker searches for eligible save sets to clone in the media database. The **Save Set** filter list includes the following options to define how NetWorker determines save set eligibility, based on the save set filter criteria:
- **Do Not Filter**—NetWorker inspects the save sets in the media database to create a clone save set list that meets the save set filter criteria.
 - **Accept**—The clone save set list includes eligible ProtectPoint save sets or Snapshot save sets, when you also enable the ProtectPoint checkbox or Snapshot checkbox.
 - **Reject**—The clone save set list does not include eligible ProtectPoint save sets and Snapshot save sets when you also enable the ProtectPoint checkbox or Snapshot checkbox.
-  **Note:** For NAS device, only Snapshot save set is applicable.
- c. **Clients filter**—In the **Client** section, specify a list of clients to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The **Client** filter list includes the following options, which define how NetWorker determines save set eligibility, based on the client filter criteria:
- **Do Not Filter**—NetWorker inspects the save sets that are associated with the clients in the media database, to create a clone save set list that meets the client filter criteria.
 - **Accept**—The clone save set list includes eligible save sets for the selected clients.
 - **Reject**—The clone save set list does not include eligible save sets for the selected clients.
- d. **Levels filter**—In the **Levels** section, specify a list of backup levels to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The **Levels** filter list includes the following options to define how NetWorker determines save set eligibility, based on the level filter criteria:
- **Do Not Filter**—NetWorker inspects the save sets regardless of the level in the media database, to create a clone save set list that meets all the level filter criteria.
 - **Accept**—The clone save set list includes eligible save sets with the selected backup levels.
 - **Reject**—The clone save set list does not include eligible save sets with the selected backup levels.

 **Note:** For NAS device, only full backup level is applicable.

13. Click **Next**.

The **Specify the Advanced Options** page appears.

14. Configure advanced options, including notifications and schedule overrides.

Note: Although the **Retries**, **Retry Delay**, or the **Inactivity Timeout** options appear, the clone action does not support these options and ignores the values.

15. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

16. From the **Failure Impact** list, specify what to do when a job fails:

- To continue the workflow when there are job failures, select **Continue**.
- To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.
- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

Note: The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.

Note: If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

17. From the **Send Notifications** list box, select whether to send notifications for the action:

- To use the notification configuration that is defined in the Policy resource to send the notification, select **Set at policy level**.
- To send a notification on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

18. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:

```
nsrlog -f policy_notifications.log
```

- On Linux, to send an email notification, type the following command:

```
mail -s subject recipient
```

- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

```
/usr/sbin/sendmail -v recipient_email "subject_text"
```

- On Window, to send a notification email, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- **-s subject**—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- **-h mailserver**—Specifies the hostname of the mail server to use to relay the SMTP email message.
- **recipient1@mailserver**—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

19. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.

20. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.

21. (Optional) In the **Start Time** option, specify the time to start the action.

Use the spin boxes to set the hour and minute values, and select one of the following options from the list box:

- **Disabled**—Do not enforce an action start time. The action will start at the time defined by the workflow.
- **Absolute**—Start the action at the time specified by the values in the spin boxes.
- **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.

22. (Optional) Configure overrides for the task that is scheduled on a specific day.

To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:

- Select the day in the calendar, which changes the action task for the specific day.
- Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.

- To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

Note:

- You can edit or add the rules in the **Override** field.
- To remove an override, delete the entry from the **Override** field.
- If a schedule is associated to an action, then override option is disabled.

23. Click **Next**.

The **Action Configuration Summary** page appears.

24. Review the settings that you specified for the action, and then click **Configure**.

After you finish

(Optional) Create a clone action to automatically clone the save sets again after this clone action. Another clone action is the only supported action after a clone action in a workflow.

Visual representation of workflows

After you create actions for a workflow, in the Administration interface, you can see a map provides a visual representation of the actions on the right side of the **Protection** window.

The following figure illustrates the visual representation of a sample workflow for a traditional backup.

Figure 20 Visual representation of a workflow



The oval icon specifies the group to which the workflow applies. The rounded rectangle icons identify actions. The parallelogram icons identify the destination pool for the action.

- You can adjust the display of the visual representation by right-clicking and selecting one of the following options:
 - **Zoom In**—Increase the size of the visual representation.
 - **Zoom Out**—Decrease the size of the visual representation.
 - **Zoom Area**—Limit the display to a single section of the visual representation.
 - **Fit Content**—Fit the visual representation to the window area.
 - **Reset**—Reset the visual representation to the default settings.
 - **Overview**—View a separate dialog box with a high-level view of the visual representation and a legend of the icons.
- You can view and edit the properties for the group, action, or destination pool by right-clicking the icon for the item, and then select **Properties**.
- You can create a group, action, or destination pool by right-clicking the icon for the item, and then select **New**.

Strategies for server backup and maintenance

When you install or upgrade the NetWorker server, the installation or upgrade process creates a default Server Protection policy for server backup and maintenance activities. You can edit the default policy, workflows, groups, and actions, or create a set of policies for server backup and maintenance.

After you install or upgrade the NMC server and then connect to the NMC GUI for the first time, the **Console Configuration** wizard prompts you to configure the NetWorker server that will backup the NMC server database.

When you define the database backup server, the **Console Configuration** wizard:

- Creates a Client resource for the NMC Server database backup. The **Save set** field for the client contains the path to the database staging directory. By default, the staging directory is in `C:\Program Files\EMC NetWorker\Management\nmcd_b_stage` on Windows and `/opt/lgtomc/nmcd_b_stage` on Linux.

Note:

The file system that contains the staging directory must have free disk space that is at least equal to the size of the current NMC database. The section "Changing the staging directory for NMC database backups" describes how to change the staging directory location.

`cst` folder is not listed as save set under **Server Protection - NMC server backup** for Linux servers. Only `/nsr/nmc/nmcd_b_stage` is listed under backed up save set. Backup of `cst` folder is taken internally, which can be verified in Recovery Wizard after the Server Protection policy backup is succeeded.

- Creates a group called NMC server.
- Adds the Client resource to the NMC server group.
- Creates a workflow that is called NMC server backup in the Server Protection policy. The workflow contains the NMC server backup action, which performs a full backup of the NMC server database every day at 2 P.M.
- Adds the NMC server group to the NMC server backup workflow.

-  **Note:** The NMC Server database backup only supports the full and skip backup levels. If you edit the NMC Server backup action and change the levels in the backup schedule to a different level, for example synthetic full, NetWorker performs a full backup of the database.

Scheduling server backup and maintenance

Server backup and maintenance activities are configured in the default workflows to start at 9 p.m. To optimize performance, ensure that the workflows start at times of minimal backup activity or other system activity.

Protection groups for NetWorker and NMC server backup and maintenance

When you install or upgrade the NetWorker server, the installation or upgrade process creates a default protection group for the NetWorker server workflows in the Server Protection policy.

Server Protection group

The Server Protection group is a default protection group to back up the NetWorker server bootstrap and client file indexes. The Server Protection group is assigned to the Server backup workflow in the default Server Protection policy. The Server backup workflow performs a bootstrap backup, which includes the NetWorker server

resource files, media database, NetWorker Authentication Service database, and client indexes for disaster recovery. The group is a dynamic client group that automatically generates a list of Client resources for the NetWorker server.

NMC server group

The NMC server group is a default protection group to back up the NMC database, which the **Console Configuration** wizard prompts you to create the first time you log in to the NMC server. The group is a client group that contains the Client resource for the NMC server and is created during the initial login and configuration of NMC server. The NMC server group is assigned to the NMC server backup workflow in the default Server Protection policy.

 **Note:** If you create custom groups for server backup and maintenance, ensure that they include both the NetWorker server and the NMC server.

Server Protection policy and workflows

When you install or upgrade the NetWorker server, the installation or upgrade process creates a Server Protection policy with default workflows to support NetWorker and NMC backup and maintenance activities.

The Server Protection policy includes the following default workflows:

Server backup

The workflow performs two actions:

- **Expiration**—An expire action to mark expired save sets as recyclable.
- **Server database backup**—A backup of the NetWorker server media database, authentication service database, and the client file indexes. The data in this backup, also called a bootstrap backup, enables you to perform a disaster recovery of the NetWorker server.

The workflow is scheduled to start daily at 10 a.m. The workflow is assigned to the default Server Protection group, which contains a dynamically generated list of the Client resources for the NetWorker server.

NMC server backup

The workflow performs a traditional backup of the NMC database. The workflow is scheduled to start a full backup daily at 2 p.m. The workflow is assigned to the default NMC server group, which contains the NMC server.

Supported actions in a server backup workflow

The NetWorker server backup workflow supports the following action types.

Server database backup

A server database backup action performs a bootstrap backup and can also include the client file indexes.

A bootstrap backup contains the following NetWorker server components:

- Media database
- Server resource files. For example, the resource (res) database and the Package Manager database (nsrccd)
- NetWorker Authentication Service database

NetWorker automatically creates a server backup action in the Server Backup workflow of the Server Protection policy. By default, a full backup of the media database, resource files, and the NetWorker Authentication Service database occurs daily. A full backup of the client file indexes occur on the first day of the month. An

incremental backup of the client file indexes occur on the remaining days of the month. The default retention policy for the server database backup is one month.

Expiration

The expiration action expires save sets in the media database based on retention time of the save set. When the retention time of the save set has been reached, NetWorker uses the `nsrim` process to expire the save set. When a save set expires, the `nsrim` process performs the following actions:

- Removes information about the save set from the client file index.
- If the save set data resides on an AFTD, removes the save set information from the media database and removes the save set data from the AFTD.
- If the save set data resides on a tape device, the `nsrim` process marks the save set as recyclable in the media database. When all save sets on a tape volume have expired, the volume is eligible for reuse.

An expiration action is created automatically in the Server maintenance workflow of the Server Protection policy. An expiration action only supports Execute and Skip backup levels.

Clone

A clone action creates a copy of one or more save sets. Cloning enables secure offsite storage, the transfer of data from one location to another, and the verification of backups.

You can configure a clone action to occur after a backup in a single workflow, or concurrently with a backup action in a single workflow. You can use save set and query groups to define a specific list of save sets to clone, in a separate workflow.

Note: The clone action clones the scheduled backup save sets only, and it does not clone the manual backup save sets. Some NetWorker module backups might appear to be scheduled backups that are initiated by a policy backup action, but they are manual backups because they are initiated or converted by a database or application. The *NetWorker Module for Databases and Applications Administration Guide* and the *NetWorker Module for SAP Administration Guide* provides more details.

Actions supported in an NMC server backup workflow

The NMC server backup workflow supports the following action types.

NMC server backup

An NMC server backup action performs a backup of the Postgres NMC database.

An NMC server backup action is created automatically in the NMC server backup workflow of the Server Protection policy. The NMC server backup action only supports the full and skip backup levels.

You can add the following action after the NMC server backup action:

Clone

A clone action creates a copy of one or more save sets. Cloning enables secure offsite storage, the transfer of data from one location to another, and the verification of backups.

You can configure a clone action to occur after a backup in a single workflow, or concurrently with a backup action in a single workflow. You can use save set and query groups to define a specific list of save sets to clone, in a separate workflow.

Note: The clone action clones the scheduled backup save sets only, and it does not clone the manual backup save sets. Some NetWorker module backups might appear to be scheduled backups that are initiated by a policy backup action, but they are manual backups because they are initiated or converted by a database or application. The *NetWorker Module for Databases and Applications Administration Guide* and the *NetWorker Module for SAP Administration Guide* provides more details.

You can add the following actions before the NMC server backup action:

Probe

A probe action runs a user-defined script on a NetWorker client before the start of a backup. A user-defined script is any program that passes a return code. If the return code is 0 (zero), then a client backup is required. If the return code is 1, then a client backup is not required.

Only a backup action can follow a probe action.

Check connectivity

A check connectivity action tests the connectivity between the clients and the NetWorker server before the start of a probe or backup action occurs. If the connectivity test fails, then the probe action and backup action does not start for the client.

Actions in the server database backup and NMC server backup workflows


Workflows enable you to chain together multiple actions and run them sequentially or concurrently.

The following supported actions can follow the lead action and other actions in a workflow.

Workflow path from a server database backup action

The Clone action is the only supported action after a server database backup action. You cannot insert an action before a server database backup action.


Figure 21 Workflow path from a server database backup action



Workflow path from an NMC server backup action

A clone action is the only supported action after an NMC server backup action. You cannot insert an action before an NMC server backup action.

Figure 22 Workflow path from an NMC server backup action



Workflow path from an expiration action

The expiration action is an independent action, which means that you can add any other action after the expiration action. It is recommended that you do not add actions


after an expiration action in the server maintenance workflow. To use the expiration action with other actions, create or modify a workflow.

Visual representation of workflows

After you create actions for a workflow, in the Administration interface, you can see a map provides a visual representation of the actions on the right side of the **Protection** window.

The following figure illustrates the visual representation of the Server Protection workflows.

Figure 23 Visual representation of the Server Protection workflows



The oval icon specifies the group to which the workflow applies. The rounded rectangle icons identify actions. The parallelogram icons identify the destination pool for the action.

You can work directly in the visual representation of a workflow to perform the following tasks:

- You can adjust the display of the visual representation by right-clicking and selecting one of the following options:
 - **Zoom In**—Increase the size of the visual representation.
 - **Zoom Out**—Decrease the size of the visual representation.
 - **Zoom Area**—Limit the display to a single section of the visual representation.
 - **Fit Content**—Fit the visual representation to the window area.
 - **Reset**—Reset the visual representation to the default settings.
 - **Overview**—View a separate dialog box with a high-level view of the visual representation and a legend of the icons.
- You can view and edit the properties for the group, action, or destination pool by right-clicking the icon for the item, and then select **Properties**.
- You can create a group, action, or destination pool by right-clicking the icon for the item, and then select **New**.

Strategies for cloning

You can use scheduled cloning or action based (automatic) cloning to manage your data.


- **Scheduled cloning**—You can have a policy, and a workflow followed by a clone action. The workflow is associated with a dynamic group. In other words, a Query or Save set protection group.
- **Action based (automatic) cloning**—You can have a policy, and a workflow followed by a backup and a clone action. The clone action can be configured as concurrent or sequential.
 - **Sequential**—When the backup action configured for a policy or workflow is triggered, backup copies are created in the selected backup pool. However, the clone action is triggered only after backup copies are created for all the selected save sets. For example, If there are save sets numbered 1 to 100, backup copies are created in order. The clone action is triggered only after the backup copy is created for save set 100.
 - i **Note:** Sequential cloning is the preferred cloning method.
 - **Concurrent**—When the backup action configured for a policy or workflow is triggered, backup copies are created in the selected backup pool. The clone action is triggered even if only a single back up copy is created from the selected save sets. For example, If there are save sets numbered 1 to 100, backup copies are created in order. The clone action for save set 1 is triggered as soon as the backup copy for save set 1 is created. However, for performance optimization, clones for save sets are triggered in batches.

You can also use automated multi-streaming (AMS) when cloning your data to speed up the replication process.

If you are replicating save sets between two Data Domain devices on different machines, replication using NetWorker takes longer because each save set uses a single stream. The use of automated multi-streaming (AMS) splits up large files (files larger than 3.5 GiB) into multiple smaller 2 GiB slices, replicates those slices individually, and recreates the original large file on the destination DDR using those slices.

The following diagram illustrates replication using AMS.

Figure 24 Replication using AMS



AMS is supported only if:

- Both the source and destination Data Domain systems support the virtual synthetic capability (DDOS 5.5 and later). This can be validated through `ddboost option show` command as shown below:

```

ddboost@localhost# ddboost option show
Option                               Value
-----
distributed-segment-processing      enabled
virtual-synthetics                  enabled
fc                                   enabled
global-authentication-mode          none
global-encryption-strength          none

```

- The save set file being copied is large enough for the use of AMS to provide an improvement over normal replication.
- All save set types other than VBA, vProxy, Hyper-V, BBB, and synthetic full. The exception is for Microsoft NMM Exchange module save sets, where AMS is used even though it uses BBB and synthetic full.

Enable AMS, if the underlying bandwidth between two DDRs is 10Gbps. Because the use of AMS creates multiple streams, there must be enough bandwidth between the two DDRs being used for the clone workflow.

The `nsrcloneconfig` file enables you to add debug flags, control cloning sessions, and use the AMS functionality. It must be manually created under the `/nsr/debug` folder.

By default, AMS is disabled. To enable AMS, ensure that the `ams_enabled` flag is set to Yes.

The following table describes the `nsrcloneconfig` file details and their default values.

Table 41 `nsrcloneconfig` file details

Settings	Default value	Description
<code>ams_enabled</code>	Yes	Enables or disables AMS support. The value can be Yes or No.
<code>ams_slice_size_factor</code>	31	Allows you to change the slice size factor value. The slice size factor corresponds to the size of the slices desired, specified by a number of bits. For example, if the slice size factor is 28, the desired slice size is 2^{28} , or 256 MiB. The default value is 31, meaning the desired slice size is 2^{31} , or 2 GiB. The default value of 31 provides the best performance during chopping and joining.
<code>ams_preferred_slice_count</code>	0	Allows you to change the preferred slice count. There is no maximum value.
<code>ams_min_concurrent_slice_count</code>	1	Allows you to increase the minimum number of concurrent file copies. If the specified value is less than the default minimum value, the default value is used.
<code>ams_max_concurrent_slice_count</code>	20	Allows you to decrease the maximum number of concurrent file copies. If the specified value exceeds the default maximum value, the default value is used.
<code>ams_force_multithreaded</code>	No	Force AMS to use threads even when the DDRs support multi-file copies. Because the

Table 41 nsrcloneconfig file details (continued)

Settings	Default value	Description
		multi-file workflow is faster, this is only useful for explicitly testing the multithreaded workflow. The value can be Yes or No.
Debug	9	

 **Note:** The Backup Data Management chapter describes how you can clone save sets manually by using the `nsrclone` command.

Road map for configuring a new cloning data protection policy

This road map provides a high level overview of how to configure a new policy for clone operations.

Before you begin

Configure the backup policy to back up the data that is cloned.

Procedure

1. Create a group to define the data to clone.
2. Create a policy. When you create a policy, you specify the name and notification settings for the policy.
3. Within the policy, create a workflow. When you create a workflow, you specify the name of the workflow, the schedule for running the workflow, notification settings for the workflow, and the protection group to which the workflow applies.
4. Create one or more clone actions for the workflow.

Protection groups for a cloning workflow

You can use two types of protection groups to clone save sets in a workflow that are separate from backup workflows. The type of protection group that you use depends on the way that you plan to configure the workflow.

Use a save set group or a query group to specify a list of save sets if cloning occurs as the head action in a cloning workflow:

- **Save set group**—Use a save set group in clone-only workflows where you want to clone a specific list of save sets. Save set groups are similar to the manual clone operations in NetWorker 8.2.x and earlier.
- **Query group**—Use a query group in clone-only workflows where you want to clone save sets on an ongoing basis, based on the save set criteria that you define. Query groups are similar to the scheduled clone operations in NetWorker 8.2.x and earlier.

 **Note:** To clone save sets in a backup workflow, use basic client group or a dynamic client group. [Strategies for traditional backups](#) provides detailed information about how to create clone actions in a traditional backup workflow.

Create multiple protection groups to perform cloning in different ways as part of separate workflows, or to perform cloning for different save sets on different schedules. For example:

- Create a basic client group for a workflow that performs a traditional backup of the a client file system followed by cloning of the save sets that result from the backup. In this case, concurrent cloning can be enabled.
- Create a query group that identifies full save sets in the last two days to clone.

Creating a save set group

A save set group defines a static list of save sets for cloning or for snapshot index generation.

Before you begin

Determine the save set ID or clone ID (ssid/clonid) of the save sets for the group by using the **Administration > Media** user interface or the `mminfo` command.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Groups**, and then select **New**.
The **Create Group** dialog box appears, starting with the **General** tab.
3. In the **Name** field, type a name for the group.

The maximum number of characters for the group name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[]()\$!^;'""`~><&|{ }

 **Note:** After you create a group, the **Name** attribute is read-only.

4. From the **Group Type** list, select **Save Set ID List**.
5. In the **Comment** field, type a description of the group.
6. (Optional) To associate the group with a workflow, from the **Workflow (Policy)** list, select the workflow.
You can also assign the group to a workflow when you create or edit a workflow.
7. In the **Clone specific save sets (save set ID/clone ID)** field, type the save set ID/clone ID (ssid/clonid) identifiers.
To specify multiple entries, type each value on a separate line.
8. To specify the Restricted Data Zone (RDZ) for the group, select the **Restricted Data Zones** tab, and then select the RDZ from the list.
9. Click **OK**.

Creating a query group

A query group defines a list of save sets for cloning or snapshot index generation, based on a list of save set criteria.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Groups**, and then select **New**.
The **Create Group** dialog box appears, starting with the **General** tab.
3. In the **Name** field, type a name for the group.
The maximum number of characters for the group name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: / \ * ? [] () \$! ^ ; ' " ` ~ > < & | { }

 **Note:** After you create a group, the **Name** attribute is read-only.

4. From the **Group Type** list, select **Save Set Query**.
5. In the **Comment** field, type a description of the group.
6. (Optional) To associate the group with a workflow, from the **Workflow (Policy)** list, select the workflow.

You can also assign the group to a workflow when you create or edit a workflow.

7. Specify one or more of the save set criteria in the following table.

 **Note:** When you specify more than one save set criteria, the list of save sets only includes save sets that match all the specified criteria.

Table 42 Save set criteria

Criteria	Description
Date and time range	Specify the start date and time range for the save sets. To specify the current date and time as the end date for the range, select Up to now . To specify a time period, select Up to .
Backup level	In the Filter save sets by level section, next to the backup level for the save set, select the full checkbox.  Note: Only the full backup level is applicable for network-attached storage (NAS) devices.
Limit the number of clones	Specify the number for the limit in the Limit number of clones list. The clone limit is the maximum number of clone instances that can be created for the save set. By default, the value is set to 1, and cannot be changed for NAS or Block.  Note: When this criteria is set to 1, which is the default value, you may experience volume outage issues with Data Domain and advanced file type devices.
Client	Next to one or more client resources that are associated with the save set in the Client list, select the checkbox.
Policy	Next to the policy used to generate the save set in the Policy list, select the checkbox.
Workflow	Next to the workflow used to generate the save set in the Workflow list, select the checkbox.
Action	Next to the action used to generate the save set in the Action list, select the checkbox.
Group	Next to the group associated with the save set in the Group list, select the checkbox.
Pools	Next to the media pool on which the save set is stored in the Pools list, select the checkbox.

Table 42 Save set criteria (continued)

Criteria	Description
	 Note: You cannot select Pools for NAS.
Name	In the Filter save sets by name field, specify the name of the save set.  Note: You cannot use wildcards to specify the save set name.

If you specify multiple criteria, the save set must match all the criteria to belong to the group.

8. To specify the Restricted Data Zone (RDZ) for the group, select the **Restricted Data Zones** tab, and then select the RDZ from the list.
9. Click **OK**.

Creating a policy

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Policies**, and then select **New**.
The **Create Policy** dialog box appears.
3. On the **General** tab, in the **Name** field, type a name for the policy.
The maximum number of characters for the policy name is 64.
 - Legal Characters: _ : - + = # , . % @
 - Illegal Characters: /*?[]()\$!^;'""`~><&|{ } **Note:** After you create a policy, the **Name** attribute is read-only.
4. In the **Comment** field, type a description for the policy.
5. From the **Send Notifications** list, select whether to send notifications for the policy:
 - To avoid sending notifications, select **Never**.
 - To send notifications with information about each successful and failed workflow and action, after the policy completes all the actions, select **On Completion**.
 - To send a notification with information about each failed workflow and action, after the policy completes all the actions, select **On Failure**.
6. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

To send email messages or the `smtpmail` application on Windows, use the default mailer program on Linux:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:

```
nsrlog -f policy_notifications.log
```

- On Linux, to send an email notification, type the following command:

```
mail -s subject recipient
```

- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

```
/usr/sbin/sendmail -v recipient_email "subject_text"
```

- On Windows, to send a notification email, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.
- `recipient1@mailserver`—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

7. To specify the Restricted Data Zone (RDZ) for the policy, select the **Restricted Data Zones** tab, and then select the RDZ from the list.
8. Click **OK**.

After you finish

Create the workflows and actions for the policy.

Create a workflow for a new policy in NetWorker Administration

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the left pane, expand **Policies**, and then select the policy that you created.
3. In the right pane, select **Create a new workflow**.
4. In the **Name** field, type the name of the workflow.

The maximum number of allowed characters for the **Name** field is 64.

- Legal Characters: `_ : - + = # , . % @`
- Illegal Characters: `/\*?[]()$!^';"~`~><&|{ }`

5. In the **Comment** box, type a description for the workflow.

The maximum number of allowed characters for the **Comment** field is 128.

6. From the **Send Notifications** list, select how to send notifications for the workflow:
 - To use the notification configuration that is defined in the policy resource to specify when to send a notification, select **Set at policy level**.

- To send notifications with information about each successful and failed workflow and action, after the workflow completes all the actions, select **On Completion**.
 - To send notifications with information about each failed workflow and action, after the workflow completes all the actions, select **On Failure**.
7. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages, or use the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:

`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:

`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver
recipient2@mailserver...
```

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.
- `recipient1@mailserver`—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

8. In the **Running** section, perform the following steps to specify when and how often the workflow runs:
- a. To ensure that the actions that are contained in the workflow run when the policy or workflow starts, in the **Enabled** box, leave the option selected. To prevent the actions in the workflow from running when the policy or workflow that contains the action starts, clear this option.
 - b. To start the workflow at the time that is specified in the **Start time** attribute, on the days that are defined in the action resource, in the **AutoStart Enabled** box, leave the option selected. To prevent the workflow

from starting at the time that is specified in the **Start time** attribute, clear this option.

- c. To specify the time to start the actions in the workflow, in the **Start Time** attribute, use the spin boxes.

The default value is 9:00 PM.

- d. To specify how frequently to run the actions that are defined in the workflow over a 24-hour period, use the **Interval** attribute spin boxes. If you are performing transaction log backup as part of application-consistent protection, you must specify a value for this attribute in order for incremental transaction log backup of SQL databases to occur.

The default **Interval** attribute value is 24 hours, or once a day. When you select a value that is less than 24 hours, the **Interval End** attribute appears. To specify the last start time in a defined interval period, use the spin boxes.

- e. To specify the duration of time in which NetWorker can manually or automatically restart a failed or canceled workflow, in the **Restart Window** attribute, use the spin boxes.

If the restart window has elapsed, NetWorker considers the restart as a new run of the workflow. NetWorker calculates the restart window from the start of the last incomplete workflow. The default value is 24 hours.

For example, if the **Start Time** is 7:00 PM, the **Interval** is 1 hour, and the **Interval End** is 11:00 PM., then the workflow automatically starts every hour beginning at 7:00 PM. and the last start time is 11:00 PM.

9. To create the workflow, click **OK**.

After you finish

Create the actions that will occur in the workflow, and then assign a group to the workflow. If a workflow does not contain a group, a policy does not perform any actions.


Workflows for scheduled cloning

A workflow can contain one or more clone actions.

Supported workflow path from a clone action

Another clone action is the only supported action after a clone action.

Figure 25 Workflow path from a clone action



Creating a clone action

A clone action creates a copy of one or more save sets. Cloning allows for secure offsite storage, the transfer of data from one location to another, and the verification of backups.

Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:

- If the action is the first action in the workflow, select **Create a new action**.
- If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: / \ * ? [] () \$! ^ ; ' " ` ~ > < & | { }

3. In the **Comment** field, type a description for the action.
4. To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.

 **Note:** When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

5. From the **Action Type** list, select **Clone**.
6. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
7. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
8. Specify a weekly, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
 - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
9. Specify the days to perform cloning:
 - To clone on a specific day, click the **Execute** icon on the day.
 - To skip a clone on a specific day, click the **Skip** icon on the day.
 - To check connectivity every day, select **Execute** from the list, and then click **Make All**.

The following table provides details on the icons.

Table 43 Schedule icons

Icon	Label	Description
	Execute	Perform cloning on this day.

Table 43 Schedule icons (continued)

Icon	Label	Description
	Skip	Do not perform cloning on this day.

10. Click **Next**.

The **Specify the Clone Options** page appears.

11. In the **Data Movement** section, define the volumes and devices to which NetWorker sends the cloned data:
- From the **Destination Storage Node** list, select the storage node with the devices on which to store the cloned save sets.
 - In the **Delete source save sets after clone completes** box, select the option to instruct NetWorker to move the data from the source volume to the destination volume after clone operation completes. This is equivalent to staging the save sets.
 - From the **Destination Pool** list, select the target media pool for the cloned save sets.
 - From the **Retention** list, specify the amount of time to retain the cloned save sets.
- After the retention period expires, the save sets are marked as recyclable during an expiration server maintenance task.
12. In the **Filters** section, define the criteria that NetWorker uses to create the list of eligible save sets to clone. The eligible save sets must match the requirements that are defined in each filter. NetWorker provides the following filter options:
- Time filter—In the **Time** section, specify the time range in which NetWorker searches for eligible save sets to clone in the media database. Use the spin boxes to specify the start time and the end time. The **Time** filter list includes the following options to define how NetWorker determines save set eligibility, based on the time criteria:
 - Do Not Filter**—NetWorker inspects the save sets in the media database to create a clone save set list that meets the time filter criteria.
 - Accept**—The clone save set list includes save sets that are saved within the time range and meet all the other defined filter criteria.
 - Reject**—The clone save set list does not include save sets that are saved within the time range and meet all the other defined filter criteria.
 - Save Set filter—In the **Save Set** section, specify whether to include or exclude ProtectPoint and Snapshot save sets, when NetWorker searches for eligible save sets to clone in the media database. The **Save Set** filter list includes the following options to define how NetWorker determines save set eligibility, based on the save set filter criteria:
 - Do Not Filter**—NetWorker inspects the save sets in the media database to create a clone save set list that meets the save set filter criteria.
 - Accept**—The clone save set list includes eligible ProtectPoint save sets or Snapshot save sets, when you also enable the ProtectPoint checkbox or Snapshot checkbox.

- **Reject**—The clone save set list does not include eligible ProtectPoint save sets and Snapshot save sets when you also enable the ProtectPoint checkbox or Snapshot checkbox.

 **Note:** For NAS device, only Snapshot save set is applicable.

- c. **Clients filter**—In the **Client** section, specify a list of clients to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The **Client** filter list includes the following options, which define how NetWorker determines save set eligibility, based on the client filter criteria:
 - **Do Not Filter**—NetWorker inspects the save sets that are associated with the clients in the media database, to create a clone save set list that meets the client filter criteria.
 - **Accept**—The clone save set list includes eligible save sets for the selected clients.
 - **Reject**—The clone save set list does not include eligible save sets for the selected clients.
- d. **Levels filter**—In the **Levels** section, specify a list of backup levels to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The **Levels** filter list includes the following options define how NetWorker determines save set eligibility, based on the level filter criteria:
 - **Do Not Filter**—NetWorker inspects the save sets regardless of the level in the media database, to create a clone save set list that meets all the level filter criteria.
 - **Accept**—The clone save set list includes eligible save sets with the selected backup levels.
 - **Reject**—The clone save set list does not include eligible save sets with the selected backup levels.

 **Note:** For NAS device, only full backup level is applicable.

13. Click **Next**.

The **Specify the Advanced Options** page appears.

14. Configure advanced options, including notifications and schedule overrides.

 **Note:** Although the **Retries**, **Retry Delay**, or the **Inactivity Timeout** options appear, the clone action does not support these options and ignores the values.

15. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

16. From the **Failure Impact** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.

 **Note:** The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

 **Note:** If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.
17. From the **Send Notifications** list box, select whether to send notifications for the action:
 - To use the notification configuration that is defined in the Policy resource to send the notification, select **Set at policy level**.
 - To send a notification on completion of the action, select **On Completion**.
 - To send a notification only if the action fails to complete, select **On Failure**.
18. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:


```
nsrlog -f policy_notifications.log
```
- On Linux, to send an email notification, type the following command:


```
mail -s subject recipient
```
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:


```
/usr/sbin/sendmail -v recipient_email "subject_text"
```
- On Window, to send a notification email, type the following command:


```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- **-s *subject***—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
 - **-h *mailserver***—Specifies the hostname of the mail server to use to relay the SMTP email message.
 - ***recipient1@mailserver***—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.
19. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
 20. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
 21. (Optional) In the **Start Time** option, specify the time to start the action.
Use the spin boxes to set the hour and minute values, and select one of the following options from the list box:
 - **Disabled**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
 22. (Optional) Configure overrides for the task that is scheduled on a specific day.
To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:
 - Select the day in the calendar, which changes the action task for the specific day.
 - Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.
 - To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.
-  **Note:**

 - You can edit or add the rules in the **Override** field.
 - To remove an override, delete the entry from the **Override** field.
 - If a schedule is associated to an action, then override option is disabled.
23. Click **Next**.
The **Action Configuration Summary** page appears.
 24. Review the settings that you specified for the action, and then click **Configure**.

After you finish

(Optional) Create a clone action to automatically clone the save sets again after this clone action. Another clone action is the only supported action after a clone action in a workflow.

Visual representation of a clone workflow

After you create actions for a workflow, in the Administration interface, you can see a map provides a visual representation of the actions on the right side of the **Protection** window.

The following figure illustrates the visual representation of a clone workflow.

Figure 26 Visual representation of a clone workflow



The oval icon specifies the group to which the workflow applies. The rounded rectangle icons identify actions. The parallelogram icons identify the destination pool for the action.

You can work directly in the visual representation of a workflow to perform the following tasks:

- You can adjust the display of the visual representation by right-clicking and selecting one of the following options:
 - **Zoom In**—Increase the size of the visual representation.
 - **Zoom Out**—Decrease the size of the visual representation.
 - **Zoom Area**—Limit the display to a single section of the visual representation.
 - **Fit Content**—Fit the visual representation to the window area.
 - **Reset**—Reset the visual representation to the default settings.
 - **Overview**—View a separate dialog box with a high-level view of the visual representation and a legend of the icons.
- You can view and edit the properties for the group, action, or destination pool by right-clicking the icon for the item, and then select **Properties**.
- You can create a group, action, or destination pool by right-clicking the icon for the item, and then select **New**.

Road map to add a clone workflow to an existing policy

This road map provides a high level overview of how to create a clone workflow and add the workflow to an existing backup policy.

Before you begin

Configure the backup policy to back up the data that is cloned.


Procedure

1. Create a query or save set group to define the data to clone.
2. Add the new group to an existing policy.
3. Create a workflow in the existing policy.
4. Create one or more clone actions for the workflow.

Example: Creating a policy that has a separate workflow for cloning

The following figure provides a high level overview of the configuration of a policy that contains two workflows, one for backups and one to clone a list of save sets.

Figure 27 Example of a policy with separate workflows for backup and cloning



Note: The amount of data and length of time that is required to complete the backup can impact the ability to clone data when the backup and clone workflows are in the same policy. For example, if the clone action starts before the backup action completes, there might not be any data yet to clone, or in other cases, only the save sets that completed at the start time of the workflow is taken into account. In both cases, NetWorker marks the Clone Workflow as successful, but there is no guarantee that all the data from the backup workflow was cloned.

Editing an existing policy to create a workflow and clone action

Use the **Policies** window to create a workflow and create the clone action.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, expand **Policies**, and then select the existing policy.
3. In the right pane, right-click in the workflow section and select **New**, and select **Properties**.

The **New Workflow** dialog box appears.

4. In the **Name** field, type the name of the workflow.

The maximum number of allowed characters for the **Name** field is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: / \ * ? [] () \$! ^ ; ' " ` ~ > < & | { }

5. In the **Comment** box, type a description for the workflow.

The maximum number of allowed characters for the **Comment** field is 128.

6. From the **Send Notifications** list, select how to send notifications for the workflow:
 - To use the notification configuration that is defined in the policy resource to specify when to send a notification, select **Set at policy level**.
 - To send notifications with information about each successful and failed workflow and action, after the workflow completes all the actions, select **On Completion**.
 - To send notifications with information about each failed workflow and action, after the workflow completes all the actions, select **On Failure**.
7. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:


```
nsrlog -f policy_notifications.log
```
- On Linux, to send an email notification, type the following command:


```
mail -s subject recipient
```
- On Windows, to send a notification email, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver
recipient2@mailserver...
```

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.
- `recipient1@mailserver`—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

8. In the **Running** section, perform the following steps to specify when and how often the workflow runs:
 - a. To ensure that the actions that are contained in the workflow run when the policy or workflow starts, in the **Enabled** box, leave the option selected. To prevent the actions in the workflow from running when the policy or workflow that contains the action starts, clear this option.
 - b. To start the workflow at the time that is specified in the **Start time** attribute, on the days that are defined in the action resource, in the **AutoStart Enabled** box, leave the option selected. To prevent the workflow

from starting at the time that is specified in the **Start time** attribute, clear this option.

- c. To specify the time to start the actions in the workflow, in the **Start Time** attribute, use the spin boxes.

The default value is 9:00 PM.

- d. To specify how frequently to run the actions that are defined in the workflow over a 24-hour period, use the **Interval** attribute spin boxes. If you are performing transaction log backup as part of application-consistent protection, you must specify a value for this attribute in order for incremental transaction log backup of SQL databases to occur.

The default **Interval** attribute value is 24 hours, or once a day. When you select a value that is less than 24 hours, the **Interval End** attribute appears. To specify the last start time in a defined interval period, use the spin boxes.

- e. To specify the duration of time in which NetWorker can manually or automatically restart a failed or canceled workflow, in the **Restart Window** attribute, use the spin boxes.

If the restart window has elapsed, NetWorker considers the restart as a new run of the workflow. NetWorker calculates the restart window from the start of the last incomplete workflow. The default value is 24 hours.

For example, if the **Start Time** is 7:00 PM, the **Interval** is 1 hour, and the **Interval End** is 11:00 PM., then the workflow automatically starts every hour beginning at 7:00 PM. and the last start time is 11:00 PM.

9. In the **Groups** group box, specify the protection group to which the workflow applies.

To use a group, select a protection group from the **Groups** list. To create a protection group, click the + button that is located to the right of the **Groups** list.

10. Click **Add**.

The Policy Action Wizard appears.

11. In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[]()\$!^;"'`~><&|{ }

12. In the **Comment** field, type a description for the action.

13. To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.

 **Note:** When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

14. From the **Action type** list, select **Clone**.

15. Specify the order of the action in relation to other actions in the workflow:

- If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.

- If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
16. Specify a weekly, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
 - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
 17. Specify the days to perform cloning:
 - To clone on a specific day, click the **Execute** icon on the day.
 - To skip a clone on a specific day, click the **Skip** icon on the day.
 - To check connectivity every day, select **Execute** from the list, and then click **Make All**.

The following table provides details on the icons.

Table 44 Schedule icons

Icon	Label	Description
	Execute	Perform cloning on this day.
	Skip	Do not perform cloning on this day.

18. Click **Next**.

The **Specify the Clone Options** page appears.

19. In the **Data Movement** section, define the volumes and devices to which NetWorker sends the cloned data:
 - a. From the **Destination Storage Node** list, select the storage node with the devices on which to store the cloned save sets.
 - b. In the **Delete source save sets after clone completes** box, select the option to instruct NetWorker to move the data from the source volume to the destination volume after clone operation completes. This is equivalent to staging the save sets.
 - c. From the **Destination Pool** list, select the target media pool for the cloned save sets.
 - d. From the **Retention** list, specify the amount of time to retain the cloned save sets.

After the retention period expires, the save sets are marked as recyclable during an expiration server maintenance task.
20. In the **Filters** section, define the criteria that NetWorker uses to create the list of eligible save sets to clone. The eligible save sets must match the requirements that are defined in each filter. NetWorker provides the following filter options:

- a. **Time filter**—In the **Time** section, specify the time range in which NetWorker searches for eligible save sets to clone in the media database. Use the spin boxes to specify the start time and the end time. The **Time** filter list includes the following options to define how NetWorker determines save set eligibility, based on the time criteria:
 - **Do Not Filter**—NetWorker inspects the save sets in the media database to create a clone save set list that meets the time filter criteria.
 - **Accept**—The clone save set list includes save sets that are saved within the time range and meet all the other defined filter criteria.
 - **Reject**—The clone save set list does not include save sets that are saved within the time range and meet all the other defined filter criteria.
 - b. **Save Set filter**—In the **Save Set** section, specify whether to include or exclude ProtectPoint and Snapshot save sets, when NetWorker searches for eligible save sets to clone in the media database. The **Save Set** filter list includes the following options to define how NetWorker determines save set eligibility, based on the save set filter criteria:
 - **Do Not Filter**—NetWorker inspects the save sets in the media database to create a clone save set list that meets the save set filter criteria.
 - **Accept**—The clone save set list includes eligible ProtectPoint save sets or Snapshot save sets, when you also enable the ProtectPoint checkbox or Snapshot checkbox.
 - **Reject**—The clone save set list does not include eligible ProtectPoint save sets and Snapshot save sets when you also enable the ProtectPoint checkbox or Snapshot checkbox.
-  **Note:** For NAS device, only Snapshot save set is applicable.
- c. **Clients filter**—In the **Client** section, specify a list of clients to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The **Client** filter list includes the following options, which define how NetWorker determines save set eligibility, based on the client filter criteria:
 - **Do Not Filter**—NetWorker inspects the save sets that are associated with the clients in the media database, to create a clone save set list that meets the client filter criteria.
 - **Accept**—The clone save set list includes eligible save sets for the selected clients.
 - **Reject**—The clone save set list does not include eligible save sets for the selected clients.
 - d. **Levels filter**—In the **Levels** section, specify a list of backup levels to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The **Levels** filter list includes the following options to define how NetWorker determines save set eligibility, based on the level filter criteria:
 - **Do Not Filter**—NetWorker inspects the save sets regardless of the level in the media database, to create a clone save set list that meets all the level filter criteria.
 - **Accept**—The clone save set list includes eligible save sets with the selected backup levels.
 - **Reject**—The clone save set list does not include eligible save sets with the selected backup levels.

 **Note:** For NAS device, only full backup level is applicable.

21. Click **Next**.

The **Specify the Advanced Options** page appears.

22. Configure advanced options, including notifications and schedule overrides.

 **Note:** Although the **Retries**, **Retry Delay**, or the **Inactivity Timeout** options appear, the clone action does not support these options, and ignores the values.

23. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

24. From the **Failure Impact** list, specify what to do when a job fails:

- To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.
-  **Note:** The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.
- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

 **Note:** If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

25. From the **Send Notifications** list box, select whether to send notifications for the action:

- To use the notification configuration that is defined in the Policy resource to send the notification, select **Set at policy level**.
- To send a notification on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

26. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on

Linux and in the C:\Program Files\EMC NetWorker\nsr\logs folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:

```
nsrlog -f policy_notifications.log
```

- On Linux, to send an email notification, type the following command:

```
mail -s subject recipient
```

- On Window, to send a notification email, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.
- `recipient1@mailserver`—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

27. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
28. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
29. (Optional) Configure overrides for the task that is scheduled on a specific day.

To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:

- Select the day in the calendar, which changes the action task for the specific day.
- Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.
 - To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

 Note:

- You can edit or add the rules in the **Override** field.
- To remove an override, delete the entry from the **Override** field.

- If a schedule is associated to an action, then override option is disabled.

30. Click **Next**.

The **Action Configuration Summary** page appears.

31. Review the settings that you specified for the action, and then click **Configure**.

Policy Notifications

You can define how a Data Protection Policy sends notifications in the Policy, Workflow, and Action resources.

The following table summarizes how the notification settings in each resource work together.

In the Policy resource, the following notification choices are available:

- **Never**—Select this option when you do not want to send any notifications.
- **On Completion**—Select this option when you want to send a notification on completion of the workflows and actions in the policy.
- **On Failure**—Select this option when you want to send a notification only if one or more of the workflows in the policy fail.

When you configure a notification at the policy level, NetWorker applies the notification to all workflows and actions in the policy that are not configured to send out notifications.

In the Workflow resource, the following notification choices are available:

- To use the notification configuration that is defined in the policy resource to send the notification, select **Set a policy level**.
- To send a workflow notification on completion of all the actions in the workflow, select **On Completion**.
- To send a workflow notification only if an action fails to complete, select **On Failure**.

When you configure a notification at the workflow level, the setting overrides what you defined at the policy level.

In the Action resource, the following notification choices are available:

- To use the notification configuration that is defined in the policy resource to send the notification, select **Set a policy level**.
- To send a notification on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

When you configure a notification at the action level, the setting overrides what you defined at the policy level. If you configured the Workflow resource to send out notifications, you will receive workflow notifications in addition to action notifications.

Monitoring policy activity

The **Monitoring** window in the **NetWorker Administration** window enables you to monitor activities for specific policies, workflows, and actions.

Policies/Actions pane

The **Policies/Actions** pane at the top of the **Monitoring** window lists the policies on the NetWorker server by default. Click the + (plus) sign next to a policy in the list to

view the workflows in the policy, and the + (plus) sign next to a workflow to view the actions for a workflow.

The **Policies** pane provides the following information for each item (where applicable):

- Overall status

The following table provides details on the status icons that may appear in the **Policies** pane.

Table 45 Policy status icons

Icon	Status
	Never run
	Running
	Succeeded
	Failed
	Probing

- Most recent start time
- Duration of the most recent run
- Next scheduled runtime
- Name of the assigned save set
- Device on which the save set is stored
- Backup level
- Data transfer rate
- Size of the save set
- Messages that resulted from an action

Right-click an action in the **Policies** pane, and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

When you sort the items on the **Policy/Actions** pane by using the **Status** column, NetWorker sorts the items in alphabetical order that is based on the label of the icon.

Consider the following when a policy/action is in a probing state:

- A message is sent when the group starts and finishes the probe operation.
- The results of the probe operation (run backup/do not run backup) are also logged.
- Probes do not affect the final status of the group, and the group status does not indicate the results of the probe.
- If probing indicates that a backup should not run, then the group status reverts to its state before the group running.
- Check the results of the probe in the **Log** window to ensure that the probe indicates that the backup can be taken.

Actions pane

To view a list of all actions, click the **Actions** tab at the bottom of the **Policies** pane. The **Policies** pane becomes the **Actions** pane.

The **Actions** pane provides the following information for each action:

- Overall status



Note: The **Actions** pane displays the same status icons as the **Policies** pane.

- Name
- Assigned policy
- Assigned workflow
- Type
- Date and time of the most recent run
- Duration of the most recent run
- Percent complete, for actions that are in progress
- Next scheduled runtime

Right-click an action in the **Actions** pane, and select **Show Details** to view details on currently running, completed, and failed activities for the action.

Monitoring cloning

You can view the status of scheduled clone jobs in the **Monitoring** window. Status information includes the last start time of the policy, workflow, or clone action, the duration of the action, the size of the save set, and the target device, pool, and volume.

To determine whether a save set on a volume has been cloned, or is itself a clone, check the search for the save set by using the **Query Save Set** tab when you select **Save Sets** in the **Media** window.

Policy log files

The NetWorker server contains the log files for all data protection Policy resources.

Policy log directory structure

The policy-related resource log files are found in the following directory:

- **Windows:**
C:\Program Files\EMC NetWorker\nsr\logs*policy_name*
 workflow_name\action_name
- **Linux:**
/nsr/logs/*policy_name*/*workflow_name*/*action_name*

where:

- *Policy_name*—is the name of the Policy resource. One folder per policy.
- *Workflow_name*—is the name of the workflow directory. One folder per action sequence.
- *Action_name*—is the name of the action log file within the workflow.

Workflow log files

The policy subdirectory contains raw log files for each workflow and one subdirectory for each action.

The location and format of the log file on Linux is:

```
/nsr/logs/policy/policy_name/workflow_name_jobid.raw
```

where *name_jobid* is the name of the workflow and the job id of the workflow. Job id is a value that uniquely identifies a workflow job record in the jobdb.

For example, the log file for a workflow that is called server backup, with a job id of 0010072 appears as follows:

```
/nsr/logs/policy/server protection/workflow_server  
backup_0010072.raw
```

Use the job id to perform queries of the jobdb with the `jobquery` command. A workflow log file can be unrendered or rendered. An unrendered log file has the file name extension `.raw`. A rendered log file's extension is `.log`. Unrendered log files contain internationalized messages that can be rendered into the local language. The content of rendered log files has been localized to a single country's language.

[View log files](#) provides more information about viewing rendered and unrendered log files.

Action log files

NetWorker creates a workflow directory for each workflow within the policy directory. The workflow directory contains log files for each action that is assigned to the workflow.

The location of the workflow directory on Linux is:

```
/nsr/logs/policy/policy_name/workflow_name
```

where:

- *policy_name*—is the name of the policy that contains the workflow.
- *workflow_name*—is the name of the workflow.

The workflow directory contains log files for each action that is assigned to the workflow. The file name appears in the following format:

```
action_name_job_id.raw
```

where:

- *action_name*—is the name of the action.
- *job_id*—is the job id of the action in the jobdb.

For example, the server backup workflow has three actions: Backup, Clone, and Clone more. There are three log files in `/nsr/logs/policy/server protection/server backup` directory with the following names:

```
Backup_1408063.raw
```

```
Clone_1408080.raw
```

```
Clone more_1408200.raw
```

Child action log files

Some actions create child actions, for example a backup action creates a save job and a savefs job. Each child action has a unique job record.

Each of these child jobs have a log file. When the parent action starts a child action, NetWorker creates a directory for the action that contains the log file for child activities.

The location of the action directory on Linux is:

```
/nsr/logs/policy/policy_name/workflow_name/  
action_name_job_id_logs
```

where:

- *policy_name*— is the name of the policy that contains the workflow.
- *workflow_name*— is the name of the workflow.
- *action_name*—is the name of the action.
- *job_id*—is the job id of the action in the jobdb.

The action directory contains log files for each child action started by the action. The file name appears in the following format:

job_id.log

where *job_id* is the job id of the child action in the jobdb.

For example, an action whose log file name is `Backup_1408063.raw` might have a directory that is named `Backup_1408063_logs`, which contains three log files:

- 1408066.log
- 1408067.log
- 1408070.log

 **Note:** The .log files are localized to a specific country or the language of the region.

NetWorker clears the information about a job from the jobsdb and deletes the associated log files at the interval that is defined by the **Jobsdb retention in hours** attribute in the properties of the NetWorker Server resource. In NetWorker 9.0.1, the default jobsdb retention is 72 hours.

Starting, stopping, and restarting policies

The workflows in a policy can run automatically, based on a schedule. You can also manually start, stop, and restart specific workflows by using the the NMC **NetWorker Administration Monitoring** window.

About this task

You can restart any failed or canceled workflow. Note, however, that the restart must occur within the restart window that you specified for the workflow. Additionally, for a VMware backup, if you cancel a workflow from **NetWorker Administration** and then want to restart the backup, ensure that you restart the workflow from the **NetWorker Administration** window. If a workflow that was started from **NetWorker Administration** is restarted from the **vSphere Web Client**, the backup fails.

Procedure

1. In the **Monitoring** window, select the workflow or actions.
2. Right-click and then select **Start**, **Stop**, or **Restart**.

A confirmation message appears.

 **Note:** You cannot stop, restart, or start individual actions.

3. Click **Yes**.

Protection period

You can use protection period to retain the last valid copy of a saveset for a specified period of time even after it expires and the protection period is more than the retention policy. This is a nsr data protection attribute and disabled by default. Protection period is disabled by default. In order to enable this feature, you can check the enable check box in the policy configuration tab.

Note: The protection period applies to scheduled backup save sets only, and it does not apply to manual backup save sets. Some NetWorker module backups might appear to be scheduled backups that are initiated by a policy backup action but they are manual backups because they are initiated or converted by a database or application. The *NetWorker Module for Databases and Applications Administration Guide* and the *NetWorker Module for SAP Administration Guide* provides more details.

Enabling protection period in CLI

The following instructions describe how to enable protection period in CLI:

About this task

Use `-P` while creating or updating the data protection policy in CLI.

```
nsrpolicy create -p <policy_name> -q <yes> -P <protection_period>
```

Enabling protection period in NMC

The following instructions describe how to enable protection period in NMC:

Procedure

1. Create a new policy or open an already existing policy.

Figure 28 Creating a new policy

2. Click **Enable Protection Period** from NMC.
3. Set the protection period at a policy level.

Figure 29 Policy properties

Policy Properties - Platinum

General | Restricted Data Zones

Identity

Name: Platinum

Comment: Template policy: Platinum

Notifications

Send notification: On Completion

Command: nsrlog -f policy_notifications.log

Protection Period

☒ Enable Protection Period

Protection Period: 9 Months

OK Reset Cancel

Starting actions in a workflow for an individual client

When you start a workflow, NetWorker performs all the actions in the workflow for all the clients that are defined in the groups that are associated with the workflow. You can also start the actions for specific clients in a workflow.

About this task

Perform the following steps to start the actions for an individual client.

Note: You cannot start the actions for specific clients in the Server backup workflow.

Procedure

1. From the **Administration** window, click **Monitoring**.
2. In the Policies pane, expand the policy.
3. Right-click the workflow, and select **Start Individual Client**. The **Start Workflow** dialog box appears.
4. Optionally, from the **Workflow** list, select a different workflow.
5. Select the checkbox next to the names of the clients on which you want to perform all the actions in the workflow.
6. Click **Start**.

Modifying data protection Policy resources

This section describes how to modify existing Policy, Workflow, Group, and Action resources.

Policies

Policies enable you to manage all data protection tasks and the data protection lifecycle from a central location.

A policy contains one or more workflows, which define the actions that should be performed, the order for the actions to occur, and the group of Client resources or save sets on which to perform the actions.

Actions include backups, cloning, client/server connectivity checks, and NetWorker server maintenance activities.

Editing a policy

You can edit the description, notification setting, and RDZ for a policy.

About this task

You cannot edit the name of a policy. To rename a policy, first delete the policy, and then re-create it with the new name.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Policies**.
3. Right-click the policy, and select **Properties**.
The **Policy Properties** dialog box appears.
4. Edit the properties for the policy. The properties are the same properties that you specified when you created the policy.
5. Click **OK**.

Deleting a policy

When you delete a policy, the deletion process also deletes all workflows and actions for the policy.

About this task

Groups that are assigned to the workflows in the policy are not deleted, however. The workflow assignment for the group is removed from the group properties. You can assign the group to a workflow in a different policy, or delete the group.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Policies**.
3. Right-click the policy, and select **Delete**.
A confirmation message appears.
4. Click **Yes**.

Note: The Policy resource remains in the **Monitoring** window until all the information about the workflows and actions within the policy expire in the jobs database. The default job expiration time is 72 hours. [Modifying the retention period for jobs in the jobs database](#) describes how to change the default job expiration time.

Workflows

Workflows define a list of actions to perform sequentially or concurrently, a schedule window during which the workflow can run, and the protection group to which the workflow applies.

A workflow can be as simple as a single action that applies to a finite list of Client resources, or it can be a complex chain of actions that apply to a dynamically changing list of resources, with some actions occurring sequentially and others occurring concurrently.

You can also define notification settings for a workflow.

Supported workflow paths


Workflows enable you to chain together multiple actions and run them either sequentially or concurrently. However, the sequence of actions in a workflow is limited by certain logical constraints.

The following sections provide details on supported actions that can follow the lead action in a workflow.

Workflow path from a snapshot backup action

You can perform a generate index action (to generate an index of the snapshot) or a clone action after a snapshot backup action.


Figure 30 Workflow path from a snapshot backup action



Workflow path from a probe action


You can perform either a traditional backup or a snapshot backup after a probe action.

Figure 31 Workflow path from a probe action




Workflow path from a server backup action


A clone action is the only supported action after a server backup action.

Figure 32 Workflow path from a server backup action**Workflow path from a check connectivity action**

You can perform a traditional backup, snapshot backup, or probe action after a check connectivity action.

Figure 33 Workflow path from a check connectivity action**Workflow path from a clone action**


Another clone action is the only supported action after a clone action.

Figure 34 Workflow path from a clone action**Workflow path from an expire action**


The expire action must be the only action in a workflow. No other actions are supported either before or after an expire action.

Workflow path from a discover action

You can perform a generate index or clone action after a discover action.

Figure 35 Workflow path from a discover action**Workflow path from a generate index action**


The only supported action after a generate index action is a clone action.

Figure 36 Workflow path from a generate index action

Workflow path from a VBA checkpoint discover action

The only supported action after a VBA checkpoint discover action is a VBA checkpoint backup action.

Figure 37 Workflow path from a VBA checkpoint discover action



Workflow path from a VBA checkpoint backup action

VBA checkpoint backup cannot be the lead action in a workflow. You must precede the VBA checkpoint backup action with a VBA checkpoint discover action.

Visual representation of traditional backup workflows

Figure 38 Traditional backup workflow



After you create actions for a workflow, in the Administration interface, you can see a map provides a visual representation of the actions on the right side of the **Protection** window.

The oval icon specifies the group to which the workflow applies. The rounded rectangle icons identify actions. The parallelogram icons identify the destination pool for the action.

You can work directly in the visual representation of a workflow to perform the following tasks:

- You can adjust the display of the visual representation by right-clicking and selecting one of the following options:
 - **Zoom In**—Increase the size of the visual representation.
 - **Zoom Out**—Decrease the size of the visual representation.
 - **Zoom Area**—Limit the display to a single section of the visual representation.
 - **Fit Content**—Fit the visual representation to the window area.
 - **Reset**—Reset the visual representation to the default settings.
 - **Overview**—View a separate dialog box with a high-level view of the visual representation and a legend of the icons.
- You can view and edit the properties for the group, action, or destination pool by right-clicking the icon for the item, and then select **Properties**.
- You can create a group, action, or destination pool by right-clicking the icon for the item, and then select **New**.

Create a workflow for an existing policy in NetWorker Administration

A policy can contain one or more unique workflows.

Before you begin

- Create a policy for the workflow.
- (Optional but recommended) Create a group of client resources or save sets to assign to the workflow.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left pane, select **Policies**.
3. Select the policy for the workflow.
4. In the right pane of the window, select the **Workflows** tab.
5. Right-click an empty area of the **Workflows** tab and select **New**.

The **New Workflow** dialog box appears.

6. In the **Name** field, type the name of the workflow.

The maximum number of allowed characters for the **Name** field is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: / \ * ? [] () \$! ^ ; ' " ` ~ > < & | { }

7. In the **Comment** box, type a description for the workflow.

The maximum number of allowed characters for the **Comment** field is 128.

8. From the **Send Notifications** list, select how to send notifications for the workflow:
 - To use the notification configuration that is defined in the policy resource to specify when to send a notification, select **Set at policy level**.
 - To send notifications with information about each successful and failed workflow and action, after the workflow completes all the actions, select **On Completion**.
 - To send notifications with information about each failed workflow and action, after the workflow completes all the actions, select **On Failure**.
9. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:

```
nsrlog -f policy_notifications.log
```

- On Linux, to send an email notification, type the following command:

```
mail -s subject recipient
```

- On Windows, type the following command: `smtpmail -s subject -h mailserver recipient1@mailserver recipient2@mailserver...` where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the

`smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.

- `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.
- `recipient1@mailserver`—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

10. In the **Running** section, perform the following steps to specify when and how often the workflow runs:
 - a. To ensure that the actions that are contained in the workflow run when the policy or workflow starts, in the **Enabled** box, leave the option selected. To prevent the actions in the workflow from running when the policy or workflow that contains the action starts, clear this option.
 - b. To start the workflow at the time that is specified in the **Start time** attribute, on the days that are defined in the action resource, in the **AutoStart Enabled** box, leave the option selected. To prevent the workflow from starting at the time that is specified in the **Start time** attribute, clear this option.
 - c. To specify the time to start the actions in the workflow, in the **Start Time** attribute, use the spin boxes.
The default value is 9:00 PM.
 - d. To specify how frequently to run the actions that are defined in the workflow over a 24-hour period, use the **Interval** attribute spin boxes. If you are performing transaction log backup as part of application-consistent protection, you must specify a value for this attribute in order for incremental transaction log backup of SQL databases to occur.
The default **Interval** attribute value is 24 hours, or once a day. When you select a value that is less than 24 hours, the **Interval End** attribute appears. To specify the last start time in a defined interval period, use the spin boxes.
 - e. To specify the duration of time in which NetWorker can manually or automatically restart a failed or canceled workflow, in the **Restart Window** attribute, use the spin boxes.
If the restart window has elapsed, NetWorker considers the restart as a new run of the workflow. NetWorker calculates the restart window from the start of the last incomplete workflow. The default value is 24 hours.
For example, if the **Start Time** is 7:00 PM, the **Interval** is 1 hour, and the **Interval End** is 11:00 PM., then the workflow automatically starts every hour beginning at 7:00 PM. and the last start time is 11:00 PM.
11. In the **Groups** group box, specify the protection group to which the workflow applies.
To use a group, select a protection group from the **Groups** list. To create a protection group, click the + button that is located to the right of the **Groups** list.
12. The **Actions** table displays a list of actions in the workflow. To edit or delete an action in the workflow, select the action and click **Edit** or **Delete**. To create one or more actions for the workflow, click **Add**.
The **Actions** table organizes the information in sortable columns. Right-click in the table to customize the attributes that appear.

13. To create the workflow, click **OK**.

Create a workflow for a new policy in NetWorker Administration

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the left pane, expand **Policies**, and then select the policy that you created.
3. In the right pane, select **Create a new workflow**.
4. In the **Name** field, type the name of the workflow.

The maximum number of allowed characters for the **Name** field is 64.

- Legal Characters: `_ : - + = # , . % @`
- Illegal Characters: `/ \ * ? [ ] ( ) $ ! ^ ; ' " ` ~ > < & | { }`

5. In the **Comment** box, type a description for the workflow.

The maximum number of allowed characters for the **Comment** field is 128.

6. From the **Send Notifications** list, select how to send notifications for the workflow:
 - To use the notification configuration that is defined in the policy resource to specify when to send a notification, select **Set at policy level**.
 - To send notifications with information about each successful and failed workflow and action, after the workflow completes all the actions, select **On Completion**.
 - To send notifications with information about each failed workflow and action, after the workflow completes all the actions, select **On Failure**.
7. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages, or use the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:


```
nsrlog -f policy_notifications.log
```
- On Linux, to send an email notification, type the following command:


```
mail -s subject recipient
```
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:


```
/usr/sbin/sendmail -v recipient_email "subject_text"
```
- On Windows, type the following command:


```
smtpmail -s subject -h mailserver recipient1@mailserver
recipient2@mailserver...
```

where:

- *-s subject*—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- *-h mailserver*—Specifies the hostname of the mail server to use to relay the SMTP email message.
- *recipient1@mailserver*—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

8. In the **Running** section, perform the following steps to specify when and how often the workflow runs:
 - a. To ensure that the actions that are contained in the workflow run when the policy or workflow starts, in the **Enabled** box, leave the option selected. To prevent the actions in the workflow from running when the policy or workflow that contains the action starts, clear this option.
 - b. To start the workflow at the time that is specified in the **Start time** attribute, on the days that are defined in the action resource, in the **AutoStart Enabled** box, leave the option selected. To prevent the workflow from starting at the time that is specified in the **Start time** attribute, clear this option.
 - c. To specify the time to start the actions in the workflow, in the **Start Time** attribute, use the spin boxes.
The default value is 9:00 PM.
 - d. To specify how frequently to run the actions that are defined in the workflow over a 24-hour period, use the **Interval** attribute spin boxes. If you are performing transaction log backup as part of application-consistent protection, you must specify a value for this attribute in order for incremental transaction log backup of SQL databases to occur.
The default **Interval** attribute value is 24 hours, or once a day. When you select a value that is less than 24 hours, the **Interval End** attribute appears. To specify the last start time in a defined interval period, use the spin boxes.
 - e. To specify the duration of time in which NetWorker can manually or automatically restart a failed or canceled workflow, in the **Restart Window** attribute, use the spin boxes.
If the restart window has elapsed, NetWorker considers the restart as a new run of the workflow. NetWorker calculates the restart window from the start of the last incomplete workflow. The default value is 24 hours.
For example, if the **Start Time** is 7:00 PM, the **Interval** is 1 hour, and the **Interval End** is 11:00 PM., then the workflow automatically starts every hour beginning at 7:00 PM. and the last start time is 11:00 PM.
9. To create the workflow, click **OK**.

After you finish

Create the actions that will occur in the workflow, and then assign a group to the workflow. If a workflow does not contain a group, a policy does not perform any actions.

Editing a workflow

You can edit all the properties for a workflow, including the name, description, schedule, notification settings, group, and actions.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Policies**.
3. Select the policy for the workflow.
4. In the right pane of the window, select the **Workflows** tab.
5. In the right pane, perform one of the following tasks:
 - To modify multiple attributes in a single configuration resource by using the **Workflow Properties** window, right-click the staging configuration and select **Properties**.
 - To modify a specific attribute that appears in the resource window, place the mouse in the cell that contains the attribute that you want to change, then right-click. The menu displays an option to edit the attribute. For example, to modify the **Comment** attribute, right-click the resource in the **Comment** cell and select **Edit Comment**.
6. Edit the properties for the workflow. The properties are the same properties that you specified when you created the workflow.

Note: To modify a specific attribute for multiple resources, press and hold the **Ctrl** key, select each resource, and then right-click in the cell that contains the attribute that you want to change. The menu displays an option to edit the attribute.
7. Click **OK**.

Deleting a workflow

When you delete a workflow, the deletion process also deletes all actions for the workflow.

About this task

The group that is assigned to the workflow is not deleted, however. The workflow assignment for the group is removed from the group properties. You can assign the group to a different workflow or delete the group.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Policies**.
3. Select the policy for the workflow.
4. In the right pane of the window, select the **Workflows** tab.

5. Right-click the workflow, and select **Delete**.

A confirmation message appears.

6. Click **Yes**.

Protection groups

Protection groups enable you to define a set of Client resources or save sets.

Assigning a protection group to a workflow

You can assign a protection group to a workflow either when you create or edit the group, or when you create or edit the workflow.

About this task

Each workflow applies to only one protection group, and each protection group can be assigned to only one workflow.

Procedure

- To assign a protection group to a workflow when you create or edit the group, select the workflow from the **Workflow(Policy)** list in the **Create Group** or **Edit Group** dialog box.
- To assign a protection group to a workflow when you create or edit the workflow, select the group from the **Groups** list in the **New Workflow** or **Workflow Properties** dialog box.

Editing a protection group

You can edit all properties for a protection group except for the group name and group type.

About this task

To rename a protection group, first delete the group, and then re-create it with the new name.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Groups**.
3. Right-click the group, and select **Properties**.

The **Edit Group** dialog box appears.

4. Edit the properties for the protection group.

The properties are the same properties that you specified when you created the group. To modify the clients in a protection group, perform one of the following tasks:

- To modify the clients in a dynamic group, in the **Dynamic clients** table, specify the criteria that NetWorker uses to select clients for the group:
 - To back up all the Client resources that are configured on the NetWorker server and have the **Scheduled backup** attribute enabled, select **Choose all clients**.
 - To generate a list of clients that is based on the value that is defined in the **Tag** attribute of the Client resource, select the **Clients with these tags** option. Specify the matching tag value in the **Tags** field and specify one tag on each line.

Note: When you specify multiple tag values, the query uses an OR operation to match the tags. For example, if you specify Sales and Support tag values, then the query builds a list of clients that contain the tag Sales or Support.

- To modify the clients in a Client group, from the **Clients** table, perform one of the following actions in the **Selected Clients** column:
 - To add a Client resource to the group, select the checkbox beside the name of the Client resource.
 - To remove Client resources from the group, clear the checkbox next to the name of the Client resource.

5. Click **OK**.

Deleting a protection group

Before you begin

Delete the workflow that is assigned to the protection group, or assign the workflow to a different protection group. You cannot delete a protection group if it is assigned to a workflow.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Groups**.
3. Right-click the group, and select **Delete**.

A confirmation message appears.

4. Click **Yes**.

Actions

Actions are the key resources in a workflow for a data protection policy. An action is a task that occurs on a work list. A work list is a list of pending work items, such a group of Client resources or save sets.

You can chain multiple actions together to occur sequentially or concurrently in a workflow.

Creating an action

The **Policy Action** wizard walks you through the steps to create an action. You can create an action either when you are creating or editing a workflow, or as a separate process from the workflow configuration.

Before you begin

Create the policy and workflow that contains the action.

Procedure

1. Open the **Policy Action** wizard by using one of the methods in the following table.

Table 46 Methods to create an action

Method	Steps
To create an action during the workflow configuration	Click Add in either the New Workflow dialog box or the Workflow Properties dialog box.
To add additional actions after the last action in an existing workflow	- In the Administration window, click Protection. - In the expanded left pane select Policies. - Select the policy. - Select the workflow. - In the right pane, select the Actions tab. - Right-click an empty area of the Actions tab and select New. Note: When you add actions to an existing workflow that is associated with a group, you only see the action types that are allowed in the action sequence.
To create the first action in a workflow	- In the Administration window, click Protection. - In the expanded left pane select Policies. - Select the policy. - Select the workflow. - In the right pane, select Create a new action.
To add an action before an action in an existing workflow	- In the Administration window, click Protection. - In the expanded left pane select Policies. - Select the policy. - Select the workflow. - In the right pane, select the action that you want the new action to precede and select Insert before. Note: When you add actions to an existing workflow that is associated with a group, you only see the action types that are allowed in the action sequence.

- In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: \ * ? [] () \$! ^ ; ' " ` ~ > < & | { }

- In the **Comment** field, type a description for the action.

4. To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.



Note: When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

5. From the **Action Type** list, select the action.
6. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
7. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
8. To select a predefined schedule, click **Select** radio button, then select the schedule from the drop down menu. If you do not want to associate a schedule to an action, then click **Define**.
9. The steps to go through the wizard depend on the action type that you select.

Editing an action

You can edit all the properties of an existing action.

About this task

Perform one of the following tasks to edit an action.

Procedure

- Open the **Policy Action** wizard for the action by using one of the methods in the following table.

Table 47 Methods to open the Policy Action wizard

Method	Steps
During workflow configuration	Select the action and then click Edit in either the New Workflow dialog box or the Workflow Properties dialog box.
From the Actions tab of the workflow	1. In the Administration window, click Protection. 2. In the expanded left pane select Policies. 3. Select the policy. 4. Select the workflow. 5. In the right pane, select the Actions tab. 6. Right-click the action, and select Properties.
From the visual representation of the workflow	Right-click the action in the visual representation of the workflow, and select Properties .

Edit the properties for the action, then click **Configure**.

- Use the quick edit option in the Actions window of a Workflow resource. To modify a specific attribute that appears in the resource window, place the mouse in the

cell that contains the attribute that you want to change, then right-click. The menu displays an option to edit the attribute. For example, to modify the **Comment** attribute, right-click the resource in the **Comment** cell and select **Edit Comment**.

 **Note:** To modify a specific attribute for multiple resources, press and hold the **Ctrl** key, select each resource, and then right-click in the cell that contains the attribute that you want to change. The menu displays an option to edit the attribute.

Deleting an action

You can delete an action in a workflow either when you are creating or editing a workflow, or as a separate process from the workflow configuration.

About this task

If the action that you delete is part of a sequence of actions in a workflow, then you can only delete the action if the removal of the action from the sequence would still result in a valid workflow. The properties for other actions in a sequence are updated to reflect the new sequence of actions after the deletion.

Procedure

- To delete an action when you are creating or editing a workflow:
 - a. Select the action in either the **New Workflow** dialog box or the **Workflow Properties** dialog box.
 - b. Click **Delete**.

A confirmation message appears.
 - c. Click **Yes**.
- To delete an action as a separate process from workflow configuration:
 - a. In the **Administration** window, click **Protection**.
 - b. In the expanded left pane, select **Policies**.
 - c. Select the policy.
 - d. Select the workflow.
 - e. In the right pane, select the **Actions** tab.
 - f. Right-click the action and select **Delete**.

A confirmation message appears.
 - g. Click **Yes**.

Configuring NSR Protection Policies from nsradmin

The `nsradmin` command is a command-line based administrative program for the NetWorker. Normally `nsradmin` monitors and modifies NetWorker resources over the network. Commands are entered on standard input, and output is produced on standard output.

If `nsradmin` command is started without a query argument, it uses a default query. By default, if the daemon being administered is `nsrd`, then all resources will be selected, but for all other daemons, no resources will be selected.

Create, Retrieve, Update, and Delete operations are supported through `nsradmin` for following resources:

- NSR protection policy
- NSR protection policy workflow
- NSR protection policy action

The NSR Protection Policy resource contains all related (nested) Workflows and Actions. CRUD operations can be carried out through directly executing on `nsradmin` prompt or specifying in an input file.

- `nsradmin`
- `nsradmin -i input_file.txt`

 **Note:** `nsradmin` visual mode is not supported for Create, Retrieve, Update, and Delete operations for NSR data protection policy, workflow, and action.

Commands

At each input prompt, `nsradmin` command expects a command name and some optional arguments. Command names can be shortened to the smallest unique string (for example, `p` for `print`). Command arguments are always specified in the form of an attribute list.

Table 48 Command descriptions

Command	Description
create or c	Create a resource with the given attributes. One of the attributes must be type to specify a NetWorker Type resource that can be created. The types command can be used to find out which NetWorker resource types a server supports.
delete or d	Delete the resources that match the current query.
print or p	Print the resources that match the current query. If a query is specified, it becomes the current query. If a name has been specified for the current show list, only the attributes for the specified name in the show list is displayed.
update or up	Update the resources given by the current query to match attributes. Set current query for print, update, or delete actions.

 **Note:** Comma separated inputs for resource names are supported for print, update, and delete operations only.

The number of characters allowed for policy, action, and workflow are 64.

- Legal Characters: `_ : - + = # , . % @`
- Illegal Characters: `/\*?[]()$!^;'""`~><&|{}`

Managing NSR Data Protection Policy

A NSR Data protection policy is a container for workflows. Each policy has one or more workflows. A workflow contains a set of actions and a list of data sources to run those actions.

The following table lists `nsradmin` commands to manage NSR Protection Policy resources:

Table 49 Commands to manage NSR Protection Policy resources

Actions	Example Commands
To create policy with default settings	<code>nsradmin> create type: NSR Protection Policy; name: Policy1</code>
To create policy with specific settings	<code>nsradmin> create type: NSR Protection Policy; name: Policy1; Notification execute on: completion; Notification action:"nsrlog -f policy_notifications_latest.log"; policy protection period:3 Months</code>
To print (display) a policy	<code>nsradmin> print type: NSR Protection Policy; NSR Protection Policy Workflow; NSR Protection Policy Action; name: Policy1</code>
To print (display) all policies with specific attribute value	<code>nsradmin> . type: NSR Protection Policy; Notification action: "nsrlog -f policy_notifications_updated.log" nsradmin> print</code>
To update a policy	<code>nsradmin> . type: NSR Protection Policy; name: Policy1 nsradmin> update Notification action: "nsrlog -f policy_notifications_updated.log"</code>
To update all policies that match specific comment value	<code>nsradmin> . type: NSR Protection Policy; comment: Gold policy nsradmin> update Notification action:"nsrlog -f policy_notifications_updated.log"</code>
To delete a policy	<code>nsradmin> delete type: NSR Protection Policy; name: Policy1</code>
To delete multiple policies with comma separated names	<code>nsradmin> delete type: NSR Protection Policy; name: Policy1, Policy3, Policy3</code>
To set current query to all policies with specific comment attribute value	<code>nsradmin>. type: NSR Protection Policy; comment: Backup</code>

Managing Workflows

The following table lists the `nsradmin` commands to manage NSR Protection Policy Workflow resources:

Table 50 Commands to manage NSR Protection Policy Workflow resources

Actions	Example Commands
To create workflow with default settings	<code>nsradmin> create type: NSR Protection Policy Workflow; name: Policy1-Workflow1; policy name: Policy1</code>
To create workflow with specific settings	<code>nsradmin> create type: NSR Protection Policy Workflow; name: Policy1-Workflow1; policy name:</code>

Table 50 Commands to manage NSR Protection Policy Workflow resources (continued)

Actions	Example Commands
	<code>Policy1; Auto Start enabled: yes; Start time:"08:00"</code>
To print (display) workflow with query name	<code>nsradmin> . type: NSR Protection Policy Workflow; name: Policy1-Workflow1 nsradmin> print</code>
To print all workflows	<code>nsradmin> print type: NSR Protection Policy Workflow</code>
To update query set workflow	<code>nsradmin> . type: NSR Protection Policy Workflow; policy name: Policy1; name: Policy1-Workflow1 nsradmin> update Notification action: "nsrlog -f workflow_notifications_SAP.log"</code>
To delete workflow	<code>nsradmin> delete type: NSR Protection Policy Workflow; policy name: Policy1; name: Policy1-Workflow1</code>
To delete all workflows with query set as name	<code>nsradmin> delete type: NSR Protection Policy Workflow; name: Policy1-Workflow1</code>
To delete multiple workflows	<code>nsradmin> delete type: NSR Protection Policy Workflow; name: policy1-Workflow1; name: Policy1-Workflow2; name: Policy1-Workflow3</code>
To delete all workflows from multiple policies, policy names given as comma separated	<code>nsradmin> delete type: NSR Protection Policy Workflow; policy name: policy1, Policy3, Policy3</code>
To delete all workflows from query set policy	<code>nsradmin> delete type: NSR Protection Policy Workflow; name: Policy1-Workflow1, Policy1-Workflow2, Policy1-Workflow3</code>
To set current query to a workflow	<code>nsradmin> . type: NSR Protection Policy Workflow; policy name: Policy1; name: Policy1-Workflow1</code>

Managing Actions

The following table lists `nsradmin` commands that manage NSR Protection Policy Action resources:

Table 51 Commands to manage NSR Protection Policy Action resources

Actions	Example Commands
To create action with default settings	<code>nsradmin> create type: NSR Protection Policy Action; policy name: Policy1; Workflow name: Policy1-Workflow1; name: BackupAction1</code>
To create action with specific settings	<code>nsradmin> create type: NSR Protection Policy Action; policy name: Policy1; Workflow name: Policy1-Workflow1; name: BackupAction1; action type:backup; backup subtype:traditional; period:week;</code>

Table 51 Commands to manage NSR Protection Policy Action resources (continued)

Actions	Example Commands
	<pre>actions:"full,full,full,incr,incr,incr,incr"; destination storage node:"nsrserverhost"; destination pool: Data Domain Default;apply dd retention lock:yes; dd retention lock time:"16 Days"; client can override:no; overrides:"full third wednesday every month"; action retries:3; retry delay:10; action inactivity timeout:44; action parallelism:50; failure impact:"abort action"; soft limit:"21:23"; hard limit:"18:14"; start time:"14:14"</pre>
To create a clone action	<pre>create type:NSR Protection Policy Action;name:clone1;policy name:Policy1;workflow name:Policy1-Workflow1;driven by: BackupAction1;action type:clone;retention:"12 days";destination storage node:"StorageNode";source storage node:"nsrserverhost";destination pool:"default clone";exclude level:yes;filter level:"full";exclude client:yes;filter client:"client1";exclude saveset type:no;filter save set type:snapshot;exclude time range:yes;Filter time range start:"12:00";Filter time range end:"23:00"</pre>
To create action with subtype snapshot	<pre>nsradmin> create type: NSR Protection Policy Action; name: Snapshot1 ; policy name: Policy1; workflow name: Policy1-Workflow1; action type: backup; backup subtype: snapshot</pre>
To create a VMWARE action	<pre>create type:NSR Protection Policy Action; name:VMAction1;policy name:Policy1;workflow name:Policy1-Workflow1; action type:backup;backup subtype:vmware; destination storage node:"StorageNode";retention:"12 days";vproxy name:"vproxy"</pre> Note: The following operations are supported only via the NMC and NWUI for VMware backup action subtype: • Addition of the VMs to the NSR protection group • Creation of dynamic tags and policies • Update VMware action attributes to non-default values
To create action with type clone following action type backup	<pre>nsradmin> create type: NSR Protection Policy Action; policy name: Policy1; workflow name: Policy1- Workflow1; action type: clone; driven by: BackupAction1; name: Clone1; Retention: 1 Years; destination pool: Data Domain Default Clone</pre>
To create a dependent clone action driven by earlier clone action	<pre>nsradmin> create type: NSR Protection Policy Action; policy name: Policy1; workflow name: Policy1- Workflow1; action type: clone; driven by: Clone1; name: Clone2; Retention: 5 Years; destination pool: Data Domain Default Clone</pre>

Table 51 Commands to manage NSR Protection Policy Action resources (continued)

Actions	Example Commands
To print (display) query set action	nsradmin> print type: NSR Protection Policy Action; policy name: Policy1; workflow name: Policy1-Workflow1; name: BackupAction1
To update a query set action	nsradmin> . type: NSR Protection Policy Action; policy name: Policy1; workflow name: Policy1-Workflow1; name: BackupAction1 nsradmin> update retention: 6 Months; dd retention lock time:3 Months; destination storage node: networker; client can override: yes; comment "This is special action with 3 Months Retention Lock!" nsradmin> update period: month; actions: "full,incr,full,incr,full,incr,full,full,incr,full,incr,full,incr,full,full,incr,full,full,incr,full,incr,full,full,incr,full,incr,full,incr,full,full,incr,full,incr,full,incr,full,incr"
To delete a query set action	nsradmin> delete type: NSR Protection Policy Action; policy name: Policy1; workflow name: Policy1-Workflow1; name: BackupAction1
To delete all actions with query set name from all workflows	nsradmin> delete type: NSR Protection Policy Action; name: Snapshot1
To delete multiple actions from single workflow specifying action names	nsradmin> delete type: NSR Protection Policy Action; policy name: Policy1; workflow name: Policy1-Workflow1; name: Clone1; name: Clone2; name: BackupAction1
To delete multiple actions with comma separated names	nsradmin> delete type: NSR Protection Policy Action; name: BackupAction1, Clone1, Clone2
To delete all actions from query set workflow	nsradmin> delete type: NSR Protection Policy Action; policy name: Policy1; workflow name: Policy1-Workflow1
To set current query to an action	nsradmin> . type: NSR Protection Policy Action; policy name: Policy1; workflow name: Policy1-Workflow1; name: BackupAction1

Configuring Filters for Clone and Index Actions

The following filters can be applied on clone and index action from nsradmin:

- Time Range filter
- Saveset filter
- Client filter
- Level filter

Filters can be applied using Exclude option for each filter.

- **Exclude** - Do not filter (default value for attribute). No savesets are filtered based on the filter type (Time, Client, Saveset, Level).

- Exclude - Yes. The filter is applied and this option is equivalent to Reject option on NMC.
- Exclude - No. The filter is applied and this option is equivalent to Accept option on NMC.

Table 52 Filter values for clone and index actions

Filters	Values
Time Range Filter	Exclude Time Range: Yes/No/Do Not Filter(Default Value of the Attribute) Specify the time range in which NetWorker searches for eligible save sets to clone in the media database. Filter Time Range Start: Start time of the saveset in 24-hours format Filter Time Range End: End time of saveset in 24-hours format
Save set Filter	Exclude saveset type: Yes/No/Do not filter (Default value for attribute) Filter Save set type: Protectpoint, Snapshot
Client Filter	Exclude client: Yes/No/Do not filter (Default value for attribute) Filter Client: Comma separated list of clients
Level Filter	Exclude level: Yes/No/Do not filter (Default value for attribute) Filter level: full, incr, 1, txnlog

When the value for Exclude is set to Do not Filter, filters (Time, Saveset, Client, Level) attribute values are not applied. Any changes to filters do not take effect unless exclude level is set to Yes or No. When Exclude is updated to Do not filter, all filters attribute values are set to Null. When Exclude is set to Do not filter, filter values are not updated and processed and no error messages are printed.

Deleting multiple actions from multiple workflows and policies

The order in which action is deleted applies when all actions are specified or when all actions inside a policy or all actions inside a workflow are selected. The order of deletion is reverse of the order of “Driven By” chain.

- Deletion of actions when action name is not specified- The `nsradmin` command iterates through all the actions in a workflow in reverse order. When you specify a particular query, based on matching criteria, it displays all actions that have to be deleted. If the action is independent then it is automatically deleted. If an action becomes independent depending upon the previous delete and it matches the query, then it is deleted. Different attributes can behave differently among actions. Therefore, deletion differs. For example, destination pool differs for backup and clone and action is fetched accordingly. Random deletion of dependent actions fails when no action name is specified. It is recommended that you specify the action name in order to match the query criteria.
- Deletion of actions when action name is specified- The `nsradmin` command deletes the random dependent actions if it matches the query. When a random action is deleted in a workflow containing a set of actions, then the “driven by” attribute of the deleted action is updated in the “driven” attribute of the action succeeding it.

 **Note:** If the head action has more than one action dependent on it, then the head action cannot be deleted.

Querying NSR protection policy workflow and NSR protection policy action resources without specifying RAP resource types

NSR Protection Policy Workflow and NSR Protection Policy Action are pseudo RAP resources and encapsulated under NSR Protection Policy RAP resource. Querying these pseudo RAP resources in `nsradmin` requires resource type. The `nsradmin` command displays the following warning message for these pseudo resources, if queried without resource type:

```
nsradmin> . comment:"This is Backup."
Current query set
nsradmin> p
RAP resource type is required for print operation of NSR protection
policy Workflow and NSR protection policy Action resources.
type: NSR client;
name: mars;
comment: "This is Backup.";

type: NSR Protection Policy;
name: Policy1;
comment: "This is Backup.";

type: NSR pool;
name: Pool1;
comment: "This is Backup.";
```

All other RAP resources including NSR Protection Policy resource can be queried without specifying RAP resource Types.

Managing policies from the command prompt

The `nsrpolicy` command enables you to create, start, stop, and display the attribute of policy, workflow, action, and group resources.

The `nsrpolicy` command requires specific privileges which are assigned based on session authentication. NetWorker supports two types of session authentication. Token-based authentication, which requires you to run the `nsrlogin` before you run the command and authenticates the user that runs the command against entries that are defined in the External Roles attribute of a User Group resource. Classic authentication, which is based on user and host information and uses the user attribute of a User Group resource to authenticate a user. Classic authentication does not require an authentication token to run the command. For example, if you run the command without first running `nsrlogin`, NetWorker assigns the privileges to the user based on the entries that are specified in the Users attribute of the User Group resource. When you use `nsrlogin` to log in as a NetWorker Authentication Service user, NetWorker assigns the privileges to the user based on the entries that are specified in the External Roles attributes of the user Group resource. The *NetWorker Security Configuration Guide* provides more information about privileges.

This section provides some examples of how to manage data protection policies from a command prompt.

The UNIX man pages and the *NetWorker Command Reference Guide* provide detailed information about how to use the `nsrpolicy` command.

Creating Data Protection Policy resources from a command prompt

Use the `nsrpolicy` command to create Policy, Protection Group, Workflow and action resources.

Procedure

1. Optionally, use the `nsrlogin` command to authenticate a user and generate a token for the `nsrpolicy` command.

[Using nsrlogin for authentication and authorization](#) provides more information.

2. Use the `nsrpolicy` command to create each Data Protection Policy resource.

- a. To create the Policy resource, type: `nsrpolicy policy create -- policy_name policy_name`.

where *policy_name* is a unique name for the Policy resource.

- b. To create a protection Group resource and add existing clients to the Group resource, type: `nsrpolicy group create client -g group_name -C "client_name1,client_name2,client_name3..."`

where:

- *group_name* is a unique name of the Group resource.
- *client_name1,client_name2,client_name3...* is a comma separated list of client names to add to the group.

- c. To create a workflow and associate the workflow with the new Policy and Group resources, type: `nsrpolicy workflow create --policy_name policy_name --workflow_name workflow_name --group_name group_name`

where:

- *policy_name* is the name of the Policy resource.
- *group_name* is the name of the Group resource.
- *workflow_name* is a unique name for the Workflow resource.
- Legal Characters: `_ : - + = # , . % @`
- Illegal Characters: `\*?[]()$!^;"'`~><&|{}`

3. Use the `nsrpolicy display` command to display the attributes for the new Data Protection Policy resource.

- To display a Policy resource, type: `nsrpolicy action display -- policy_name policy_name`
Where *policy_name* is the name of the Policy resource.
- To display a Workflow resource, type: `nsrpolicy action display -- workflow_name workflow_name`
Where *workflow_name* is the name of the Workflow resource.
- To display a Group resource, type: `nsrpolicy action display -- group_name group_name`

Creating Action resources from a command prompt

Use the `nsrpolicy action create` command to create Action resources

Procedure

1. Optionally, use the `nsrlogin` command to authenticate a user and generate a token for the `nsrpolicy` command.
2. Use the `nsrpolicy action create` command to create the Action resource.

For example: `nsrpolicy action create action_type --policy_name policy_name --workflow_name workflow_name -A backup_action_name [-M "start_time"] [-d preceding_action_name]`

Where:

- *action_types* are one of the following: check-connectivity, probe, backup traditional, backup snapshot, clone, discover-nas-snap, index-nas-snap, server-backup, expire, vba-checkpoint-discover, vba-checkpoint-backup.
- *policy_name* is the name of an existing Policy resource that contains this action.
- *workflow_name* is the name of an existing Workflow resource in the Policy resource that contains the action.
- *action_name* is a unique name for the new Action resource.
 - Legal Characters: `_ : - + = # , . % @`
 - Illegal Characters: `/ \ * ? [ ] ( ) $ ! ^ ; ' " ` ~ > < & { | }`
- *start_time* is the time to start the action, in one of the following formats:
 - `-M "hh:mm"`—To start the action at a specific time. For example, to create a new action in an existing workflow that starts at 11:15 PM, type `-M "23:15"`
 - `-M "+hh:mm"`—To start the action after period of time has elapsed since the start of the workflow. For example, to create a new action that starts 3 hours after the start of a workflow, type `-M "+3:00"`
- *preceding_action_name* is the name of the Action that precedes the new action in the Workflow.

For example:

- To create a traditional backup action and add this action to the SQL workflow in the `SQL_hosts` policy resource, type: `nsrpolicy action create backup traditional --policy_name SQL_hosts --workflow_name SQL -A SQL_backup`.
- To create a clone action and insert the clone action immediately after a backup action created in the SQL workflow, type: `nsrpolicy action create backup traditional --policy_name policy_name SQL_hosts --workflow_name SQL -A SQL_clone -d SQL_backup`.
- To create a new action in an existing workflow that starts at 11:15 PM, type the following command:


```
nsrpolicy action create backup traditional -p policy_name -w
workflow_name -A action_name -M "23:25"
```

- To create a new action that starts 3 hours after the start of a workflow, type:

```
nsrpolicy action create backup traditional -p policy_name -w
workflow_name -A action_name -M "+3:00"
```

Starting, stopping, and restarting workflows from a command prompt

Use the `nsrpolicy` command to start, stop, and restart the actions in a workflow.

Starting a workflow from a command prompt

You can start all actions that are contained in one workflow in a policy, or start all actions for one client in a workflow.

- To start all actions in a specific workflow in a Policy resource, type the following command: `nsrpolicy start --policy_name "policy_name" --workflow_name "workflow_name"`
 **Note:** You cannot start another instance of a workflow that is already running.
- To start all actions for a specific client in a workflow, type the following command: `nsrpolicy start --policy_name "policy_name" --workflow_name "workflow_name" --client_list client_list`
 **Note:** You can use this command to start actions for failed clients in a workflow that is currently running.

where:

- "*policy_name*" is the name of the Policy resource that contains the workflow that you want to start.
- "*workflow_name*" is the name of the Workflow resource that you want to start.
- *client_list* is a comma-separated list of host names for the clients in the workflow whose actions you want to start.

Stopping all actions in a workflow from a command prompt

To stop all actions in a specific workflow in a policy, type the following command:

```
nsrpolicy stop --policy_name "policy_name" --workflow_name
"workflow_name"
```

where:

- "*policy_name*" is the name of the Policy resource that contains the workflow that you want to stop.
- "*workflow_name*" is the name of the Workflow resource that you want to stop.

Restarting a workflow from a command prompt

To restart all actions in a workflows that a Policy resource contains, type the following command: `nsrpolicy restart --policy_name "policy_name" --workflow_name "workflow_name"`

where:

- "*policy_name*" is the name of the Policy resource that contains the workflow that you want to restart.
- "*workflow_name*" is the name of the Workflow resource that you want to restart.

Running a workflow with action overrides

Before an action starts NetWorker defines how to run the action by reviewing the attributes values of the policy, workflow, and action resources. The `nsrworkflow` command line option `-A` enables you to override attribute values that NetWorker uses to run the action. Actions which support override values are: traditional and snapshot backups, probe, and clone.

Specify the `-A` option in the format `-A "action_name cmd_line_flags"`, where:

- *action_name*—Specifies the name of the action resource.
- *cmd_line_flags*—Defines a list of command line flags and the new parameter value. For more options, refer to `savegrp` command.

Use escaped double quotes or single quotes for action names or parameters that contain spaces or special characters. For example: `-A "\"action name\" -l full"` or `-A "'action name' -l full"`

For example, to specify an override on the level of a backup action and the retention time of the backup and clone actions in the workflow, type the following command:

```
nsrworkflow -p Backup -w workflow_name -A "action_name -l level -y \"retention_period\"" -A "action_name -y \"retention_period\""
```

To specify a backup level override of 3 and a retention period of 3 years for the backup and clone actions for a workflow named `fs_backup_clone`, an backup action named `backup` and a clone action named `clone`, type the following command:

```
nsrworkflow -p Backup -w fs_backup_clone -A "backup -l 3 -y \"3 years\"" -A "clone -y \"3 years\""
```

Running an adhoc workflow outside of the backup schedule

The `nsrworkflow -a` flag, which allows you to override a backup schedule and run an adhoc backup.

The following list describes workflow changes in adhoc mode:

- The action start time setting for all actions in a workflow (if set) is ignored.
- The action schedule activity of 'skip' is converted to the default schedule activity for the action. This conversion to the default action allows adhoc execution of workflows on days where the schedule is configured to level 'skip'. You can use the `-A` option to specify a different schedule activity if the action supports it. The following list describes the default schedule activity for various actions:
 - Database-level backup actions: 'incr'
 - Server-level backup actions: '1' (cumulative incremental)
 - All other actions: 'exec'

Note: There are different backup-level override flags for file-level backups and virtual machine backups:

- For virtual machine backups, the backup-level override flag is `-L`.
- For file-level backups, the backup-level override flag is `-l`.

The following example commands use the `nsrworkflow -a` flag, the `-A` flag to specify the schedule activity, and the backup-level override flag:

- Single saveset backup:

```
nsrworkflow -p <Policy> -w <Workflow> -A "backup -l full -c
10.63.101.77:<file_path>" -a
```

- Virtual machine backup: /

```
nsrworkflow -p "VM Backup" -w "2-Weekly VMCluster1" -c
"vm:client1,vm:client2" -A "'backup' -L full" -a
```

Displaying Data Protection Policy resource configurations

NetWorker stores Data Protection Policy resource configuration information in a JavaScript Object Notation (JSON) string. Displaying the contents of the JSON string provides you with the ability to view the hierarchical relationship between the resources.

Use the `nsrpolicy policy display` command to display the configuration attributes for a Policy resource and all the Workflow and Action resources that are associated with the Policy resource:

```
nsrpolicy policy display -p policy_name
```

where *policy_name* is the name of the Policy resource. Enclose Policy names that contain spaces in quotation marks.

For example, to display the resources in the Server Protection Policy resource, type the following command:

```
nsrpolicy policy display -p "Server Protection"
```

Output similar to the following appears

```
{
  "policyName": "Server Protection",
  "policyComment": "Default policy for server that includes
server backup and maintenance",
  "policySummaryNotification": {
    "policyCompletionNotificationAction": "nsrlog -f
policy_notifications.log",
    "policyCompletionNotificationExecuteOn": "completion"
  },
  "policyWorkflows": [
    {
      "workflowName": "Server backup",
      "synthesisRoot": [
        "NSR group/Server backup",
        "NSR Snapshot Policy/Server backup"
      ],
      "workflowActions": [
        {
          "actionName": "Server db backup",
          "actionSpecific": {
            "actions": {
              "actionType": "server backup",
              "asbDestinationPool": "Default",
              "asbDestinationStorageNode": "nsrserverhost",
              "asbPerformBootstrap": true,
              "asbPerformCFI": true,
              "asbRetentionPeriod": "1 Months"
            }
          }
        }
      ],
      "actionSchedulePeriod": "month",
```

```

"actionScheduleActivity": [
  "full",
  "1","1","1","1","1","1","1","1","1","1","1","1","1","1","1","1",
  "1","1","1","1","1","1","1","1","1","1","1","1","1","1","1","1"
],
"actionComment": "Perform server database backup that is
required for disaster r
ecover",
"actionCompletionNotification": {
  "policyCompletionNotificationAction": "",
  "policyCompletionNotificationExecuteOn": "ignore"
},
"actionConcurrent": false,
"actionDrivenBy": "",
"actionEnabled": true,
"actionFailureImpact": "continue",
"actionHardLimit": "00:00",
"actionInactivityTimeout": 30,
"actionParallelism": 0,
"actionRetries": 1,
"actionRetryDelay": 30,
"actionSoftLimit": "00:00"
},
{
  "actionName": "Expiration",
  "actionSpecific": {
    "actions": {
      "actionType": "expire"
    }
  },
  "actionSchedulePeriod": "week",
  "actionScheduleActivity": [
    "exec","exec","exec","exec","exec","exec","exec"
  ],
  "actionComment": "Expire the savesets",
  "actionCompletionNotification": {
    "policyCompletionNotificationAction": "",
    "policyCompletionNotificationExecuteOn": "ignore"
  },
  "actionConcurrent": false,
  "actionDrivenBy": "Server db backup",
  "actionEnabled": true,
  "actionFailureImpact": "continue",
  "actionHardLimit": "00:00",
  "actionInactivityTimeout": 30,
  "actionParallelism": 0,
  "actionRetries": 1,
  "actionRetryDelay": 30,
  "actionSoftLimit": "00:00"
}
],
"workflowAutostartEnabled": true,
"workflowComment": "Perform server backup",
"workflowCompletionNotification": {
  "policyCompletionNotificationAction": "",
  "policyCompletionNotificationExecuteOn": "ignore"
}

```

```

},
"workflowDescription": "server backup action;expire action;",
"workflowEnabled": true,
"workflowGroups": [
  "Server Protection"
],
"workflowInterval": "24:00",
"workflowNextstart": "2015-06-13T10:00:00-0400",
"workflowRestartWindow": "12:00",
"workflowStarttime": "10:00"
},
{
  "workflowName": "NMC server backup",
  "synthesisRoot": [
    "NSR group/NMC server backup",
    "NSR Snapshot Policy/NMC server backup"
  ],
  "workflowActions": [
    {
      "actionName": "NMC server backup",
      "actionSpecific": {
        "actions": {
          "actionType": "backup",
          "actionBackupSubtypeSpecific": {
            "backupSubtypes": {
              "abBackupSubtype": "traditional",
              "abtDestinationPool": "Default",
              "abtEstimate": false,
              "abtFileInactivityAlertThreshold": 0,
              "abtFileInactivityThreshold": 0,
              "abtRevertToFullWhenSyntheticFullFails": true,
              "abtTimestampFormat": "none",
              "abtVerifySyntheticFull": true
            }
          }
        },
        "abDestinationStorageNode": [
          "nsrserverhost"
        ],
        "abRetentionPeriod": "1 Months",
        "abOverrideRetentionPeriod": false,
        "abOverrideBackupSchedule": false,
        "abClientOverridesBehavior": "clientCanOverride"
      }
    },
    {
      "actionSchedulePeriod": "week",
      "actionScheduleActivity": [
        "full", "full", "full", "full", "full", "full", "full"
      ],
      "actionCompletionNotification": {
        "policyCompletionNotificationAction": "",
        "policyCompletionNotificationExecuteOn": "ignore"
      },
      "actionConcurrent": false,
      "actionDrivenBy": "",
      "actionEnabled": true,
      "actionFailureImpact": "continue",

```

```

"actionHardLimit": "00:00",
"actionInactivityTimeout": 30,
"actionParallelism": 100,
"actionRetries": 1,
"actionRetryDelay": 30,
"actionSoftLimit": "00:00"
}
],
"workflowAutostartEnabled": true,
"workflowComment": "Perform NMC database backup",
"workflowCompletionNotification": {
  "policyCompletionNotificationAction": "",
  "policyCompletionNotificationExecuteOn": "ignore"
},
"workflowDescription": "Traditional Backup to pool Default,
with expiration 1 Months;",
"workflowEnabled": true,
"workflowGroups": [
  "NMC server"
],
"workflowInterval": "24:00",
"workflowNextstart": "2015-06-12T14:00:00-0400",
"workflowRestartWindow": "12:00",
"workflowStarttime": "14:00"
}
]
}

```

Identifying clients that missed the workflow schedule

When a client misses the workflow schedule, it logs a message in the workflow logs. Therefore, job records are created so that DPA can easily identify them and inform the end user about this. Clients can miss workflow schedule due to the following reasons:

- Disabled clients included in workflow- Clients that are disabled from NetWorker, still remains part of a workflow. When a client is disabled, it is removed from that workflow run. To identify this, **disabled clients** attribute is added to the workflow job record that has the list of clients that were disabled.
- Workflow does not start at the scheduled time because the previous run still remains active- If a previous workflow is still running at the time of the next scheduled run, the scheduled run does not start. This long running workflow can be due to a slow client, large data change rate, and L0 backups for large hosts. When the workflow does not run, there is no indication in NetWorker and DPA fails to report it. To identify this, a new utility type job record is created which corresponds to the missed workflow. The new completion status displays **missed the schedule** and the completion report displays **previous instance is still running**.
- Workflow is not started since the server is down at the scheduled time- Workflow is scheduled to run at a specific time and does not cover the need for a server maintenance. If the server is down during that time, there is no indication once it restarts that the workflow was not started on time. To identify this, a new utility type job record is created which corresponds to the missed workflow. The new completion status displays **missed the schedule** and the completion report displays **server was down**.

Note: The **disabled clients** attribute is populated irrespective of whether a workflow is missed or not. If the workflow has only one action from the set (check connectivity and clone) then the **disabled clients** attribute is not populated.

Managing policies using the NetWorker Management Web UI

You can use the NetWorker Management Web UI to create your data protection policy solution.

Setting up and configuring data protection policies involves the following tasks:

- Create a policy
- Create a workflow
- Create a protection group
- (Optional) Create rules
- Creating one or more action(s)

Policies, workflows, and actions in NetWorker Management Web UI

In NetWorker Management Web UI, you must use the **Policy** wizard to create policies, workflows, and actions.

The following topics provide more information:

- [Create a policy using the NetWorker Management Web UI](#)
- [Create a workflow using the NetWorker Management Web UI](#)
- [Create an action using the NetWorker Management Web UI](#)

Create a policy using the NetWorker Management Web UI

You can use the NetWorker Management Web UI to create data protection policies.

Procedure

1. Select **Protection > Policies**.

The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.

2. Click **ADD**.

The **Create Policy** wizard appears.

3. Under **Basic Configuration**, in the **Name** field, type a name for the policy.
4. In the **Description** box, type a description for the policy.
5. Select **Enable Protection Period** to specify the protection period.
6. Enter the protection period value in minutes, hours, days, months, or years.
7. Click **NEXT**.
8. Under **Notifications**, from the **Notify** list, select an appropriate notification option.
 - To avoid sending notifications, select **Ignore**.

- To send notifications with information about each successful and failed workflow and action after all the actions in the policy complete, select **On Completion**.
 - To send a notification with information about each failed workflow and action after all the actions in the policy complete, select **On Failure**.
9. Under **Notify**, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver recipient2@mailserver...`

where:

- **-s subject**—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- **-h mailserver**—Specifies the hostname of the mail server to use to relay the SMTP email message.
- **recipient1@mailserver**—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

10. Click **FINISH**.

After you finish

You can now create the workflow, group, and actions for the policy.

Edit a policy using the NetWorker Management Web UI

To edit a policy using the NetWorker Management Web UI, perform the following:

Procedure

1. Select **Protection > Policies**.

The available data protection policies that you can use appears.

2. Select the policy that you want to modify, and click **EDIT**.

- Under **Basic Configuration**, make the required changes.
 - Under **Notification**, make the required changes.
3. Click **Save**.

 **Note:** If you want to delete a policy, select the policy that you want to delete, and click **Delete**.

Create a workflow using the NetWorker Management Web UI

You can create a workflow after you create a new policy, or you can create a workflow for an existing policy.

Procedure

1. Select **Protection > Policies**.
The available data protection policies that you can use appears.
2. Click on an existing policy or the policy that you created.
The details of the selected policy appear in the right pane.
3. Click **ADD**.
4. Under **Basic Configuration**, in the **Name** field, type a name for the workflow, or use the default name.
The maximum number of allowed characters for the Name field is 64. This name cannot contain spaces or special characters such as + or %.
5. In the **Description** box, type a description for the workflow.
6. Under **Schedule**, do the following:
 - a. Select **AutoStart** to start the workflow at the time that is specified in the Start time attribute.
 - b. Use the **Start Time** spin boxes to specify the time to start the actions in the workflow.
 - c. Use the **Interval** attribute spin boxes to specify how frequently to run the actions that are defined in the workflow over a 24-hour period. The default Interval attribute value is 24 hours, or once a day. When you select a value that is less than 24 hours, the **Interval End** attribute appears. To specify the last start time in a defined interval period, use the spin boxes.
 - d. Use the **Restart Window** attribute spin boxes to specify the duration of time in which NetWorker can manually or automatically restart a failed or canceled workflow.
7. Under **Notifications**, from the **Notify** list, select an appropriate notification option.
 - To use the notification configuration that is defined in the policy resource to specify when to send a notification, select **Set at policy level**.
 - To send notifications with information about each successful and failed workflow and action after all the actions in the policy complete, select **On Completion**.
 - To send a notification with information about each failed workflow and action after all the actions in the policy complete, select **On Failure**.
8. Under **Notify**, when you select the **On Completion** option or **On failure** option, the **Command** field appears. Use this field to configure how NetWorker sends

the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver recipient2@mailserver...`

where:

- *-s subject*—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- *-h mailserver*—Specifies the hostname of the mail server to use to relay the SMTP email message.
- *recipient1@mailserver*—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

9. Click **NEXT**.
10. Under **Associate Group**, do one of the following to assign the workflow to a group:
 - Select an existing group, and click **FINISH** to create a workflow.
 - Click **ADD** to create a group. The section [Create a VMware group using the NetWorker Management Web UI](#) provides more information on creating groups.

After you finish

Create the actions that will occur in the workflow.

Edit a workflow using the NetWorker Management Web UI

To edit a workflow using the NetWorker Management Web UI, perform the following:

Procedure

1. Select **Protection > Policies**.

The available data protection policies that you can use appears.

2. Click on an existing policy to view the workflows associated with the policy.

3. Select the workflow that you want to modify, and click **EDIT**.
 - Under **Basic Configuration**, make the required changes.
 - Under **Notification**, make the required changes.
 - Under **Associate Group**, make the required changes.
4. Click **Save**.

 **Note:** If you want to delete a workflow, select the workflow that you want to delete, and click **Delete**.

Start a workflow using the NetWorker Management Web UI

When you start a workflow, NetWorker performs all the actions in the workflow for all the clients that are defined in the groups that are associated with the workflow.

Procedure

1. Select **Protection > Policies**.
The available data protection policies that you can use appears.
2. Click on an existing policy to view the workflows associated with the policy.
3. Select the workflow that you want to start, and click **START**.

 **Note:** The Start button is disabled when the workflow has no association with a group or has no action, or both. When the workflow starts, you see a notification that the workflow has started successfully. When you click on a workflow which is in running state, a notification stating that the workflow failed to start because the workflow has already started appears.

4. When prompted to confirm, click **START**.
The selected workflow is started.

Protection groups

Use the NetWorker Management Web UI to create the following protection groups:

- VMware
- Save Set
- NAS Device
- Client

Create a VMware group using the NetWorker Management Web UI

A VMware group allows you to define the virtual machines or virtual disk files to back up within the policy/workflow.

Procedure

1. Select **Protection > Groups**.
Groups that have already been created appears. The details of the selected group appears in the right pane.
2. Click **ADD**.
The **Create Group** wizard appears.
3. In the **Name** field, type a name for the group.

4. (Optional) In the **Description** box, type a description of the group.
 5. From the **Type** list, select **VMware**.
 6. (Optional) From the **Policy** list, select a policy that you want to assign the group to.
 7. (Optional) From the **Workflow** list, select a workflow that you want to assign the group to.
-  **Note:** The applicable workflows corresponding to the selected policy are only available for selection.
8. Select a backup optimization mode.
 - **Capacity** - Use for variable segment sizing.
 - **Performance** - Use for fixed segment sizing.
 9. (Optional) Select the Dynamic Association checkbox if you plan to apply rules that will determine which virtual machines and containers are dynamically included in the group based upon the rule criteria.
 10. Click **NEXT**.
 11. From the **vCenter** list, select a vCenter server that contains the VMware objects that you want to protect, or click **ADD VCENTER** to add a vCenter server. The section Add the vCenter server using the NetWorker Management Web UI of the NetWorker VMware Integration Guide provides instructions to add a vCenter server.
 12. Select a pre-defined rule that you want to apply for any VMware objects that are dynamically included in the group based upon the rule criteria, or click **ADD RULE** to create a new rule and click **EDIT RULE** to modify an existing rule. The section [Create or edit a rule using the NetWorker Management Web UI](#) provides more information on creating and editing rules.
 13. Under **Include/Exclude Resources**, select the objects (Datacenter, ESXi host, virtual machine, resource pool, vApp, or disk) to include in the group and to view a consolidated list of virtual machines selected statically and evaluated by a rule. You can also exclude virtual machines or disks.
 14. Click **Preview** to view a list of the static and dynamic virtual machines and objects that have been added to the group.
 15. Click **Finish**.

Create a save set group using the NetWorker Management Web UI

A save set group defines a static list of save sets for cloning or for snapshot index generation.

Procedure

1. Select **Protection > Groups**.
Groups that have already been created appears. The details of the selected group appears in the right pane.
2. Click **ADD**.
The **Create Group** wizard appears.
3. In the **Name** field, type a name for the group.
4. (Optional) In the **Description** box, type a description of the group.
5. From the **Type** list, select **Save Set**.

6. (Optional) From the **Policy** list, select a policy that you want to assign the group to.
7. (Optional) From the **Workflow** list, select a workflow that you want to assign the group to.

 **Note:** The applicable workflows corresponding to the selected policy are only available for selection.

8. Specify a **Filter Type**.
 - **Static**- Represents the save set ID list.
 - **Dynamic** - Represents the save set query.

If the filter type is **Static**:

- a. Click **NEXT**.
- b. Enter a list of comma separated save set IDs, and click **ADD**.

 **Note:** You can also use the search function to filter the save set IDs from the **Selected Saveset ID List**.

- c. Click **FINISH**.

If the filter type is **Dynamic**:

- a. Use the **Maximum number of Clones** spin boxes to specify the number of clones that can be created for the save set.
- b. Click **NEXT**.
- c. (Optional) Specify one or more of the basic save set criteria.

Filter criteria	Description
Time	Specify the start date and time range for the save sets. To specify the current date and time as the end date for the range, select Up To Now . To specify a time period, select Up to .
Levels	Specify a backup level.
Name	Specify a list of comma separated save set names.

- d. Click the + icon to add a filter, and click the - icon to remove a filter.
- e. Click **NEXT**.
- f. (Optional) Select the save set resource filter.

Resource filter	Description
Clients	Select client resources associated with the save set.
Workflows	Select the workflow used to generate the save set.

Resource filter	Description
Policies	Select the policy used to generate the save set.
Actions	Select the action used to generate the save set.
Groups	Select the group associated with the save set.
SourcePools	Select the pool on which the save set is stored.

g. Click **FINISH**.

Create a NAS device group using the NetWorker Management Web UI

A NAS device group defines a list of NAS devices for discovery.

Procedure

1. Select **Protection > Groups**.

Groups that have already been created appears. The details of the selected group appears in the right pane.

2. Click **ADD**.

The **Create Group** wizard appears.

3. In the **Name** field, type a name for the group.
4. (Optional) In the **Description** box, type a description of the group.
5. From the **Type** list, select **NAS Device**.
6. (Optional) From the **Policy** list, select a policy that you want to assign the group to.
7. (Optional) From the **Workflow** list, select a workflow that you want to assign the group to.



Note: The applicable workflows corresponding to the selected policy are only available for selection.

8. Click **NEXT**.
9. In the **NAS Devices** section, add or remove NAS devices.
10. Click **Finish**.

Create a client group using the NetWorker Management Web UI

You can create a static client group or a dynamic client group for a traditional backup or a server backup.

Procedure

1. Select **Protection > Groups**.

Groups that have already been created appears. The details of the selected group appears in the right pane.

2. Click **ADD**.

The **Create Group** wizard appears.

3. In the **Name** field, type a name for the group.
4. (Optional) In the **Description** box, type a description of the group.
5. From the **Type** list, select **Client**.
6. (Optional) From the **Policy** list, select a policy that you want to assign the group to.
7. (Optional) From the **Workflow** list, select a workflow that you want to assign the group to.

 **Note:** The applicable workflows corresponding to the selected policy are only available for selection.

8. Specify the **Filter Type**:
 - **Static** - Specify a static list of client resources for a traditional backup, a check connectivity action, or a probe action.
 - **Dynamic** - Specify a dynamic list of client resources for a traditional backup, a probe action, a check connectivity action, or a server backup action.
9. If the filter type is **Static**:
 - a. Click **NEXT**.
 - b. In the **Client Resources** section, select the clients from the list to add it to the group. The Client Resources section displays a list of configured clients on the NetWorker server.
 - c. Click **FINISH**.
10. If the filter type is **Dynamic**:
 - a. Click **NEXT**.
 - b. In the **Client Resources** section, do one of the following:
 - To include all the clients on the NetWorker server, select **All Clients**, and click **FINISH**.
 - To define a tag, select **Client with these tags** and specify a tag attribute, and then click **ADD**.

 **Note:** You can also use the search function to filter the tags from the **Selected Client Tags**.

- c. Click **FINISH**.

Edit a group using the NetWorker Management Web UI

To edit a group using the NetWorker Management Web UI, perform the following:

Procedure

1. Select **Protection > Groups**.

Groups that have already been created appears.

2. Select the group that you want to modify.

The details of the group such as the type, sub type, the backup optimization mode, and the VMware resources appears in the right pane.

 **Note:** You cannot change the name of the group.

3. Click **EDIT**, and make the required changes.
4. Click **Save**.



Note: If you want to delete a group, select the group that you want to delete, and click **Delete**.

Create an action using the NetWorker Management Web UI

Actions are the key resources in a workflow for a data protection policy and define a specific task, for example, a backup or a clone. An action is the task that occurs on the client resources in the group assigned to the workflow. You can chain multiple actions together to occur sequentially or concurrently in a workflow.

Before you begin

Create the required policy and workflow.

About this task

To create an action:

Procedure

1. Select **Protection > Policies**.

The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.

2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow to create an action. The action type can be one of the following:
 - Backup:
 - Traditional—Is a scheduled backup of the save sets defined for the Client resources in the assigned group.
 - VMware (vProxy)—Performs a backup of virtual machines in vCenter to a Data Domain system. You can only perform one VMware backup action per workflow. The VMware backup action must occur before clone actions.
 - Snapshot—Creates snapshot backups of supported storage arrays or appliances.
 - Server—Performs NetWorker server bootstrap and database backups.
 - Clone—Performs a clone of the VMware backup on a Data Domain system to any clone device that NetWorker supports (including Data Domain system or tape targets). You can specify multiple clone actions. Clone actions must occur after the Backup action.
 - Probe—Runs a user-defined script on a NetWorker client before the start of a backup.
 - Check Connectivity—Verifies the connectivity between the clients and the NetWorker server before the start of a probe or backup action. If the connectivity test fails, then the probe action and backup action does not start for the client.
 - Expire—Marks the expired save sets as recyclable.

- Discover—Discovers snapshots on a NAS device that the current NetWorker server did not create. NetWorker then creates save set entries for the discovered snapshots on a NAS client in the media database.
- Generate Index—Creates entries in the client file indexes for snapshot backups and discovered snapshots.

 Note:

- If you have a VMware group associated with the workflow, the first action that you can create is always a backup action, followed by any number of clone actions.
 - If you have a Saveset group associated with the workflow, you can create clone or generate index actions.
 - If you have a NAS device group associated with the workflow, you can create a discover action.
 - If you have a client group type associated with the workflow, you can create a traditional backup, snapshot backup, server backup, probe, and check connectivity actions.
 - If you do not have any groups associated with the workflow, you can create any of the supported actions.
 - For the expire action to be the first action in a workflow, there must not any group associated with the workflow.
- The section [Create a traditional backup action using the NetWorker Management Web UI](#) provides more information on creating traditional backup actions
 - The section [Create a VMware backup action using the NetWorker Management Web UI](#) provides more information on creating VMware backup actions.
 - The section [Create a snapshot backup action using the NetWorker Management Web UI](#) provides more information on creating snapshot backup actions
 - The section [Create a server backup action using the NetWorker Management Web UI](#) provides more information on creating server backup actions.
 - The section [Create a clone action using the NetWorker Management Web UI](#) provides more information on creating clone actions.
 - The section [Create a probe action using the NetWorker Management Web UI](#) provides more information on creating probe actions.
 - The section [Create a check connectivity action using the NetWorker Management Web UI](#) provides more information on creating check connectivity actions.
 - The section [Create an expire action using the NetWorker Management Web UI](#) provides more information on creating expire actions.
 - The section [Create a discover action using the NetWorker Management Web UI](#) provides more information on creating discover actions.
 - The section [Create a generate index action using the NetWorker Management Web UI](#) provides more information on creating generate index actions.

Create a traditional backup action using the NetWorker Management Web UI

You can use the NetWorker Management Web UI to create a traditional backup action.

Before you begin

Create the required policy and workflow.

Procedure

1. Select **Protection > Policies**.

The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.

2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow with the appropriate group association, or create a new workflow, and click **ACTION**.
4. From the **ACTION** list, select **Backup > Traditional** to create a traditional backup action.

The Create Action page appears.

5. Under **Basic Configuration**, in the **Name** field, type the name of the action, or use the default name.

The maximum number of characters for the action name is 64.

6. In the **Description** field, type a description for the action.
7. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the policy or workflow that contains the action is started, clear this option. By default, this option is selected.
8. (Optional) From the **Force Backup Level** list, select a backup level.

For workflows that have more than one scheduled backup within a 24-hour period, use the Force Backup Level attribute to allow more than one backup to occur at two different backup levels in a 24-hour period. When you select a backup level in the Force Backup Level attribute, the first backup is performed at the scheduled backup level. Each subsequent occurrence of the backup action in the next 24 hours occurs at the level defined in the Force Backup Level attribute. For example, if the level defined by the schedule is Full and the Force Backup Level attribute is set to Incr, the first backup started by the action occurs at a level full and subsequent backups, within 24 hours of the start of the full backup are incremental. By default, this option is cleared, which means that if the action runs multiple backup operations in a 24 period, all the backups occur at the scheduled backup level.

9. Click **NEXT**.
10. Under **Destination Storage Node**, add or remove storage nodes. The server storage node is selected by default. You can also order the storage nodes.
11. From the **Destination Pool**, select the media pool in which to store the backup data.
12. Specify the **Retention Period** value.

After the retention period expires, the save set is removed from the client file index and marked as recyclable in the media database during an expiration server maintenance task.

13. From the **Client Override Behavior** list, specify how NetWorker uses certain client configuration attributes that perform the same function as attributes in the Action resource:
 - **Client Can Override**—The values in the Client resource for Schedule, Pool, Retention policy, and the Storage Node attributes take precedence over the values that are defined in the equivalent Action resource attributes.

Note: If the NetWorker policy action schedule is set to the Skip backup level, the **Client can Override** option is not honored. For NetWorker to consider the **Client can Override** option, change the action schedule to a different level.
 - **Client Cannot Override**—The values in the Action resource for the Schedule, Destination Pool, Destination Storage Node, and the Retention attributes take precedence over the values that are defined in the equivalent Client resource attributes.
 - **Legacy Backup Rules**—This value only appears in actions that are created by the migration process. The updating process sets the Client Override Behavior for the migrated backup actions to Legacy Backup Rules.
14. From the **Success Threshold** list, select one of the following to define the success threshold for a save session:
 - **Warning**—Save sets that complete with warnings are reported as success with warnings.
 - **Success**—Save sets that complete with warnings are reported as failed. This is the default value. The number of times NetWorker retries a failed save set is determined by the value defined in the **Retries** attribute.
15. Select **Apply Lock** under **DD Retention Lock** to enable retention lock for the virtual machines included in this backup action. Note that the device used for backing up these virtual machines must also have DD Retention lock enabled.
16. Use the **Lock Period** spin boxes to specify the duration the virtual machines will remain on the Data Domain device before the retention lock expires.
17. Click **NEXT**.
18. Select a schedule configuration:
 - **Select**—Allows you to select a pre-defined schedule that is applicable to the selected action type, or create a new schedule. Click **ADD SCHEDULE** to create a new schedule. The section [Create a schedule using the NetWorker Management Web UI](#) provides more information on creating schedules.

Note: For traditional backup actions, you can only select or create a backup type schedule.
 - **Define**—Allows you to define a schedule.

For **Define**, do the following:

- a. Specify a weekly or monthly recurrence schedule for the action.
 - To specify a schedule for each day of the week, select **Weekly**.
 - To specify a schedule for each day of the month, select **Monthly**.
- b. Click on each day to specify the type of backup to perform.

The backup levels for NetWorker includes the following:

- **Full**—Perform a full backup on the specified day. Full backups include all files, regardless of whether the files changed.

- **Incr**—Perform an incremental backup on the specified day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
- **Cumulative Incr**—Perform a cumulative incremental backup. Cumulative incremental backups include files that have changed since the last full backup.
- **Logs Only**—Perform a backup of only database transaction logs.
- **Incremental Synthetic Full**—Perform an incremental synthetic backup on the specified day. An incremental synthetic full backup includes all data that changed since the last full backup and subsequent incremental backups to create a synthetic full backup.
- **Skip**—Do not perform a backup on the specified day.

To perform the same type of backup on each day, click , and select one of the following:

- Make All Full
- Make All Incremental
- Make All Cumulative Incremental
- Make All Logs Only
- Make All Incremental Synthetic Full
- Make All Skip

c. (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

 **Note:** If you select this option, you must configure the overrides using either fixed dates or recurring dates.

d. Click **NEXT**.

19. Under **Schedule Overrides**, do the following:

a. In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

b. To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

 **Note:**

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.

- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.
- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

20. Click **NEXT**.

21. Under **Notifications**, from the **Notify** list, select an appropriate notification option.

- To avoid sending notifications, select **Ignore**.
- To send notifications on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

22. Specify the **Command** to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver recipient2@mailserver...`

where:

- **-s subject**—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- **-h mailserver**—Specifies the hostname of the mail server to use to relay the SMTP email message.
- **recipient1@mailserver**—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

23. Click **NEXT**.

24. From the **Start Time** list, select one of the following options to specify the time to start the action. Use the spin boxes to set the hour and minute values.
 - **Set at Workflow level**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
25. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. The default value is 100.
26. From the **On Failure** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort Action**.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort Workflow**.
27. In the **Retries** field, specify the number of times that NetWorker must retry a failed probe or backup action, before NetWorker considers the action as failed. If the Retries value is 0, NetWorker does not retry a failed probe or backup action. The default value is 1.
28. In the **Retry Delay** field, specify a delay in seconds to wait before retrying a failed probe or backup action. When the Retry Delay value is 0, NetWorker retries the failed probe or backup action immediately. The default is 1 second.
29. In the **Inactivity Timeout** field, specify the maximum number of minutes that a job run by an action can try to respond to the server.

 If the job does not respond within the specified time, the server considers the job a failure and NetWorker retries the job immediately to ensure that no time is lost due to failures. The default value is 30 minutes.

 Increase the timeout value if a backup consistently stops due to inactivity. Inactivity might occur for backups of large save sets, backups of save sets with large sparse files, and incremental backups of many small static files.
30. In the **Soft Limit** field, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
31. In the **Hard Limit** field, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
32. Click **NEXT**.

 The Action Configuration Summary appears.
33. Review the settings that you have configured, and then click **FINISH**.

After you finish

(Optional) Create a clone action to automatically clone the save sets after the backup. A clone action is the only supported action after a backup action in a workflow.

Create a VMware backup action using the NetWorker Management Web UI

You can use the NetWorker Management Web UI to create a VMware backup action.

Before you begin

Create the required policy and workflow.

Procedure

1. Select **Protection > Policies**.

The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.

2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow with the appropriate group association, or create a new workflow, and click **ACTION**.
4. From the **ACTION** list, select **Backup > VMWare** to create a VMware backup action.
5. Under **Basic Configuration**, in the **Name** field, type the name of the action, or use the default name.

The maximum number of characters for the action name is 64.
6. In the **Description** field, type a description for the action.
7. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the policy or workflow that contains the action is started, clear this option. By default, this option is selected.
8. Click **NEXT**.
9. From the **Destination Storage Node** list, select the storage node that contains the devices where you want to store the backup data.

 **Note:** When you deploy the vCenter server in the Cloud, a parameter displays in the backup action logs that indicates HypervisorMode: VMC. When not deployed in the Cloud, the parameter indicates HypervisorMode: vSphere.

10. From the **Destination Pool**, select the media pool in which to store the backup data.
11. Specify the **NetWorker Retention Period** value.

After the retention period expires, the save set is removed from the client file index and marked as recyclable in the media database during an expiration server maintenance task.
12. Select **Apply Lock** under **DD Retention Lock** to enable retention lock for the virtual machines included in this backup action. Note that the device used for backing up these virtual machines must also have DD Retention lock enabled.
13. Use the **Lock Period** spin boxes to specify the duration the virtual machines will remain on the Data Domain device before the retention lock expires.
14. Click **NEXT**.
15. Under **VMware-specific Configuration**, select one of the following vProxy options:

- **Automatic**—Select this option to allow NetWorker to choose the vProxy host for backups.
 - **Manual**—Specify this option to select the vProxy host that NetWorker users for backups. Provide the name of the vProxy host in the **vProxy Name** field.
16. Under **Application Consistency**, select **Quiesce Application** to enable application-consistent protection as part of the policy backup action, which includes protection of the Microsoft SQL Server. You can then select from the **Basic** and **Advanced** options.
- Select the **Basic** option to create a backup copy for applications during virtual machine quiescing. No additional processing is performed.
 - Select the **Advanced** option to create an SQL server application-consistent backup during virtual machine quiescing, and optionally create a transaction log backup for all SQL Server instances.

When you select the **Advanced** option, the following additional options appear:

- **Transaction Log Backup**—Select this option, if you want to perform a transaction log backup of SQL databases in the virtual machine as part of the policy backup action.

 **Note:** During SQL Server configuration, the NT AUTHORITY\SYSTEM login must be granted SQL login and SQL sysadmin role rights in order to perform transaction log backups.
 - **Quiesce Timeout**—Specify the amount of time, in minutes, to wait for the quiesce operation on the virtual machine to time out before failing. If not selected, the backup action proceeds even if quiescing was not performed, unless a validation problem occurs. If an application-consistent backup cannot complete due to a problem with validation, the backup action fails even if this option is not selected.
 - **System Administrator Username and System Administrator Password**—Specify the virtual machine credentials for a user with administrative privileges. All virtual machines in the workflow must use the same System Administrator username/password.
-  **Note:** If you select the Advanced option, application-consistent processing is applied for all virtual machines within the parent workflow. When selecting this option, ensure that the policy's workflow and client groups are provisioned specifically for virtual machines that require advanced application-consistent protection. NetWorker will always attempt to perform advanced application processing for virtual machines in a workflow that contains a backup action with advanced application processing enabled.

17. Click **NEXT**.

18. Select a schedule configuration:

- **Select**—Allows you to select a pre-defined schedule that is applicable to the selected action type, or create a new schedule. Click **ADD SCHEDULE** to create a new schedule. The section [Create a schedule using the NetWorker Management Web UI](#) provides more information on creating schedules.

 **Note:** For VMware backup actions, you can only select or create a VMware type schedule.
- **Define**—Allows you to define a schedule.

For **Define**, do the following:

- a. Specify a weekly or monthly recurrence schedule for the action.
 - To specify a schedule for each day of the week, select **Weekly**.
 - To specify a schedule for each day of the month, select **Monthly**.

- b. Click on each day to specify the type of backup to perform.

The backup levels for NetWorker includes the following:

- **Full**—Perform a full backup on the specified day. Full backups include all files, regardless of whether the files changed.
- **Incr**—Perform an incremental backup on the specified day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
- **Logs Only**—Perform a backup of only database transaction logs.
- **Skip**—Do not perform a backup on the specified day.

To perform the same type of backup on each day, click , and select one of the following:

- Make All Full
- Make All Incremental
- Make All Logs Only
- Make All Skip

- c. (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

 **Note:** If you select this option, you must configure the overrides using either fixed dates or recurring dates.

- d. Click **NEXT**.

19. Under **Schedule Overrides**, do the following:

- a. In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

- b. To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

 **Note:**

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.

- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.
- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

20. Click **NEXT**.

21. Under **Notifications**, from the **Notify** list, select an appropriate notification option.

- To avoid sending notifications, select **Ignore**.
- To send notifications on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

22. Specify the **Command** to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver recipient2@mailserver...`

where:

- **-s subject**—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- **-h mailserver**—Specifies the hostname of the mail server to use to relay the SMTP email message.
- **recipient1@mailserver**—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

23. Click **NEXT**.

24. From the **Start Time** list, select one of the following options to specify the time to start the action. Use the spin boxes to set the hour and minute values.
 - **Set at Workflow level**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
25. Although the **Retries**, **Retry Delay**, and the **Inactivity Timeout** options appear, this action does not support these options and you can ignore these values.
26. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. The default value is 100.
27. From the **On Failure** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort Action**.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort Workflow**.
28. In the **Soft Limit** field, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
29. In the **Hard Limit** field, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
30. Click **NEXT**.
The Action Configuration Summary appears.
31. Review the settings that you have configured, and then click **FINISH**.

After you finish

(Optional) Create a clone action to automatically clone the save sets after the backup. A clone action is the only supported action after a backup action in a workflow.

Create a snapshot backup action using the NetWorker Management Web UI

You can use the NetWorker Management Web UI to create a snapshot backup action.

Before you begin

Create the required policy and workflow.

Procedure

1. Select **Protection > Policies**.
The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.
2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow with the appropriate group association, or create a new workflow, and click **ACTION**.
4. From the **ACTION** list, select **Backup > Snapshot** to create a snapshot backup action.

The Create Action page appears.

5. Under **Basic Configuration**, in the **Name** field, type the name of the action, or use the default name.

The maximum number of characters for the action name is 64.

6. In the **Description** field, type a description for the action.
7. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the policy or workflow that contains the action is started, clear this option. By default, this option is selected.
8. (Optional) From the **Force Backup Level** list, select a backup level.

For workflows that have more than one scheduled backup within a 24-hour period, use the Force Backup Level attribute to allow more than one backup to occur at two different backup levels in a 24-hour period. When you select a backup level in the Force Backup Level attribute, the first backup is performed at the scheduled backup level. Each subsequent occurrence of the backup action in the next 24 hours occurs at the level defined in the Force Backup Level attribute. For example, if the level defined by the schedule is Full and the Force Backup Level attribute is set to Incr, the first backup started by the action occurs at a level full and subsequent backups, within 24 hours of the start of the full backup are incremental. By default, this option is cleared, which means that if the action runs multiple backup operations in a 24 period, all the backups occur at the scheduled backup level.

9. Click **NEXT**.
10. Under **Destination Storage Node**, add or remove storage nodes. The server storage node is selected by default. You can also order the storage nodes.
11. From the **Destination Pool**, select the media pool in which to store the backup data.
12. Specify the **Retention Period** value.
After the retention period expires, the save set is removed from the client file index and marked as recyclable in the media database during an expiration server maintenance task.
13. From the **Client Override Behavior** list, specify how NetWorker uses certain client configuration attributes that perform the same function as attributes in the Action resource:

- **Client Can Override**—The values in the Client resource for Schedule, Pool, Retention policy, and the Storage Node attributes take precedence over the values that are defined in the equivalent Action resource attributes.

 **Note:** If the NetWorker policy action schedule is set to the Skip backup level, the **Client can Override** option is not honored. For NetWorker to consider the **Client can Override** option, change the action schedule to a different level.

- **Client Cannot Override**—The values in the Action resource for the Schedule, Destination Pool, Destination Storage Node, and the Retention attributes take precedence over the values that are defined in the equivalent Client resource attributes.
- **Legacy Backup Rules**—This value only appears in actions that are created by the migration process. The updating process sets the Client Override Behavior for the migrated backup actions to Legacy Backup Rules.

14. Use the **Minimum Retention Time** spin boxes to specify the minimum amount of time to retain the backup data.

When the minimum retention time passes, a snapshot action that is in progress can remove the snapshot from the storage device to ensure that there is sufficient disk space for the new snapshot.

15. Click **NEXT**.

16. Select a schedule configuration:

- **Select**—Allows you to select a pre-defined schedule that is applicable to the selected action type, or create a new schedule. Click **ADD SCHEDULE** to create a new schedule. The section [Create a schedule using the NetWorker Management Web UI](#) provides more information on creating schedules.



Note: For snapshot backup actions, you can only select or create a backup type schedule.

- **Define**—Allows you to define a schedule.

For **Define**, do the following:

- a. Specify a weekly or monthly recurrence schedule for the action.

- To specify a schedule for each day of the week, select **Weekly**.
- To specify a schedule for each day of the month, select **Monthly**.

- b. Click on each day to specify the type of backup to perform.

The backup levels for NetWorker includes the following:

- **Full**—Perform a full backup on the specified day. Full backups include all files, regardless of whether the files changed.
- **Incr**—Perform an incremental backup on the specified day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
- **Cumulative Incr**—Perform a cumulative incremental backup. Cumulative incremental backups include files that have changed since the last full backup.
- **Logs Only**—Perform a backup of only database transaction logs.
- **Incremental Synthetic Full**—Perform an incremental synthetic backup on the specified day. An incremental synthetic full backup includes all data that changed since the last full backup and subsequent incremental backups to create a synthetic full backup.
- **Skip**—Do not perform a backup on the specified day.

To perform the same type of backup on each day, click , and select one of the following:

- Make All Full
- Make All Incremental
- Make All Cumulative Incremental
- Make All Logs Only
- Make All Incremental Synthetic Full

- Make All Skip

c. (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

 **Note:** If you select this option, you must configure the overrides using either fixed dates or recurring dates.

d. Click **NEXT**.

17. Under **Schedule Overrides**, do the following:

a. In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

b. To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

 **Note:**

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.
- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.
- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

18. Click **NEXT**.

19. Under **Notifications**, from the **Notify** list, select an appropriate notification option.

- To avoid sending notifications, select **Ignore**.
- To send notifications on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

20. Specify the **Command** to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver
recipient2@mailserver...`

where:

- *-s subject*—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- *-h mailserver*—Specifies the hostname of the mail server to use to relay the SMTP email message.
- *recipient1@mailserver*—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

21. Click **NEXT**.
22. From the **Start Time** list, select one of the following options to specify the time to start the action. Use the spin boxes to set the hour and minute values.
 - **Set at Workflow level**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
23. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. The default value is 100.
24. From the **On Failure** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort Action**.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort Workflow**.
25. In the **Retries** field, specify the number of times that NetWorker must retry a failed probe or backup action, before NetWorker considers the action as failed. If the Retries value is 0, NetWorker does not retry a failed probe or backup action. The default value is 1.
26. In the **Retry Delay** field, specify a delay in seconds to wait before retrying a failed probe or backup action. When the Retry Delay value is 0, NetWorker retries the failed probe or backup action immediately. The default is 1 second.

27. In the **Inactivity Timeout** field, specify the maximum number of minutes that a job run by an action can try to respond to the server.

If the job does not respond within the specified time, the server considers the job a failure and NetWorker retries the job immediately to ensure that no time is lost due to failures. The default value is 30 minutes.

Increase the timeout value if a backup consistently stops due to inactivity. Inactivity might occur for backups of large save sets, backups of save sets with large sparse files, and incremental backups of many small static files.
28. In the **Soft Limit** field, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
29. In the **Hard Limit** field, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
30. Click **NEXT**.

The Action Configuration Summary appears.
31. Review the settings that you have configured, and then click **FINISH**.

Create a server backup action using the NetWorker Management Web UI

You can use the NetWorker Management Web UI to create a server backup action.

Before you begin

Create the required policy and workflow.

Procedure

1. Select **Protection > Policies**.

The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.
2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow with the appropriate group association, or create a new workflow, and click **ACTION**.
4. From the **ACTION** list, select **Backup > Server** to create a server backup action.

The Create Action page appears.
5. Under **Basic Configuration**, in the **Name** field, type the name of the action, or use the default name.

The maximum number of characters for the action name is 64.
6. In the **Description** field, type a description for the action.
7. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the policy or workflow that contains the action is started, clear this option. By default, this option is selected.
8. Click **NEXT**.
9. From the **Destination Storage Node** list, select the storage node that contains the devices where you want to store the backup data.

10. From the **Destination Pool**, select the media pool in which to store the backup data.
11. Specify the **Retention Period** value.

After the retention period expires, the save set is removed from the client file index and marked as recyclable in the media database during an expiration server maintenance task.

12. Select the **Backup Options**.

- **Client File Index** —Select this option to backup client file indexes residing on the NetWorker Server of NetWorker Clients selected as part of the group.
- **Bootstrap** — Select this option to backup media and authentication databases of the NetWorker Server .

 **Note:** NetWorker requires that the bootstrap backup write to a device that is local to the NetWorker server. If you enable bootstrap, ensure that you select the NetWorker server in the Destination Storage Node attribute.

13. Select **Apply Lock** under **DD Retention Lock** to enable retention lock for the virtual machines included in this backup action. Note that the device used for backing up these virtual machines must also have DD Retention lock enabled.
14. Use the **Lock Period** spin boxes to specify the duration the virtual machines will remain on the Data Domain device before the retention lock expires.
15. Click **NEXT**.
16. Select a schedule configuration:

- **Select**—Allows you to select a pre-defined schedule that is applicable to the selected action type, or create a new schedule. Click **ADD SCHEDULE** to create a new schedule. The section [Create a schedule using the NetWorker Management Web UI](#) provides more information on creating schedules.

 **Note:** For server backup actions, you can only select or create a server backup type schedule.

- **Define**—Allows you to define a schedule.

For **Define**, do the following:

- a. Specify a weekly or monthly recurrence schedule for the action.
 - To specify a schedule for each day of the week, select **Weekly**.
 - To specify a schedule for each day of the month, select **Monthly**.

- b. Click on each day to specify the type of backup to perform.

The backup levels for NetWorker includes the following:

- **Full**—Perform a full backup on the specified day. Full backups include all files, regardless of whether the files changed.
- **Incr**—Perform an incremental backup on the specified day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
- **Cumulative Incr**—Perform a cumulative incremental backup. Cumulative incremental backups include files that have changed since the last full backup.

- **Skip**—Do not perform a backup on the specified day.

To perform the same type of backup on each day, click , and select one of the following:

- Make All Full
- Make All Incremental
- Make All Cumulative Incremental
- Make All Skip

- c. (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

 **Note:** If you select this option, you must configure the overrides using either fixed dates or recurring dates.

- d. Click **NEXT**.

17. Under **Schedule Overrides**, do the following:

- a. In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

- b. To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

 **Note:**

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.
- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.
- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

18. Click **NEXT**.

19. Under **Notifications**, from the **Notify** list, select an appropriate notification option.

- To avoid sending notifications, select **Ignore**.

- To send notifications on completion of the action, select **On Completion**.
 - To send a notification only if the action fails to complete, select **On Failure**.
20. Specify the **Command** to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver recipient2@mailserver...`

where:

- *-s subject*—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- *-h mailserver*—Specifies the hostname of the mail server to use to relay the SMTP email message.
- *recipient1@mailserver*—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

21. Click **NEXT**.
22. From the **Start Time** list, select one of the following options to specify the time to start the action. Use the spin boxes to set the hour and minute values.
- **Set at Workflow level**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
23. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. The default value is 100.
24. From the **On Failure** list, specify what to do when a job fails:
- To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort Action**.

- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort Workflow**.
25. In the **Soft Limit** field, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
 26. In the **Hard Limit** field, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
 27. Click **NEXT**.
The Action Configuration Summary appears.
 28. Review the settings that you have configured, and then click **FINISH**.

After you finish

(Optional) Create a clone or an expire action after the server backup action.

Create a clone action using the NetWorker Management Web UI

You can use the NetWorker Management Web UI to create a clone action.

Procedure

1. Select **Protection > Policies**.
The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.
2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow with the appropriate group association, or create a new workflow, and click **ACTION**.
4. From the **ACTION** list, select **Clone** to create a clone action.
5. Under **Basic Configuration**, in the **Name** field, type the name of the action, or use the default name.
The maximum number of characters for the action name is 64.
6. In the **Description** box, type a description for the action.
7. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the policy or workflow that contains the action is started, clear this option. By default, this option is selected.
8. If the action is part of a sequence of actions in a workflow path, from the **Driven By** list, select the action that should precede this action.
9. Click **NEXT**.
10. Select the **Delete source save sets after clone completes** option to instruct NetWorker to delete the data from the source volume after cloning to the destination volume completes. This is equivalent to staging the save sets.
11. Under **Devices and Volumes**, define the volumes and devices to which NetWorker sends the cloned data:
 - a. From the **Source Storage Node** list, select the source storage node for a clone action, that is, the storage node from which clone data is read.
 - b. From the **Destination Storage Node** list, select the storage node with the devices on which to store the cloned save sets.

- c. From the **Destination Pool** list, select the target media pool for the cloned save sets.

12. Specify the **NetWorker Retention Period** value.

After the retention period expires, the save set is removed from the client file index and marked as recyclable in the media database during an expiration server maintenance task.

13. If you want to have the same retention period as that of the backup action, select **Keep Retention Same as Backup**.

 **Note:** If you select this option, you cannot configure the NetWorker retention period.

14. Select **Apply Lock** under **DD Retention Lock** to enable retention lock for the virtual machines included in this clone action.

15. Use the **Lock Period** spin boxes to specify the duration the virtual machines will remain on the Data Domain device before the retention lock expires.

16. Click **NEXT**.

17. (Applicable only for a second clone action) In the **Filter Savesets** section, you can do the following:

- a. To define the criteria that NetWorker uses to create the list of eligible save sets to clone, select **Define Filters**. The eligible save sets must match the requirements that are defined in each filter. NetWorker provides the following filter options:

- **Time** —In the Time section, specify the time range in which NetWorker searches for eligible save sets to clone in the media database. Use the spin boxes to specify the start time and the end time. The Time filter list includes the following options to define how NetWorker determines save set eligibility, based on the time criteria:
 - **Accept**—The clone save set list includes save sets that are saved within the time range and meet all the other defined filter criteria.
 - **Reject**—The clone save set list does not include save sets that are saved within the time range and meet all the other defined filter criteria.
- **Backup Levels** —In the Backup Levels section, specify the backup levels that you want to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The Levels filter list includes the following options, which define how NetWorker determines save set eligibility, based on the backup level filter criteria:
 - **Accept**—The clone save set list includes eligible save sets with the selected backup levels.
 - **Reject**—The clone save set list does not include eligible save sets with the selected backup levels.
- **Save Sets** —In the Save Sets section, specify whether to include or exclude ProtectPoint and Snapshot save sets, when NetWorker searches for eligible save sets to clone in the media database. The Save Set filter list includes the following options, which define how NetWorker determines save set eligibility, based on the save set filter criteria:
 - **Accept**—The clone save set list includes eligible ProtectPoint save sets or Snapshot save sets, when you enable the ProtectPoint checkbox or Snapshot checkbox.

- **Reject**—The clone save set list does not include eligible ProtectPoint save sets and Snapshot save sets when you enable the ProtectPoint checkbox or Snapshot checkbox.
- **Clients** —In the Client section, specify a list of clients to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The Client filter list includes the following options, which define how NetWorker determines save set eligibility, based on the client filter criteria:
 - **Accept**—The clone save set list includes eligible save sets for the selected clients.
 - **Reject**—The clone save set list does not include eligible save sets for the selected clients.

18. If you do not want to define a filter criteria, select **Do Not Filter**, and click **NEXT**.

 **Note:** The Filter Savesets feature is available only in the case of a second clone action.

19. Select a schedule configuration:

- **Select**—Allows you to select a pre-defined schedule that is applicable to the selected action type, or create a new schedule. Click **ADD SCHEDULE** to create a new schedule. The section [Create a schedule using the NetWorker Management Web UI](#) provides more information on creating schedules.

 **Note:** For clone actions, you can only select or create an execute type schedule.
- **Define**—Allows you to define a schedule.

For **Define**, do the following:

- a. Specify a weekly or monthly recurrence schedule for the action.
 - To specify a schedule for each day of the week, select **Weekly**.
 - To specify a schedule for each day of the month, select **Monthly**.
- b. Click on each day to specify the type of backup to perform.
 - **Execute**—Perform a clone action on the specified day.
 - **Skip**—Do not perform a clone action on the specified day.

To perform the same type of backup on each day, click , and select one of the following:

- Make All Execute
- Make All Skip

c. (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

 **Note:** If you select this option, you must configure the overrides using either fixed dates or recurring dates.

d. Click **NEXT**.

20. Under **Schedule Overrides**, do the following:

- a. In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

- b. To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

 Note:

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.
- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.
- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

21. Click **NEXT**.

22. Under **Notifications**, from the **Notify** list, select an appropriate notification option.

- To avoid sending notifications, select **Ignore**.
- To send notifications on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

23. Specify the **Command** to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`

- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

```
/usr/sbin/sendmail -v recipient_email "subject_text"
```

- On Windows, to send a notification email, type the following command:
smtpmail -s subject -h mailserver recipient1@mailserver
recipient2@mailserver...

where:

- *-s subject*—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the smtpmail program assumes that the message contains a correctly formatted email header and nothing is added.
- *-h mailserver*—Specifies the hostname of the mail server to use to relay the SMTP email message.
- *recipient1@mailserver*—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

24. Click **NEXT**.
25. From the **Start Time** list, select one of the following options to specify the time to start the action. Use the spin boxes to set the hour and minute values.
 - **Set at Workflow level**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
26. (Optional) Select **Concurrent** to enable concurrent operations for the action.
27. From the **On Failure** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort Action**.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort Workflow**.
28. In the **Soft Limit** field, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
29. In the **Hard Limit** field, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
30. Click **NEXT**.
The Action Configuration Summary appears.
31. Review the settings that you have configured, and then click **FINISH**.

After you finish

(Optional) Create a clone action to automatically clone the save sets again after this clone action.

Create a probe action using the NetWorker Management Web UI

A probe action runs a user-defined script on a NetWorker client before the start of a backup. A user-defined script is any program that passes a return code. If the return

code is 0 (zero), then a client backup is required. If the return code is 1, then a client backup is not required.

Before you begin

Create the required policy and workflow.

Procedure

1. Select **Protection > Policies**.

The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.

2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow with the appropriate group association, or create a new workflow, and click **ACTION**.
4. From the **ACTION** list, select **Probe** to create a probe action.
The Create Action page appears.
5. Under **Basic Configuration**, in the **Name** field, type the name of the action, or use the default name.
The maximum number of characters for the action name is 64.
6. In the **Description** field, type a description for the action.
7. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the policy or workflow that contains the action is started, clear this option. By default, this option is selected.
8. To start the backup only if all the probes associated with client resources in the assigned group succeed, select **Start backup after all probes succeed**. By default, this option is selected.
9. Click **NEXT**.
10. Select a schedule configuration:

- **Select**—Allows you to select a pre-defined schedule that is applicable to the selected action type, or create a new schedule. Click **ADD SCHEDULE** to create a new schedule. The section [Create a schedule using the NetWorker Management Web UI](#) provides more information on creating schedules.

 **Note:** For probe actions, you can only select or create an execute type schedule.

- **Define**—Allows you to define a schedule.

For **Define**, do the following:

- a. Specify a weekly or monthly recurrence schedule for the action.
 - To specify a schedule for each day of the week, select **Weekly**.
 - To specify a schedule for each day of the month, select **Monthly**.
- b. Click on each day to specify the type of backup to perform.

The backup levels includes the following:

- **Execute**—Execute the action on the specified day.
- **Skip**—Skip the action on the specified day.

To perform the same type of backup on each day, click , and select one of the following:

- Make All Execute
- Make All Skip

c. (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

 **Note:** If you select this option, you must configure the overrides using either fixed dates or recurring dates.

d. Click **NEXT**.

11. Under **Schedule Overrides**, do the following:

a. In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

b. To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

 **Note:**

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.
- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.
- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

12. Click **NEXT**.

13. Under **Notifications**, from the **Notify** list, select an appropriate notification option.

- To avoid sending notifications, select **Ignore**.
- To send notifications on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

14. Specify the **Command** to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver
recipient2@mailserver...`

where:

- *-s subject*—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- *-h mailserver*—Specifies the hostname of the mail server to use to relay the SMTP email message.
- *recipient1@mailserver*—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

15. Click **NEXT**.
16. From the **Start Time** list, select one of the following options to specify the time to start the action. Use the spin boxes to set the hour and minute values.
 - **Set at Workflow level**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
17. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. The default value is 0.
18. From the **On Failure** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort Action**.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort Workflow**.
19. In the **Retries** field, specify the number of times that NetWorker must retry a failed probe or backup action, before NetWorker considers the action as failed. If the Retries value is 0, NetWorker does not retry a failed probe or backup action. The default value is 1.

20. In the **Retry Delay** field, specify a delay in seconds to wait before retrying a failed probe or backup action. When the Retry Delay value is 0, NetWorker retries the failed probe or backup action immediately. The default is 1 second.
21. In the **Inactivity Timeout** field, specify the maximum number of minutes that a job run by an action can try to respond to the server.

If the job does not respond within the specified time, the server considers the job a failure and NetWorker retries the job immediately to ensure that no time is lost due to failures. The default is 30 minutes.

Increase the timeout value if a backup consistently stops due to inactivity. Inactivity might occur for backups of large save sets, backups of save sets with large sparse files, and incremental backups of many small static files.
22. In the **Soft Limit** field, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
23. In the **Hard Limit** field, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
24. Click **NEXT**.

The Action Configuration Summary appears.
25. Review the settings that you have configured, and then click **FINISH**.

After you finish

(Optional) Create a traditional or snapshot backup action.

Create a check connectivity action using the NetWorker Management Web UI

A check connectivity action tests the connectivity between the clients and the NetWorker server, usually before another action such as a backup occurs.

Before you begin

Create the required policy and workflow.

Procedure

1. Select **Protection > Policies**.

The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.
2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow with the appropriate group association, or create a new workflow, and click **ACTION**.
4. From the **ACTION** list, select **Check Connectivity** to create a check connectivity action.

The Create Action page appears.
5. Under **Basic Configuration**, in the **Name** field, type the name of the action, or use the default name.

The maximum number of characters for the action name is 64.
6. In the **Description** field, type a description for the action.
7. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the

policy or workflow that contains the action is started, clear this option. By default, this option is selected.

8. To specify that the connectivity check is successful only if the connectivity test is successful for all clients in the assigned group, select **Succeed only after all clients succeed**. To specify that the connectivity check is successful if the connectivity test is successful for one or more clients in the assigned group, clear this option.
9. Click **NEXT**.
10. Select a schedule configuration:
 - **Select**—Allows you to select a pre-defined schedule that is applicable to the selected action type, or create a new schedule. Click **ADD SCHEDULE** to create a new schedule. The section [Create a schedule using the NetWorker Management Web UI](#) provides more information on creating schedules.

 **Note:** For check connectivity actions, you can only select or create an execute type schedule.
 - **Define**—Allows you to define a schedule.

For **Define**, do the following:

- a. Specify a weekly or monthly recurrence schedule for the action.
 - To specify a schedule for each day of the week, select **Weekly**.
 - To specify a schedule for each day of the month, select **Monthly**.

- b. Click on each day to specify the type of backup to perform.

The backup levels include the following:

- **Execute**—Execute the action on the specified day.
- **Skip**—Skip the action on the specified day.

To perform the same type of backup on each day, click , and select one of the following:

- Make All Execute
- Make All Skip

- c. (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

 **Note:** If you select this option, you must configure the overrides using either fixed dates or recurring dates.

- d. Click **NEXT**.

11. Under **Schedule Overrides**, do the following:

- a. In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

- b. To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

 Note:

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.
- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.
- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

12. Click **NEXT**.

13. Under **Notifications**, from the **Notify** list, select an appropriate notification option.

- To avoid sending notifications, select **Ignore**.
- To send notifications on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

14. Specify the **Command** to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver recipient2@mailserver...`

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the

smtpmail program assumes that the message contains a correctly formatted email header and nothing is added.

- *-h mailserver*—Specifies the hostname of the mail server to use to relay the SMTP email message.
- *recipient1@mailserver*—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

15. Click **NEXT**.
16. From the **Start Time** list, select one of the following options to specify the time to start the action. Use the spin boxes to set the hour and minute values.
 - **Set at Workflow level**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
17. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. The default value is 8.
18. From the **On Failure** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort Action**.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort Workflow**.
19. In the **Retries** field, specify the number of times that NetWorker must retry a failed probe or backup action, before NetWorker considers the action as failed. The default value is 3.
20. In the **Retry Delay** field, specify a delay in seconds to wait before retrying a failed probe or backup action. When the Retry Delay value is 0, NetWorker retries the failed probe or backup action immediately. The default is 1 second.
21. In the **Soft Limit** field, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
22. In the **Hard Limit** field, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
23. Click **NEXT**.
The Action Configuration Summary appears.
24. Review the settings that you have configured, and then click **FINISH**.

After you finish

(Optional) Create one of the following actions after the check connectivity action:

- Probe
- Traditional backup
- Snapshot backup

Create an expire action using the NetWorker Management Web UI

The expire action removes all expired save sets from the client file index and marks the save sets as recyclable in the media database. Save sets expire when the retention period for the save set is exceeded. You can create an expiration action in an existing workgroup only after a server backup action.

Before you begin

Create the required policy and workflow. The expire action must be the first action in the workflow or you can create the expire action after a server backup action.

Procedure

1. Select **Protection > Policies**.

The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.

2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow that does not have any group associated with it, or create a new workflow, and click **ACTION**.
4. From the **ACTION** list, select **Expire** to create an expire action.

The Create Action page appears.

5. Under **Basic Configuration**, in the **Name** field, type the name of the action, or use the default name.

The maximum number of characters for the action name is 64.

6. In the **Description** field, type a description for the action.
7. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the policy or workflow that contains the action is started, clear this option. By default, this option is selected.
8. Click **NEXT**.
9. Select a schedule configuration:

- **Select**—Allows you to select a pre-defined schedule that is applicable to the selected action type, or create a new schedule. Click **ADD SCHEDULE** to create a new schedule. The section [Create a schedule using the NetWorker Management Web UI](#) provides more information on creating schedules.

 **Note:** For expire actions, you can only select or create an execute type schedule.

- **Define**—Allows you to define a schedule.

For **Define**, do the following:

- a. Specify a weekly or monthly recurrence schedule for the action.
 - To specify a schedule for each day of the week, select **Weekly**.
 - To specify a schedule for each day of the month, select **Monthly**.

- b. Click on each day to specify the type of backup to perform.

The backup levels include the following:

- **Execute**—Execute the action on the specified day.
- **Skip**—Skip the action on the specified day.

To perform the same type of backup on each day, click , and select one of the following:

- Make All Execute
- Make All Skip

- c. (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

 **Note:** If you select this option, you must configure the overrides using either fixed dates or recurring dates.

- d. Click **NEXT**.

10. Under **Schedule Overrides**, do the following:

- a. In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

- b. To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

 **Note:**

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.
- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.
- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

11. Click **NEXT**.

12. Under **Notifications**, from the **Notify** list, select an appropriate notification option.

- To avoid sending notifications, select **Ignore**.
- To send notifications on completion of the action, select **On Completion**.

- To send a notification only if the action fails to complete, select **On Failure**.

13. Specify the **Command** to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver
recipient2@mailserver...`

where:

- *-s subject*—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- *-h mailserver*—Specifies the hostname of the mail server to use to relay the SMTP email message.
- *recipient1@mailserver*—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

14. Click **NEXT**.
15. From the **Start Time** list, select one of the following options to specify the time to start the action. Use the spin boxes to set the hour and minute values.
 - **Set at Workflow level**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
16. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. The default value is 0.
17. From the **On Failure** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort Action**.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort Workflow**.

18. In the **Soft Limit** field, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
19. In the **Hard Limit** field, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
20. Click **NEXT**.
The Action Configuration Summary appears.
21. Review the settings that you have configured, and then click **FINISH**.

Create a discover action using the NetWorker Management Web UI

A discover action discovers snapshots on a NAS device that the current NetWorker server did not create. NetWorker then creates save set entries for the discovered snapshots on a NAS client in the media database. The discover action also compares the save sets in the media database for the NAS client against the actual snapshots stored on the client. If the media database entry does not have a corresponding snapshot on the NAS client, then the media database entry is removed.

Before you begin

Create the required policy and workflow. Ensure that the group assigned to the workflow contains only NAS devices.

Procedure

1. Select **Protection > Policies**.
The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.
2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow with the appropriate group association, or create a new workflow, and click **ACTION**.
4. From the **ACTION** list, select **Discover** to create a discover action.
The Create Action page appears.
5. Under **Basic Configuration**, in the **Name** field, type the name of the action, or use the default name.
The maximum number of characters for the action name is 64.
6. In the **Description** field, type a description for the action.
7. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the policy or workflow that contains the action is started, clear this option. By default, this option is selected.
8. By default, in the **Discover Type** list, **NAS Snapshot** is selected.
9. Click **NEXT**.
10. Select a schedule configuration:
 - **Select**—Allows you to select a pre-defined schedule that is applicable to the selected action type, or create a new schedule. Click **ADD SCHEDULE** to create a new schedule. The section [Create a schedule using the NetWorker Management Web UI](#) provides more information on creating schedules.

 **Note:** For discover actions, you can only select or create an execute type schedule.

- **Define**—Allows you to define a schedule.

For **Define**, do the following:

- Specify a weekly or monthly recurrence schedule for the action.
 - To specify a schedule for each day of the week, select **Weekly**.
 - To specify a schedule for each day of the month, select **Monthly**.

- Click on each day to specify the type of backup to perform.

The backup levels include the following:

- **Execute**—Execute the action on the specified day.
- **Skip**—Skip the action on the specified day.

To perform the same type of backup on each day, click , and select one of the following:

- Make All Execute
- Make All Skip

- (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

 **Note:** If you select this option, you must configure the overrides using either fixed dates or recurring dates.

- Click **NEXT**.

- Under **Schedule Overrides**, do the following:

- In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

- To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

 **Note:**

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.
- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.

- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

12. Click **NEXT**.

13. Under **Notifications**, from the **Notify** list, select an appropriate notification option.

- To avoid sending notifications, select **Ignore**.
- To send notifications on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

14. Specify the **Command** to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver recipient2@mailserver...`

where:

- **-s subject**—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- **-h mailserver**—Specifies the hostname of the mail server to use to relay the SMTP email message.
- **recipient1@mailserver**—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

15. Click **NEXT**.

16. From the **Start Time** list, select one of the following options to specify the time to start the action. Use the spin boxes to set the hour and minute values.

- **Set at Workflow level**—Do not enforce an action start time. The action will start at the time defined by the workflow.

- **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
17. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. The default value is 0.
 18. From the **On Failure** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort Action**.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort Workflow**.
 19. In the **Soft Limit** field, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
 20. In the **Hard Limit** field, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
 21. Click **NEXT**.
The Action Configuration Summary appears.
 22. Review the settings that you have configured, and then click **FINISH**.

After you finish

(Optional) Create a generate index action to create entries in the client file indexes for the snapshot save sets, or a clone action to automatically clone the save sets after the discover action.

Create a generate index action using the NetWorker Management Web UI

You can use the NetWorker Management Web UI to create a generate index action.

Before you begin

Create the required policy and workflow.

Procedure

1. Select **Protection > Policies**.
The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.
2. Click on an available policy to create a new workflow or view existing workflows.
3. Click on an existing workflow with the appropriate group association, or create a new workflow, and click **ACTION**.
4. From the **ACTION** list, select **Generate Index** to create a generate index action.
The Create Action page appears.
5. Under **Basic Configuration**, in the **Name** field, type the name of the action, or use the default name.
The maximum number of characters for the action name is 64.
6. In the **Description** field, type a description for the action.

7. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the policy or workflow that contains the action is started, clear this option. By default, this option is selected.
8. Select the **Saveset type**. The NAS Snapshot backup is the only supported type.
9. Click **NEXT**.
10. In the **Filter Savesets** section, you can do the following:
 - a. To define the criteria that NetWorker uses to create the list of eligible save sets to clone, select **Define Filters**. The eligible save sets must match the requirements that are defined in each filter. NetWorker provides the following filter options:
 - **Time** —In the Time section, specify the time range in which NetWorker searches for eligible save sets to clone in the media database. Use the spin boxes to specify the start time and the end time. The Time filter list includes the following options to define how NetWorker determines save set eligibility, based on the time criteria:
 - **Accept**—The clone save set list includes save sets that are saved within the time range and meet all the other defined filter criteria.
 - **Reject**—The clone save set list does not include save sets that are saved within the time range and meet all the other defined filter criteria.
 - **Backup Levels** —In the Backup Levels section, specify the backup levels that you want to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The Levels filter list includes the following options, which define how NetWorker determines save set eligibility, based on the backup level filter criteria:
 - **Accept**—The clone save set list includes eligible save sets with the selected backup levels.
 - **Reject**—The clone save set list does not include eligible save sets with the selected backup levels.
 - **Save Sets** —In the Save Sets section, specify whether to include or exclude ProtectPoint and Snapshot save sets, when NetWorker searches for eligible save sets to clone in the media database. The Save Set filter list includes the following options, which define how NetWorker determines save set eligibility, based on the save set filter criteria:
 - **Accept**—The clone save set list includes eligible ProtectPoint save sets or Snapshot save sets, when you enable the ProtectPoint checkbox or Snapshot checkbox.
 - **Reject**—The clone save set list does not include eligible ProtectPoint save sets and Snapshot save sets when you enable the ProtectPoint checkbox or Snapshot checkbox.
 - **Clients** —In the Client section, specify a list of clients to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The Client filter list includes the following options, which define how NetWorker determines save set eligibility, based on the client filter criteria:
 - **Accept**—The clone save set list includes eligible save sets for the selected clients.
 - **Reject**—The clone save set list does not include eligible save sets for the selected clients.

11. If you do not want to define a filter criteria, select **Do Not Filter**, and click **NEXT**.
12. Select a schedule configuration:
 - **Select**—Allows you to select a pre-defined schedule that is applicable to the selected action type, or create a new schedule. Click **ADD SCHEDULE** to create a new schedule. The section [Create a schedule using the NetWorker Management Web UI](#) provides more information on creating schedules.

Note: For generate index actions, you can only select or create an execute type schedule.
 - **Define**—Allows you to define a schedule.

For **Define**, do the following:

- a. Click on each day to specify the type of backup to perform.

The backup levels include the following:

- **Execute**—Execute the action on the specified day.
- **Skip**—Skip the action on the specified day.

To perform the same type of backup on each day, click , and select one of the following:

- Make All Execute
- Make All Skip

- b. (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

Note: If you select this option, you must configure the overrides using either fixed dates or recurring dates.

- c. Click **NEXT**.

13. Under **Schedule Overrides**, do the following:

- a. In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

- b. To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

Note:

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.

- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.
- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

14. Click **NEXT**.

15. Under **Notifications**, from the **Notify** list, select an appropriate notification option.

- To avoid sending notifications, select **Ignore**.
- To send notifications on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

16. Specify the **Command** to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the Protecting virtual machines `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:
`nsrlog -f policy_notifications.log`
- On Linux, to send an email notification, type the following command:
`mail -s subject recipient`
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:
`/usr/sbin/sendmail -v recipient_email "subject_text"`
- On Windows, to send a notification email, type the following command:
`smtpmail -s subject -h mailserver recipient1@mailserver recipient2@mailserver...`

where:

- **-s subject**—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- **-h mailserver**—Specifies the hostname of the mail server to use to relay the SMTP email message.
- **recipient1@mailserver**—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

17. Click **NEXT**.

18. From the **Start Time** list, select one of the following options to specify the time to start the action. Use the spin boxes to set the hour and minute values.
 - **Set at Workflow level**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
19. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. The default value is 0.
20. From the **On Failure** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort Action**.
 - To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort Workflow**.
21. In the **Soft Limit** field, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
22. In the **Hard Limit** field, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
23. Click **NEXT**.
The Action Configuration Summary appears.
24. Review the settings that you have configured, and then click **FINISH**.

After you finish

(Optional) Create a clone action.

Edit an action using the NetWorker Management Web UI

To modify an action, perform the following:

Procedure

1. Select **Protection > Policies**.
The available data protection policies that you can use appears. The details of the selected policy appear in the right pane.
2. Click on an available policy to view existing workflows.
3. Click on an existing workflow to view the actions associated with the workflow.
4. Select the action that you want to modify, and click **EDIT**.
5. Make the required changes, and click **SAVE**.
The action is successfully updated.

Create a rule using the NetWorker Management Web UI

You can use the NetWorker Management Web UI to create a rule.

About this task

To create a rule, perform the following:

Procedure

1. Select **Protection > Rules**.

Rules that have already been created appears. The details of the rule such as the datasource type, condition, and usage also appears in the right pane.

2. Click **ADD**.

The **Add Rules** wizard appears.

3. In the **Name** field, type a name for the rule.
4. Select the **Datasource Type** from the drop-down. The default Datasource Type is VMware.
5. (Optional) In the **Description** box, you can specify more information about the rule.
6. Specify a matching condition. You can select **All** as the match type, if the item has to meet all of the rules criteria or select **Any** to include the item, if the item meets any of the criteria.
 - a. Specify the VMWare object type. You can select one of the following:
 - Virtual Machine
 - ESXi Host/Cluster
 - Virtual App
 - Virtual Machine Folder
 - Datacenter
 - Resource Pool
 - b. Specify the object type properties that the rule uses to determine a match. It can be the object's name, path, or tag. The available properties depend on the object type.
 - Virtual Machine - name, vSphere tag
 - ESXi Host/Cluster - path, vSphere tag
 - Virtual App - name, vSphere tag
 - Virtual Machine Folder - name, path, vSphere tag
 - Datacenter - name, path, vSphere tag
 - Resource Pool - path, vSphere tag
 - c. Select an operator to further define how a match is made based on the selected object type property. The operator value can be one of the following:
 - equals
 - not equals
 - contains
 - not contains
 - starts with
 - does not start with
 - ends with
 - does not end with

- regular expression
- d. Click the **Browse** icon to select the vCenter server and vSphere tag and click **OK** to exit the dialog.

Note: This option is available only if you select the vSphere tag property type in the definition.
 - e. Click the + icon to add a rule definition, and click the - icon to remove a rule definition.
7. Click **Create**.

The Rule is successfully created.
 8. Repeat steps 1 through 7 for any additional rules that you want to create.

Note: You can associate a rule to a group, if dynamic selection is enabled when creating groups. The section [Create a VMware group using the NetWorker Management Web UI](#) provides more information on creating groups.

Edit a rule using the NetWorker Management Web UI

To edit a rule, perform the following:

Procedure

1. Select **Protection > Rules**.

Rules that have already been created appears. The details of the rule such as the datasource type, condition, and usage also appears in the right pane.
2. Select the rule that you want to modify, and click **EDIT**.
3. Make the required changes, and click **SAVE**.

The rule is successfully updated.

- Note:**

 - You cannot change the name of a rule.
 - If you want to delete a rule, select the rule that you want to delete, and click **Delete**. However, a rule cannot be deleted, if it is associated with a group.

Troubleshooting policies

This section provides information about issues related to the configuration and management of policy resources.

Remote system error - Cannot assign requested address

This message appears intermittently when a single workflow has more 2000 save sets and the backup and clone operations occur concurrently. In this scenario the number RPC connections that the configuration requires exceeds the available number of RPC ports. To resolve this issue, split the workflow with a large number of save sets into multiple workflows, up to a maximum of 2000 save sets in each workflow and stagger the workgroup start times by 30 minutes.

Unable to start because the Group for this workflow is empty

This message appears when you use the Start Individual Client option to start actions for specific clients in the Server backup workflow. NetWorker does not support the Start Individual Client option for the Server backup workflow. To resolve this issue, start all actions for all the clients in the workflow.

Running actions from the command line

NetWorker 19.1 and later provide you with the ability to run actions from a command line for debugging purposes only.

To debug an action, use the action binary, for example, `nsrworkflow`, `nsrpolicy`, `savegrp`, or `nsrnassnap_index` with the following options:

- `--policy_name`—Specifies the name of the policy that contains the action. This option is required.
- `--workflow_name`—Specifies the name of the workflow that contains the action. This option is not required when a policy only contains one workflow.
- `--action_name`—Specifies the name of the action. This option is not required when a workflow only contains one action.
- `-Z action_type`—Required for the `savegrp` binary. Specifies the action type of the action. Supported values are `backup:traditional`, `backup:snapshot` and `probe`. If you do not specify this option, `savegrp` defaults to the `backup:traditional` action type.
- `--driven_by_action`—Specifies the source of the input work items for an action, for example a list of backup save set. Sources include one of the following options:
 - `jobid`—Specifies the jobid of the driving action.
 - `stdin`—Instructs the action binary to read the items from stdin.
 - `file:absolute_path_to_file`—Instructs the action binary to read the items from a file.



Note: This option is only required when the action is not the first action in a workflow.

Example 7 Debugging an action by using stdin

In the following example, a backup of the save set `/baz` failed for host `foo.com`. The name of traditional backup action for the save set is `backup`. A workflow named `traditional1`, which is in a policy named `Backup` contains the action.

To troubleshoot the backup action, perform the following steps:

1. Connect to the NetWorker Server with an administrator account.
2. From a command prompt, start the `nsradmin` program:

```
nsradmin
```

3. From the `nsradmin` prompt, define the attributes that `nsradmin` will display for a resource, for example, the resource name and the save set value, by typing the following command:

```
show name; save set
```

4. Enable `nsradmin` to display the hidden resource ID attribute for the NetWorker resources:

Example 7 Debugging an action by using stdin (continued)

```
option resource id
```

5. Display a list of client resources, by typing the following command:

```
p type: nsr client
```

6. From the output, record the resource identifier that appears for the client resource that contains the save set associated with the action that you want to debug. For example, output similar to the following appears:

```
name:foo.com;
saveset:/baz;
resourceidentifier:70.0.77.10.0.0.0.208.36.124.87.128.222.
109.22(1);
```

```
name:foo.com;
saveset:/foo,/bar;
resourceidentifier:93.0.89.114.0.0.0.55.25.124.87.128.222.
109.22;(9)
```

 **Note:** The resource ID does not include the brackets or the number contained within the brackets.

7. Use the `savegrp` command and the resource ID to start the action:

```
echo resource_ID|savegrp --policy_name=policy_name --
workflow_name=workflow_name --action_name=action_name -v --
driven_by_action=stdin
```

For example:

```
echo 93.0.89.114.0.0.0.55.25.124.87.128.222.109.22|savegrp --
policy_name=Backup --
workflow_name=traditional1 --action_name=backup -v --
driven_by_action=stdin
```

CHAPTER 5

Backup Options

This chapter contains the following topics:

- [Overview of resources that support backups](#).....356
- [Save sets](#).....356
- [Backup levels](#).....359
- [Backup scheduling](#).....372
- [Managing Schedules using NetWorker Management Web UI](#).....383
- [Backup retention](#).....386
- [General backup considerations](#).....389
- [Directives](#).....396

Overview of resources that support backups

NetWorker provides you with resources that enable you to customize what data is in the backup, when the backup occurs, and how the backup occurs.

The following table summarizes each supporting resource. Many of the resources require planning and configuration on the NetWorker server or on the client itself before the backup occurs.

Table 53 Resource overview

Resource	Description	Example
Backup levels	Defines whether to back up all data on the client, or only data that has changed.	Perform a full backup to back up all files, regardless of whether they have changed, or an incremental backup to back up only files that changed since the last backup.
Schedules	Defines the backup level to perform on each day.	Perform a full backup on Sunday, and an incremental backup on all other days of the week.
Time policies	Defines time periods. Use time policies to define save set retention. Save set retention is how long the save set entries are maintained in the media database and client file indexes.	Backups for a client are maintained in the database, and can be browsed for recovery for a month.
Directives	Specifies resources that contain special instructions that control how the NetWorker server processes files and directories during backup and recovery. For example, encryption and compression.	A directive specifies that the backup should skip files with a <code>.tmp</code> extension.

Save sets

The collection of data items that are backed up during a backup session between the NetWorker server and a Client resource is called a *save set*.

A save set can consist of the following:

- A group of files or entire file systems.
- Application data, such as a database, or operating system settings.

You can use the predefined save sets for scheduled backups, or specify a list of save sets to back up for a client resource in the **Save set** attribute on the **General** tab of the **Client Properties** dialog box.

Predefined save sets include the `DISASTER_RECOVERY:\` save set and the `ALL` save set.

When you specify a list of save sets for a client resource, the following guidelines apply:

- For Windows operating systems, use the same pathname case that the Windows file system uses. Although most file systems are case-independent, the NetWorker software cross-platform indexing system is case-sensitive. Always specify the Windows drive letter in uppercase.
- Place multiple entries on separate lines. For example, to back up a log file directory that is named `C:\Docs\CustomerLogs`, and all data that is contained in a directory that is named `D:\accounting`, type the following entries:

```
C:\Docs\CustomerLogs
D:\accounting
```

- For clients that use non-ASCII locales on UNIX platforms, or for Windows clients that are configured from a UNIX host that uses non-ASCII locales, special considerations apply when you type a path or file name in the **Save set** attribute:
 - Type the path or file name in the locale that was used when you created the path or file. If using a different locale when you type a path or file name, backups fail with a `No such file or directory` error message.
 - Either use the `ALL` save set in this situation, or log in to the client by using the correct locale and then configure the client from that computer.
- To back up a UNIX or Linux host that contains path or file names with multiple locales, create a separate Client resource for each locale. For example, to configure a multi-locale UNIX host with data in both Japanese and French, create two different Client resources. One Client resource to define the save sets for the Japanese data, and one Client resource to define the save sets for the French data.

The ALL save set

The `ALL` save set is the default save set when you create a Client resource.

Save sets included in the ALL save set

The following table provides a list of the save sets that are in the `ALL` save set for supported operating systems.

Table 54 Data in the ALL save set

Operating system	Files
Windows	• <code>DISASTER_RECOVERY:\</code> • Noncritical volumes
Mac OS X	All local and mounted volumes
UNIX	• When the backup starts, the <code>savefs</code> process reads the contents of the <code>/etc/vfstab</code> file on Solaris clients, the <code>/etc/fstab</code> file on HP-UX and Linux clients, or the <code>/etc/filesystems</code> file on AIX clients. The contents of the file are compared to the currently mounted file systems and BTRFS sub-volumes. Only

Table 54 Data in the ALL save set (continued)

Operating system	Files
	currently mounted file systems and BTRFS sub-volumes that are configured in these files are backed up. When NetWorker encounters a sub-directory that has a sub-volume ID that differs from the parent sub-volume ID, NetWorker will not backup the contents of the subdirectory, unless you specify the <i>save -x</i> in the Backup command field in the properties of the Client resource. • For a Solaris sparse or whole root zone client, all mounted file systems in the sparse or whole root zone that are not normally skipped, such as NFS, are backed up. • ZFS file systems are backed up. • If the save set name includes a symbolic link, a save set recovery is not supported.

Save sets excluded from the ALL save set

The following directories, file systems, and files are excluded from the ALL save set:

Table 55 File systems excluded from the ALL save set

• hsfs	• sharefs	• dfs	• binfmt_mi	• nucam
• proc	• nfs2	• autofs	sc	• fdfs
• fd	• nfs3	• iso9060	• usbfs	• xx
• cachefs	• nfs3perf	• udf	• devpts	• none
• lofs	• profs	• sysfs	• smbfs	
• mntfs	• nfs4	• debugfs	• swap	
• ctfs	• nfs	• subfs	• tmp	
• objfs	• brfs	• usbdevfs	• tmpfs	
			• nucfs	

NOTICE When you use the ALL save set for a backup, the NetWorker software creates a temporary file similar to a directive under each drive. The file name uses the format *drive guid.txt* and lists the files that are excluded from the backup. The file is temporary and is automatically deleted when the backup completes.

Keywords for scheduled file system backups

You can use special keywords with the ALL save set to define the file systems to include in a backup. The following table provides a list of the special ALL save sets and the backup behavior.

Table 56 Special ALL save sets

Special ALL save set syntax	Backup behavior
<code>all-file_system</code>	- Only back up locally mounted file systems of a particular type, where <i>file_system</i> is <code>zfs</code>, <code>ntfs</code>, <code>btrfs</code>, or <code>ext3</code>. For example: <code>all-zfs</code> backs up all locally mounted ZFS file systems on a Solaris host. <code>all-btrfs</code> backs up all mounted BTRFS sub-volumes that appear in the <code>/etc/fstab</code> file. - File systems such as NFS that are normally skipped are still skipped. - The <i>NetWorker E-LAB Navigator</i> provides a list of the supported file system for each operating system.
<code>all-mounts</code>	- On UNIX clients, back up all currently mounted file systems. - On Windows clients, the <code>all-mounts</code> save set is equivalent to the ALL save set. - File systems such as NFS that are normally skipped are still skipped.
<code>all-local</code>	- For a global zone client, the file systems in the sparse or whole root zone on the physical host are backed up. File systems in the global zone are skipped. - For a sparse or whole root zone client, the <code>all-local</code> save set is equivalent to the ALL save set.
<code>all-global</code>	- For a global zone client, all file systems in the global zone are backed up. All sparse and whole root zone file systems on the physical host are skipped. - For a Solaris sparse or whole root zone client, the <code>all-global</code> save set is equivalent to the ALL save set.

Backup levels

You can specify the level of the backup to be performed during scheduled backups.

When you limit the frequency of full backups, you help maintain server efficiency while still ensuring that data is protected. Different backup levels enable you to balance the

amount of time that is required to complete a backup with the number of volumes that are required to recover from a disk failure.

The following table describes the available backup levels.

Table 57 Backup levels

Backup level	Function
Full	Results in a back up of all files, regardless of whether the files have changed.
Incremental	Results in the back up of the files that have changed since the last backup, regardless of the level of the last backup.
Cumulative incremental	Results in the back up of all files that have changed since the last full backup.
Logs only	Results in the back up of the transaction log for databases that are created by a NetWorker module. For example, the NetWorker Module for Databases and Applications, the NetWorker Module for Microsoft, or the NetWorker Module for SAP.
Synthetic full	Results in the back up of all data that has changed since the last full backup and subsequent incremental backups, to create a synthetic full backup.
Skip	Skips the scheduled backup. For example, you can skip a backup on a holiday if no one is available to change or add more media volumes.

Comparing backup levels

Evaluate the advantages and disadvantages of each backup level to develop the backup strategy for an environment.

The following table lists key advantages and disadvantages of each backup level.

Table 58 Advantages and disadvantages of backup levels

Backup level	Advantages	Disadvantages
Full	Faster recovery	- Slower backups - High server load - High load on the client and network - Uses more volume space
Incremental	- Faster than a full backup - Low server load - Uses less volume space than a full backup	- Slow recovery - Data can spread across multiple volumes
Cumulative incremental	- Faster than a full backup - Low server load	Slow recovery

Table 58 Advantages and disadvantages of backup levels (continued)

Backup level	Advantages	Disadvantages
	• Uses the least amount of volume space	• Data can spread across multiple volumes
Logs only	• Faster than a full or incremental backup • Low server load	• Slow recovery • Data can spread across multiple volumes
Synthetic full	• Faster than a full backup • Faster recovery • Low load on the server, client, and network • Requires fewer volumes for recovery	• High load on the storage node • Requires at least two volume drives • Uses the most volume space

Review the following additional considerations when selecting backup levels:

- If you have only one stand-alone storage device and the full backup does not fit on a single piece of media, an operator must be available to monitor the backup, and change the media.
- Full backups cause the online indexes to grow more rapidly than incremental or cumulative incremental backups.
- Cumulative incremental backups serve as checkpoints in schedules because they collect all the files that have changed over several days, or even weeks, into a single backup session.
- Synthetic full backups provide the same benefits at the same cost as full backups. The difference is that synthetic full backups are less taxing on the network and client because a new full backup is created from a previously created full or synthetic full backup and subsequent incremental backups.

Backup levels and data recovery requirements

The schedule and configuration of backup levels directly affects how long a recovery from a disk failure takes and how many backup volumes are needed for the recovery.

Plan the backup levels to minimize the number of volumes or the amount of disk space that is used to store the data. The fewer the number of volumes that are required to recover from a disk failure, the less the time that you require to restore the data.

Note: You can also reduce the size and the time it takes to back up data by using directives. For example, use a directive to skip certain files or file systems when performing a backup.


The following example illustrates how the backup levels affect the requirements for data recovery.

In the following figure:

- Day 1—A full backup is run.
- Day 2— An incremental backup saves all files that have changed since the full backup.
- Day 3—Another incremental backup saves all files that have changed since Day 2.

- Day 4—A cumulative incremental backup saves all files that have changed since the full backup on Day 1.

Figure 39 Incremental and cumulative incremental backup levels



To recover all data from a disk failure on Day 4, you need the data from the full backup from September 30 and the cumulative incremental backup on Day 4. You no longer need the data from Day 1, 2, and 3, because the volume with the cumulative incremental backup includes that information.

Backup levels for the online indexes

The backup of the NetWorker server online indexes (client file index and media database) occur in a separate policy.

NetWorker automatically creates a server backup action in the Server Backup workflow of the Server Protection policy. By default, a full backup of the media database, resource files, and the NetWorker Authentication Service database occurs daily. A full backup of the client file indexes occur on the first day of the month. An incremental backup of the client file indexes occur on the remaining days of the month.

Synthetic full backups

A synthetic full backup combines a full backup and subsequent incremental backups to form a new full backup. A synthetic full is equivalent to a traditional full backup and can be used in all the same ways as a traditional full backup.

A synthetic full save set includes data that was backed up between the full backup and the last incremental backup. After a synthetic full backup occurs, the next synthetic full backup includes data that was backed up between the previous synthetic full backup, and subsequent incremental backups.

During a traditional full backup, client data is sent over the network to the NetWorker storage nodes, which can have a negative effect on client network performance. For synthetic full backups, however, the NetWorker software analyzes the full backup and subsequent incremental backups, extracts the most current versions of files, and then streams the data into a new full backup. Synthesizing the new full backup does not include the client machines and localizes the network traffic to the NetWorker server and storage nodes.

Performing synthetic full backups also reduces recovery time because the data is restored from the single synthetic full backup instead of from the last full backup and the incremental backups that follow it.

Synthetic full backups do not eliminate the requirement for full backups. It is recommended to perform full backups on a monthly or quarterly basis, and limit the number of incremental backups.


How a synthetic full backup is created

When a synthetic full backup operation starts, the NetWorker software performs an incremental backup of the save set and then adds that to the full and incremental backups that are already in place for the synthetic full process. Then the synthetic full backup occurs.

The following figure illustrates how a synthetic full backup is created.

Figure 40 Synthetic full backups

Save time



In this example, the synthetic full backup operation creates the incremental backup at T4. Then a synthetic full backup is created by combining the full backup at T1 with the subsequent incremental backups at T2, T3, and T4 to form a synthetic full backup at T4 + Delta. The save set at T4 + Delta is equivalent to a full backup that is taken at T4.

The T4 + Delta represents a small time change of one or two seconds from the time of T4, since two separate save sets cannot be assigned the exact same save set time. For example, if T4 is created at 1334389404, then T4+Delta is created at 1334389405, with a difference of one second.

The synthetic full save set includes only files that are covered by save sets up to T4 at 1334389404. The incremental backup after the synthetic full backup at 1334389405 includes all changes since 1334389404. Note that the synthetic full backup does not include the changes since T4, since only one save set can exist at any particular time.

After a synthetic full backup is performed, the next synthetic full backup combines the previous synthetic full backup and subsequent incremental backups.

When to use synthetic full backups

Synthetic full backups are supported only for backups of file system data with NetWorker 8.0 and later.

Synthetic full backups provide the most benefit in the following environments:

- The backup window is less than the amount of time it takes to perform a full backup.
- A client is at a remote location, and data transfer over the network to the server is a performance issue for either the network or the client.
- Network bandwidth is limited.
- Large backups over the network are cost-prohibitive.

Synthetic full backups include only the NetWorker server and storage node. If all the data is on a few storage nodes, then the network overhead for creating the synthetic full backup can be drastically reduced when compared to a traditional full backup of the same save sets.

NOTICE Under most conditions, synthetic full backups can free network bandwidth and client resources. However, a synthetic full backup might take longer to run on the storage node than a full backup because incremental backups are combined into a synthetic full backup. Without proper planning, synthetic full backups might affect the performance of the storage node.

To manage resource usage, perform synthetic full operations outside of the normal backup window. Also, synthetic full backups do not eliminate the requirement for full backups. It is best practice to schedule and perform full backups on a monthly or quarterly basis and limit the number of incremental backups.

Requirements for synthetic full backups

Ensure that the environment meets the requirements for synthetic full backups.

Save set requirements for synthetic full backups

All save sets participating in the construction of a synthetic full save set must meet the following requirements:

- Be file system save sets.
- Retain the same client name and save set name during the incremental and full backups that combine to form the synthetic full backup.
- Be browsable in the online index.
- Be created with NetWorker 8.0 or later.

Do not perform synthetic full backups with the following types of save sets:

- NDMP, SCSI, VCB, or snapshot save sets.
- Save sets that contain backups of raw disk file partitions.
- Save sets that contain database systems such as Microsoft Exchange and Oracle.
- Save sets where the backup command with `save` is not used.
- The **Save set** attribute for the client resource contains the `DISASTER_RECOVERY: \ save set` or the `ALL save set` on Windows.

When you use the `ALL` save set with synthetic full and virtual synthetic full backups, the noncritical volumes save successfully. However, critical volumes including `DISASTER_RECOVERY: \` are not backed up. The `nsrconsolidate()` command is unable to process the `DISASTER_RECOVERY: \ save set`. The client then runs a traditional full backup for the `DISASTER_RECOVERY: \ save set`.

Backups that are performed during a checkpoint restart might be in a synthetic full backup, if the other requirements for synthetic full backups are met.

For UNIX clients, include the forward slash to designate root (/) when specifying a save set name for the client resource. Otherwise, the synthetic full backup fails. For example, specify `/tmp` instead of `tmp`.

For Windows clients, include the backslash (\) when specifying a drive letter in a save set name for the client resource. Otherwise, the synthetic full backup fails. For example, specify `D:\` instead of `D:.`

Client resource configuration requirements for synthetic full backups

Ensure that the **Backup renamed directories** attribute is enabled on the **General** tab of the **Client Properties** dialog box for the Client resource. Select **View Diagnostic Mode** in the Administration interface to access the **Backup renamed directories** attribute in the **Client Properties** dialog box.

If you configure multiple policy workflows to run concurrently, set the **Parallelism** attribute to 40 for the Client resource for the NetWorker server. The **Parallelism** attribute is available on the **Globals (1 of 2)** tab of the **Client Properties** dialog box. Setting the attribute to 20 limits the number of concurrent synthetic full operations to 20. Divide the parallelism setting by two to control the number of concurrently running synthetic full operations. The best number of concurrent synthetic full operations depends on the following criteria:

- Configuration of the NetWorker server.
- Size of the save sets and number of clients.
- Number of `nsrpolicy` instances that are concurrently running.

Backup storage for synthetic full backups

Configure a Client resource for the NetWorker storage node that you use for the synthetic full backup. A client connection license for this storage node is not used if the storage node is not backed up.

There must be at least two available attached devices to perform a synthetic full backup: one for reading the backup data, and one for writing the backup data to a synthetic full backup.

You can store synthetic full backups on any device that can be used in a traditional full backup. However, since synthetic full backups include concurrent recover and save operations, it is strongly recommended that you direct synthetic full backups to devices that can perform concurrent operations, such as Data Domain devices or Advanced File Type Devices (AFTDs). Using these device types allows the NetWorker software to automatically handle volume contention, where the same volume is required for both reading and for writing simultaneously. These devices typically offer better performance.

You can use other devices such as tape drives, VTLs, and basic file devices as the destination for synthetic full backups, but careful preparation is required for the backup to succeed. The backup must be configured so that the destination volume does not contain any of the sources save sets that are used for the synthetic full backup. Also, for tape media, ensure that there are enough available drives to allow for concurrent recovery of the source data and for saving the synthetic full backup. Without careful planning, synthetic full backups to tape, VTL, or basic file devices might stall because of volume contention.

To direct a synthetic full backup to a dedicated pool, configure a separate backup action for synthetic full backups in the data protection policy, and select the pool as the destination pool in the backup action for the synthetic full backup.

Scheduling considerations for synthetic full backups

A synthetic full backup is resource intensive because it concurrently performs both recover and save operations. As a result, it is best to perform synthetic full operations outside of the normal backup window.

You can do this by creating separate workflows in a data protection policy for synthetic full backups. When using synthetic full backups, do not exceed the time interval of one month between traditional full backups.

To maintain current resource usage, which is defined as the space usage in the backup media and client file indexes, run synthetic full backups in place of traditional full backups. Running synthetic full backups more frequently than traditional backups are currently run results in the consumption of more space in the backup media and client file indexes.

For example, if a full backup occurs once a week, you can replace the full backup with an incremental backup followed by a synthetic full backup without increasing the backup space usage.

If you perform a full backup on Sunday and then incremental backups on Monday through Saturday, then consider changing to the following schedule:

- Full backup on the first Sunday of the month.
- Incremental backups on Monday through Saturday.
- Synthetic full backups on the second, third, fourth, and fifth Sunday of the month.

Support for directives with synthetic full backups

You can use the `compressasm` and `aes` (encryption) directives with synthetic full backups.

When using directives with synthetic full backups, consider the following:

- If directives were applied to save sets during the full and incremental backups that are part of the synthetic full backup, the synthetic full backup does not remove those directives.
- Any directives, including the `compressasm` and `aes` directives, that were applied to the full and incremental backups that are part of the synthetic full backup are not applied again.
- Do not use directives for synthetic full backups that are stored on a Data Domain device.
- Unsupported directives are ignored during a synthetic full backup.

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrconsolidate` command.

 **NOTICE** Directives do not apply to virtual synthetic full backups.

Recovery storage node selection for synthetic full backups

The storage node that is used for recovery depends on whether the required volume is mounted.

If the required volume is already mounted, then the storage node where the volume is mounted is used for recovering data.

If the required volume is not mounted, then the recovery storage node is selected based on the value in the **Recover storage node** attribute on the **Globals (2 of 2)** tab of the **Client Properties** dialog box for the Client resource. Select **View Diagnostic**

Mode in the Administration interface to access the **Recover storage node** attribute in the **Client Properties** dialog box.

Performing synthetic full backups

You can schedule synthetic full backups from the Administration window, or perform a manual incremental synthetic full backup from the command prompt.

Performing scheduled synthetic full backups

Perform scheduled synthetic full backups by configuring a data protection policy with a traditional backup action.

Procedure

1. Ensure that the environment meets the requirements that are provided in [Requirements for synthetic full backups](#) on page 364.
2. Create a group to define the clients for the synthetic full backups:
 - Create a basic client group to specify a static list of clients.
 - Create a dynamic client group to specify a dynamic list of Client resources.

When the backup starts, the NetWorker policy engine dynamically generates a list of Client resources that match the tags that are specified for the group.

Create separate groups for Windows clients and UNIX clients. Do not mix clients with different operating system types in the same group.
3. Create a policy.

Policies provide a container for the workflows, actions, and groups that support and define the backup action.
4. Create a workflow.

Workflows define the start time for a series of actions, the order of actions in a sequence, and the group of client resources for which the action occurs.
5. Use the **Policy Action** wizard to create a traditional backup action with the following settings:
 - In the schedule area of the **Choose Action Type** page, click the icon on each day to specify the type of backup to perform. The following icon indicates that a synthetic full backup will occur on the selected day:



- On the **Options** page, leave the **Verify synthetic full** option selected to verify the integrity of the new index entries that are created in the client file index for the synthetic full backup.
- On the **Options** page, leave the **Revert to full when synthetic full fails** option selected to perform a full backup of the save set if the synthetic full backup fails.

Performing manual synthetic full backups

Run the `nsrconsolidate` program from the command line of the NetWorker server to perform a manual synthetic full backup of a save set for a client.

Use the `-c` option to specify the client name, and the `-N` option to specify the save set name, with the `nsrconsolidate` command. You can also use the `-C` option to specify both the client and save set name together, the `-S` option to specify the save

set ID (instead of the save set name), and the `-t` and `-e` options to specify the start time and end time for the save set, respectively.

The value that you specify for a save set name, client name, file name, or directory name with `nsrconsolidate` for a Windows client is case-sensitive because the NetWorker software cross-platform indexing system is case-sensitive. A best practice is to always specify the Windows drive letter in uppercase.

When you run multiple `nsrconsolidate` commands, run fewer commands that include many save sets instead of multiple commands with fewer save sets. This strategy helps `nsrconsolidate` to manage the number of concurrent synthetic full operations and reduce resource usage. The best number of concurrent synthetic full operations depends on the following criteria:

- Configuration of the NetWorker server.
- Size of the save sets and number of clients.
- Number of `nsrpolicy` instances that are concurrently running.

The *NetWorker Command Reference Guide* or the UNIX man pages provide details on `nsrconsolidate`.

Validating synthetic full backups

You can validate VSF backups by using the `mminfo` command, the **Media** window of the Administration interface, and the `savegrp` logs.

Validating synthetic full backups with the `mminfo` command

The following table lists the `mminfo` commands with applicable switches for validating synthetic full backups.

Table 59 `mminfo` commands for synthetic full backup validation

Command with switches	Description
<code>mminfo -aS</code>	Shows detailed information about synthetic full backups, including information about the save sets used to form the synthetic full backup.
<code>mminfo -q syntheticfull -c <i>client</i> -N <i>save_set</i></code>	Queries all synthetic full save sets for the specified <i>client</i> and <i>save_set</i> .

Validating synthetic full backups in the Media window of the Administration interface

When you search for save sets in the **Media** window of the Administration interface, you can limit the save set results to synthetic full save sets by selecting the **Synthetic Full** checkbox on the **Query Save Set** tab. [Searching for save sets](#) on page 533 provides instructions.

Validating synthetic full backups in the backup action logs

The following excerpt from the backup action log file illustrates the type of messages NetWorker displays when performing a synthetic full backup:

```
1707:97860:nsrconsolidate: Synthetic full save set hostname:/
sat-tree at savetime 1358188522 was created by using non-
virtual synthetic mode
95773:nsrrecovery: Virtual synthetic succeeded for hostname:/
test1
```

Synthetic full backup reporting

The backup statistics and backup status reports provide details on synthetic full backups. A value of `Synthetic` in the `Type` column for the `Save Sets Details` report or the `Save Sets Details by client` report indicates that the backup is a synthetic full backup. [Enterprise data reporting](#) on page 642 provides more information.

Virtual synthetic full backups

A virtual synthetic full (VSF) backup is the same as a synthetic full backup, except that it is performed on a single Data Domain system.

Similar to synthetic full, VSF uses full and partial backups to create a full backup. However, since the backup occurs on a Data Domain system using DD Boost APIs, the backup does not require save set data to be sent over the network. The result is improved performance over synthetic full and traditional full backups.

The following table compares traditional synthetic full and virtual synthetic full backups.

Table 60 Comparison of traditional synthetic full and virtual synthetic full backups

Traditional synthetic full	Virtual synthetic full
Data is read from and written to volumes.	Data movement is limited within the same Data Domain system.
Read/write for all types of volumes is supported.	Only Data Domain devices are supported, and the source and destination volumes must belong to the same Data Domain system. However, the volumes can belong to different MTrees in the same Data Domain system.
The client file index is created by <code>nsrrecopy</code> .	The client file index is created by <code>nsrconsolidate</code> .
Client Direct support is not required.	Client Direct support is required.

Requirements for VSF backups

Ensure that the environment meets the requirements for virtual synthetic full (VSF) backups.

The following table lists the requirements for VSF backups.

Table 61 Requirements for virtual synthetic full backups

Requirement	Details
DDOS version	Version 5.3 or later for both Data Domain systems and Data Domain Archivers.
DD Boost version	Version 2.6 or later.
Data Domain system configuration	Enable the <code>virtual-synthetics</code> option on the Data Domain system. To verify that <code>virtual-synthetics</code> is enabled, log in

Table 61 Requirements for virtual synthetic full backups (continued)

Requirement	Details
	to the Data Domain system and type the following command: <pre>ddboost option show</pre> Ensure that a value of <code>enabled</code> appears next to the <code>virtual-synthetics</code> option in the output for the command. ① NOTICE If <code>virtual-synthetics</code> is disabled but all other requirements for VSF are met, then the VSF backup fails with errors. NetWorker does not perform a traditional synthetic full backup in this case.
Backup storage	All constituent backups for the VSF backup must be on the same Data Domain system. The save sets can be distributed across multiple storage nodes and located in different MTrees on the Data Domain system.
Client resource configuration	• Enable the Client direct attribute on the General tab of the Client Properties dialog box for the client resource. You must select ViewDiagnostic Mode in the Administration interface to access the Client direct attribute in the Client Properties dialog box. • Enable the Data Domain backup attribute on the Apps & Modules tab of the Client Properties dialog box for the client resource. • To ensure optimal backup performance, configure the client to backup 10 or fewer save sets.
Device resource configuration	Specify a value in the volume location attribute for the device resource for the Data Domain system. NetWorker updates the volume location attribute during the device mount operation. ① NOTICE Before you update a storage node that uses Data Domain devices, unmount each device. Once the update completes, mount each device.
NetWorker upgrade requirements	If you upgrade the NetWorker client to release 8.1 or later from a release before 8.1, you must perform a full backup before you

Table 61 Requirements for virtual synthetic full backups (continued)

Requirement	Details
	perform a VSF backup. Otherwise, file-by-file recovery fails.
Cloning requirements	The <code>virtual-synthetics</code> option must be enabled for Data Domain systems being used for cloning VSF backups. Otherwise, cloning fails.

The *NetWorker Data Domain Boost Integration Guide* provides details on configuring the NetWorker environment for use with a Data Domain system.

Support for directives

Directives do not apply to VSF backups because the VSF backup is created by the Data Domain system.

Support for concurrent operations

The volume of concurrent VSF operations that a Data Domain system can handle depends on the model of the Data Domain system and the capacity of the NetWorker host. The following scenarios have been tested and verified to work:

- Concurrent VSF backups.
- A VSF backup concurrent with a cloning operation.
- A VSF backup concurrent with clone-controlled replication.

Performing VSF backups

Procedure

1. Ensure that the environment meets the requirements for virtual synthetic full (VSF) backups.

If NetWorker detects that one or more of the requirements are not met, then a traditional synthetic full backup occurs instead.

2. Perform the backup:

- For scheduled backups, select the synthetic full backup level for the backup action in the data protection policy.

The procedure for scheduled VSF backups is the same as the procedure for scheduled traditional synthetic full backups. [Performing scheduled synthetic full backups](#) on page 367 provides more information on configuring a data protection policy for a scheduled synthetic full backup.

- For manual backups at the command line, use the `nsrconsolidate` command.

The procedure for manual VSF backups is the same as the procedure for manual traditional synthetic full backups. [Performing manual synthetic full backups](#) on page 367 provides more information.

Validating VSF backups

You can validate VSF backups by using the `mminfo` command, the **Media** window of the Administration interface, and the `savegrp` logs.

Validating VSF backups with the `mminfo` command

The following table lists the `mminfo` commands with applicable switches for validating VSF backups.

Table 62 `mminfo` commands for VSF backup validation

Command with switches	Description
<code>mminfo -aS</code>	Shows detailed information about synthetic full backups, including information about the save sets used to form the synthetic full backup.
<code>mminfo -q syntheticfull -c <i>client</i> -N <i>save_set</i></code>	Queries all synthetic full save sets for the specified <i>client</i> and <i>save_set</i> .

Validating VSF backups in the Media window of the Administration interface

When you search for save sets in the **Media** window of the Administration interface, you can limit the save set results to synthetic full and VSF save sets by selecting the **Synthetic Full** checkbox on the **Query Save Set** tab. [Searching for save sets](#) on page 533 provides instructions.

Validating VSF backups in the `savegrp` logs

The following excerpt from the policy log file illustrates the type of messages NetWorker displays when performing VSF backups or traditional synthetic full backups, or when performing a traditional synthetic full backup because the VSF backup requirements are not met:

```
1707:97860:nsrconsolidate: Synthetic full save set hostname:/
sat-tree at savetime 1358188522 was created by using non-
virtual synthetic mode
95773:nsrrecovery: Virtual synthetic succeeded for hostname:/
test1
```

Backup scheduling

When you schedule backups, you define the days on which backups occur and the level of backup (full, incremental, and so on) that occurs each day.

Migration of NSR schedule from 8.x to 19.1.1

NSR schedule resource migrated to 19.1.1 has activity type set as backup except for the default VMware incremental forever schedule in 8.x which has activity type as VMware. All NSR schedule associated with the NSR group resource in 8.x is retained when NSR group is converted to NSR protection policy action in 19.1.1. This means that backup action created from NSR group has the same association of schedule which NSR group had.

For scheduled clones, NSR schedule is embedded in 8.x and remains the same in 19.1.1. When they are converted to NSR protection policy action, schedule still remains


embedded. NSR task created for the same scheduled clone in 8.x is converted to NSR schedule in 19.1.1 and the name of that schedule is same as the name of the scheduled clone in 8.x.

Scheduling backup cycles

The period from one full backup to the next full backup is called a backup cycle.

For example, the default schedule for backups is a full backup on a client each Sunday, and incremental backups on the other days of the week, as illustrated in the following figure.


Figure 41 Default weekly backup schedule



Depending on the size of a network, you could perform full backups for all clients simultaneously. For example, if no one works over the weekend you could schedule full backups during this time.

Alternatively, you may need to configure backups to balance the backup load on and increase the efficiency of a NetWorker server. Since full backups transfer large amounts of data and typically take longer than other backup levels, you might want to stagger them throughout the week. For example, you could configure backups so that full backups occur for one group of clients on Sunday, for a second group of clients on Tuesday, and a third group of clients on Thursday, as illustrated in the following figure.

Figure 42 Staggered weekly backup schedule for multiple groups of clients



Note: Consider using a synthetic full backup in environments with a short backup window period when you must create a full backup.

Considerations for scheduling backups

Planning schedules for backups in an environment requires careful consideration of several factors.

For example:

- The amount of data you must back up.
- The number of backup media volumes to use.
- The amount of time available to complete a backup.
- The number of volumes that are required to recover from a disaster such as a disk failure.

Recovery considerations

You must also determine the requirements for recovering files. For example, if users expect to recover any version of a lost file that was backed up during a three-month period (that is, the retention setting is three months), then you must maintain all the backup volumes for a three-month period. However, if users expect to be able to recover data from only the last month, you do not need to maintain as many volumes.

Considerations for large client file systems

At a moderate backup rate of 400 KB per second, a full backup for a client with 10 GB of data takes about seven hours to complete. Performing a scheduled full backup for such large client save sets may not be convenient because of the amount of time required.

For large client file systems, consider scheduling consider separate backups for each of the client disk volumes. This strategy enables you to back up all the client files, but not all at once, which is less time-consuming than a full backup of all local data at one time.

To schedule separate backups of each client disk volume, configure multiple client resources for the client, and explicitly list one disk volume as the save set for each client resource. Add each client resource to a different group. Then configure separate policy workflows to back up each group on a different schedule.

NOTICE When you create explicitly list save sets, any files or file systems not in that list are omitted from the backup, including any new disk volumes that you add to the system. Remember to configure backups for any new disk volumes after you add them.

Methods for scheduling backups

You can configure the backup schedule for a group of clients as part of data protection policy settings, or you can configure schedule overrides.

Schedules and backup levels assigned to an action

You specify the schedule and backup level as part of the backup action by selecting either **Define** or **Select** options under **Select Schedule**. The following figure illustrates the default define weekly schedule for a traditional backup action, with a full backup on Sunday and incremental backups on the remaining days of the week.

Figure 43 Selecting the schedule

Policy Action Wizard

Specify the Action Information

Specify information about the action and define the schedule for the action operation. Select an action type, and if applicable the select action sub-type. To associate this action with a different workflow in the policy, select the workflow from the Workflow drop down. If the new action should occur after another action that is associated with the policy, select an existing action from the Previous drop-down. To start the new action at the same time as the action defined in the Previous attribute, select Concurrent.

Specify the Action Information

Identity

Name: backup

Comment:

Enabled: ☒

Action Type: Backup

Backup Subtype: Traditional

Force Backup Level:

Workflow

Policy: MyPolicy1

Workflow: MyWorkflow1

Previous:

Concurrent: ☐

Select Schedule

☐ Define

☒ Select

Default

Period

☒ Weekly by day (For example Sunday, Friday, etc.)

☐ Monthly by day (For example 1st day, 5th day etc.)

Click on the day to schedule the action Make All Full

Sun	Mon	Tue	Wed	Thu	Fri	Sat
						

You can also configure the schedule for a backup action on a monthly basis instead of on a weekly basis.

Click the icon in the schedule to change the backup level that is performed on that day. The following table provides details about the backup level that each icon represents.

Table 63 Scheduled backup level icons

Icon	Label	Description
	Full	Perform a full backup on this day. Full backups include all files, regardless of whether the files changed.
	Incr	Perform an incremental backup on this day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
	Cumulative Incr	Perform a cumulative incremental backup. Cumulative incremental backups include files that

Table 63 Scheduled backup level icons (continued)

Icon	Label	Description
		have changed since the last full backup.
	Logs Only	Perform a backup of only database transaction logs.
	Synthetic Full	Perform a synthetic full backup on this day. A synthetic full backup includes all data that changed since the last full backup and subsequent incremental backups to create a synthetic full backup.
	Skip	Do not perform a backup on this day.

The following figure illustrates the Predefined NSR schedule that can be associated with the action. If a level needs to be changed, it can be done by updating the NSR schedule resource:

Figure 44 Defining the schedule

Specify the Action Information

Specify information about the action and define the schedule for the action operation. Select an action type, and if applicable the select action sub-type. To associate this action with a different workflow in the policy, select the workflow from the Workflow drop down. If the new action should occur after another action that is associated with the policy, select an existing action from the Previous drop-down. To start the new action at the same time as the action defined in the Previous attribute, select Concurrent.

Identity

Name: backup

Comment:

Enabled: ☒

Action Type: Backup

Backup Subtype: Traditional

Force Backup Level:

Workflow

Policy: MyPolicy1

Workflow: MyWorkflow1

Previous:

Concurrent: ☐

Select Schedule

☒ Define

☐ Select

Default

Period

☒ Weekly by day (For example Sunday, Friday, etc.)

☐ Monthly by day (For example 1st day, 5th day etc.)

Click on the day to schedule the action

Make All Full

Sun	Mon	Tue	Wed	Thu	Fri	Sat

Configuring multiple backup levels for frequently scheduled backups

Use the **Force Backup Level** attribute in the **Specify the Action Information** window of the Action wizard to override the backup levels of a Traditional backup action that occurs multiple times in a 24 hour period.

For workflows that have more than one scheduled backup within a 24-hour period, use the **Force Backup Level** attribute to allow more than one backup to occur at two different backup levels in a 24-hour period. When you select a backup level in the **Force Backup Level** attribute, the first backup is performed at the scheduled backup level. Each subsequent occurrence of the backup action in the next 24 hours occurs at the level defined in the **Force Backup Level** attribute. For example, if the level defined by the schedule is Full and the **Force Backup Level** attribute is set to Incr, the first backup started by the action occurs at a level full and subsequent backups, within 24 hours of the start of the full backup are incremental. By default this option is cleared, which means that if the action runs multiple backup operations in a 24 period, all the backups occur at the scheduled backup level.

The following figure provides an example of the **Force Backup Level** attribute in the **Specify the Action Information** window, with the **Cumulative Incr** option selected.

Figure 45 The Force Backup Level attribute

Specify the Action Information

Specify information about the action and define the schedule for the action operation. Select an action type, and if applicable the select the workflow from the Workflow drop down. If the new action should occur after another action that is associated with the pc action at the same time as the action defined in the Previous attribute, select Concurrent.

The screenshot shows the 'Specify the Action Information' tab in the NetWorker configuration interface. On the left, there is a sidebar with navigation links: 'Specify the Action Information' (selected), 'Specify the Backup Options', 'Specify the Advanced Options', 'Action Configuration Summary', and 'Action Wizard Results'. The main area is divided into two sections: 'Identity' and 'Workflow'. In the 'Identity' section, the 'Name' field is 'backup to DD', 'Comment' is empty, 'Enabled' is checked, 'Action Type' is 'Backup', 'Backup Subtype' is 'Traditional', and 'Force Backup Level' is 'Cumulative Incr'. In the 'Workflow' section, 'Policy' is empty, 'Workflow' is 'Full', 'Previous' is empty, and 'Concurrent' is unchecked.

Defining a schedule for a client

NetWorker allows you to override the backup level for a schedule traditional backup action by configuring a schedule for a client.

NetWorker provides you with preconfigured schedules that you can assign to a client. Review the following sections for information about preconfigured schedules, how to modify a schedule, and how to assign a schedule to a client resource.

Preconfigured schedules

When you override the policy backup schedule for a client resource, you can select or customize one of the preconfigured schedules that are available when you install or upgrade the NetWorker software.

The following table describes the preconfigured schedules.

Table 64 Preconfigured NetWorker schedules

Schedule name	NetWorker backup operation
Default	Weekly schedule that performs a full backup every Sunday and incremental backups on all other days.
Forever Incremental	Monthly schedule that performs a synthetic full backup every day.
Full Every Day	Weekly schedule that performs a full backup every day.

Table 64 Preconfigured NetWorker schedules (continued)

Schedule name	NetWorker backup operation
Full Every Friday	Weekly schedule that performs a full backup every Friday and incremental backups on all other days.
Full on 1st Friday of Month	Monthly schedule that performs a full backup on the first Friday of the month and incremental backups on all other days. You cannot edit this schedule.
Full on 1st of Month	Monthly schedule that performs a full backup on the first calendar day of the month, and incremental backups on all other days.
Quarterly	Monthly schedule that performs a full backup on the first day of a quarter, a cumulative incremental backup once a week after the full backup, and then incremental backups on all other days.
Synthetic Full 1st Friday of Month	Monthly schedule that performs a synthetic full backup on the first Friday of every month, and incremental backups on all other days.
Synthetic Full Every Friday	Weekly schedule that performs a synthetic full backup on every Friday and incremental backups on all other days.
Synthetic Full on 1st of Month	Monthly schedule that performs a synthetic full backup on the first calendar day of the month, and incremental backups on all other days.
Synthetic Full Quarterly	Monthly schedule that performs a synthetic full backup on the first day of each quarter, a cumulative incremental backup once a week after the synthetic full backup, and then incremental backups on all other days.

You can edit all preconfigured schedules except for schedules that contain overrides, which are indicated by an asterisk next to a backup level in the schedule calendar. You cannot delete a preconfigured schedule.

Managing the schedule resource

Review this section for information about how to create, edit, copy, and delete schedule resources.

Creating a backup schedule Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Schedules**.
3. From the **File** menu, select **New**.

The **Create Schedule** dialog box appears.

4. In the **Name** box, type a name for the schedule.
5. Under **Select Schedule** option, when **Define** is selected, from Period List, select **Week** or **Month** to control whether the schedule repeats on a weekly or monthly basis.
6. Optional, specify a description of the schedule in the **Comment** box.
7. Set the backup level for each day by right-clicking the day, selecting **Set Level** and then the backup level.
8. From the **Activity type** list, select **Backup** or **VMware** or **Server Backup** or **Execute** to associate the schedule to the appropriate actions. For details on activity type, refer to table 65.
9. Optional, set the override backup level for a day by right-clicking the day, selecting **Override Level** and then the backup level.

For example, to prevent a full backup from running on a holiday, override the schedule so that the full backup runs on the day before or the day after the holiday. An asterisk (*) next to a backup level indicates that an override has been set for that day.

 **Note:** If you override backup levels by using the `nsradmin` command line program, you can also specify relative date values such as `full first friday every 2 week`. The `nsr_schedule` UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about overriding backup levels.

10. Click **OK**.

Editing a schedule

You can edit all custom schedules, and all preconfigured schedules, except for preconfigured schedules that contain overrides. Overrides are indicated by an asterisk next to a backup level in the schedule calendar. You can edit all schedule settings except for the name.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Schedules**.
3. In the right pane, right-click the schedule and select **Properties**.

The **Schedule Properties** dialog box appears.

4. Edit the settings for the schedule and click **OK**.

Copying a schedule

You can create a new backup schedule by copying an existing schedule and then editing the copy.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Schedules**.
3. In the right pane, right-click the schedule to copy and select **Copy**.

The **Create Schedule** dialog box appears with the same information as the copied schedule except for the name.

4. In the **Name** box, type a name for the new schedule.
5. Edit the settings for the schedule and click **OK**.

Deleting a schedule

You can delete any custom schedules that you have created. You cannot delete preconfigured schedules.

Before you begin

Ensure that the schedule has not been applied to any Client resources by verifying the setting in the **Schedule** list on the **General** tab of the **Client Properties** dialog box for each Client resource.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Schedules**.
3. In the right pane, right-click the schedule and select **Delete**.
A confirmation message appears.
4. Click **Yes**.

Types of schedule

The following table mentions the various schedule types, the corresponding activity type that needs to be provided during create, update, and delete of NSR schedule resource and the supported levels and the actions of the NSR schedule resource that can be used for.

Table 65 Types of schedule

Schedule type	Activity type (nsradmin)	Allowed backup levels	Allowed level string in nsradmin	Actions supported
Backup	Backup	full, incremental, cumulative, logs only, incremental synthetic full and skip	full, incr, 1, txnlog, inc_synth_full, skip	Backup actions-traditional and snapshot
VMware	VMware	full, incremental, logs only, and skip	full, incr, txnlog, skip	VMware (vproxy)
Server backup	Server backup	full, incremental, cumulative incremental and skip	full, incr, 1, skip	Server backup
Execute	Execute	execute, skip	exec, skip	clone, probe, discover, check-connectivity actions, expire

Note:

- The default activity type is backup and the default period is month.
- The default values for the action attribute are populated as per the above table.
- If the action attribute is being provided, then the period attribute also needs to be provided.

- The activity type can be set only during the create operation and cannot be updated.
- All NSR schedule associations need to be removed from the workflow or action, before it can be deleted.
- The create or update operation might fail if an invalid level is specified for a given activity type.

Table 66 Types of actions

Activity type	Default action values when period is month	Default action values when period is week
Backup	full incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr	full incr incr incr incr incr incr
VMware	incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr	incr incr incr incr incr incr incr
Server backup	full incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr incr	full incr incr incr incr incr incr
Execute	exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec exec	exec exec exec exec exec exec exec

Configuring a client to override the schedule assigned to an action

You can override the backup schedule that is specified in the data protection policies that apply to a client resource by specifying a schedule for the Client resource itself.

Procedure

1. (Optional) Create or customize the schedule that you plan to assign to the Client resource.
2. In the **Administration** window, select **View > Diagnostic Mode** to enable diagnostic mode view.
A check mark next to **Diagnostic Mode** in the **View** menu indicates that diagnostic mode view is enabled.

3. In the **Administration** window, click **Protection**.
4. In the expanded left pane, select **Clients**.
5. In the right pane, right-click the client resource and select **Modify Client Properties**.

The **Client Properties** dialog box appears, starting with the **General** tab.

6. Ensure that the **Scheduled Backup** checkbox is selected.
When the checkbox is clear, scheduled backups do not occur for the client.
7. From the **Schedule** list, select the schedule to use instead of the schedule in the data protection policies that apply to the Client resource.
8. Enable **Client determines level**.
9. Click **OK**.

Managing Schedules using NetWorker Management Web UI

You can manage the schedule and define the days on which backups occur and the level of backup (full, incremental, and so on) that occurs each day.

Create a schedule using the NetWorker Management Web UI

You can use the NetWorker Management Web UI to create a schedule.

Procedure

1. Select **Protection > Schedules**.
The existing schedule details appear. The details of the selected schedule appear in the right pane.
2. Click **ADD**.
3. Under **Basic Configuration**, in the **Name** field, type a name for the schedule.
4. In the **Description** field, type a description for the schedule.
5. From the **Type** list, select one of the following action types:
 - **Backup**—Select this option for backup, traditional, and snapshot actions.
 - **VMware**—Select this option for VMware backup action.
 - **Server Backup**—Select this option for server backup action.
 - **Execute**—Select this option for clone, probe, check connectivity, expire, discover, and generate index actions.
6. Click **NEXT**.
7. Specify a weekly or monthly schedule for the action.
 - To specify a schedule for each day of the week, select **Weekly**.
 - To specify a schedule for each day of the month, select **Monthly**.
8. Click on each day to specify the type of backup to perform.

The backup levels for NetWorker includes the following:

- Backup type
 - **Full**—Perform a full backup on the specified day. Full backups include all files, regardless of whether the files changed.

- **Incr**—Perform an incremental backup on the specified day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
- **Cumulative Incr**—Perform a cumulative incremental backup. Cumulative incremental backups include files that have changed since the last full backup.
- **Logs Only**—Perform a backup of only database transaction logs.
- **Incremental Synthetic Full**—Perform an incremental synthetic backup on the specified day. An incremental synthetic full backup includes all data that changed since the last full backup and subsequent incremental backups to create a synthetic full backup.
- **Skip**—Do not perform a backup on the specified day.
- VMware type
 - **Full**—Perform a full backup on the specified day. Full backups include all files, regardless of whether the files changed.
 - **Incr**—Perform an incremental backup on the specified day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
 - **Logs Only**—Perform a backup of only database transaction logs.
 - **Skip**—Do not perform a backup on the specified day.
- Server backup type
 - **Full**—Perform a full backup on the specified day. Full backups include all files, regardless of whether the files changed.
 - **Incr**—Perform an incremental backup on the specified day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
 - **Logs Only**—Perform a backup of only database transaction logs.
 - **Skip**—Do not perform a backup on the specified day.
- Execute type
 - **Execute**—Execute the action on the specified day.
 - **Skip**—Skip the action on the specified day.

To specify the same type of schedule on each day, click , and select one of the following that appears based on the selected schedule type:

- Make All Full
 - Make All Incremental
 - Make All Cumulative Incremental
 - Make All Logs Only
 - Make All Incremental Synthetic Full
 - Make All Execute
 - Make All Skip
9. (Optional) Select **Override Options** to configure overrides for the task that is scheduled on a specific day.

 **Note:** If you select this option, you must configure the overrides using either fixed dates or recurring dates.

10. Click **NEXT**.

11. Under **Override**, do the following:

- a. In the **Recurring Pattern** attribute, if the status is **Not Available**, click **ADD**. Otherwise, click **VIEW/EDIT** to specify recurring patterns. Then, select the backup level, and define the override schedule to occur on a specific day, week, month, quarter, or year.

The selected date is highlighted with a different color and an asterisk (*).

- b. To specify a fixed date pattern, select the month and year, and then click on each day to specify the backup level.

The selected date is highlighted with a different color.

 **Note:**

- You can select multiple override dates.
- In the case of a fixed date pattern, to clear an override schedule, click on the day that you want to clear the override for and select **Clear Selection**. The Clear Selection option is not applicable in the case of a recurring pattern.
- If the fixed and the recurring pattern is on the same day, the fixed pattern gets preference. However, if you clear the fixed override schedule, the recurring pattern is displayed.
- When creating overrides, you cannot select the previous month and year. You also cannot create fixed overrides for the days before the current day because the option to change the level is disabled.
- When editing overrides, you can clear existing fixed overrides for the days before the current day. However, you cannot set any new overrides because the option to change the level is disabled.

12. Click **FINISH**.

Edit a schedule using the NetWorker Management Web UI

To modify a schedule, perform the following:

Procedure

1. Select **Protection > Schedules**.

The existing schedule details appear. The details of the selected schedule appear in the right pane.

2. Select the schedule that you want to modify, and click **EDIT**.
3. Make the required changes, and click **SAVE**.

The schedule is successfully updated.

 **Note:** If you want to delete a schedule, select the schedule that you want to delete, and click **DELETE**.

Backup retention

The retention setting for a save set determines how long the NetWorker server maintains save set entries in the media database and client file indexes. Until the retention period expires, you can recover client backup data from backup storage either by browsing the data or by recovering the entire save set.

[Removing expired save sets](#) on page 541 describes how to remove save sets from backup storage after the retention period expires.

Methods for setting retention

You can specify retention for backup save sets and clone save sets in a variety of ways. If you specify retention by using multiple methods, then the retention setting that applies depends on the scenario.

Note:

If you set a retention policy on February 29 of a leap year, the last day in which the policy applied is 1 day earlier than you might expect. For example, if you set a retention policy to 1 year on March 3, 2015, the save set will expire on March 3, 2016 as expected, which is 366 days. If you set a retention policy to 1 year on February 29, 2016, you might expect that the policy will expire March 1, 2017. However, the policy will actually expire on February 28, 2017, which is 365 days. This behavior is only seen when a retention policy is set on February 29 for one or more years.

Retention for data protection policies

You can specify retention for backup save sets and clone save sets as part of the actions in a data protection policy. Retention settings are available for the traditional backup, snapshot backup, VMware backup, server backup, VBA checkpoint backup, and clone actions.

A single Client resource can belong to multiple groups. Therefore, you can assign different retention settings for the same client and save set data by configuring different workflows and actions. Consider the following example scenario:

- A client belongs to both Client Group A and Client Group B.
- Client Group A is assigned to Workflow 1, which performs a backup with a retention setting of 1 month.
- Client Group B is assigned to Workflow 2, which performs a backup with a retention setting of 1 year.

In this case, backups for the client that are performed with Workflow 1 are retained for 1 month, and backups for the client that are performed with Workflow 2 are retained for 1 year.

Retention for Client resources

You can assign a retention policy to a client resource that overrides the retention period that is specified in an Action resource, when you configure the **Client Override Behavior** attribute value to **Client Can Override** in the Action resource. [Assigning a retention policy to a Client resource](#) provides more information.

Retention for Pool resources

Previous versions of NetWorker allowed you to define a value in the **Retention** attribute of a Pool resource. When you update a NetWorker 8.2.x or earlier server, the

update process retains the value that is defined in the **Retention** attribute of a Pool resource as a read-only value.

Order of precedence for Retention resource attributes

It is recommended that you use the configuration settings in an Action resource to determine which pool received backup data. NetWorker provides you with the ability to configure a Pool attribute in the client resource, which can override the value defined in the Action resource. Additionally, the Pool resource contains 8.2.x legacy attributes that provide you with the ability to define backup data criteria for the pool. How and when NetWorker uses the attributes values defined in the Pool, Action, and Client resources to determine which backup pool will receive data depends on the value that you select in the **Client Override Behavior** attribute of the Action resource:

- **Client Can Override**—The value in **Retention** attribute of the Client resource takes precedence over the **Retention** value that is defined in the Action resource.
- **Client Can Not Override**—The value defined **Retention** attribute in the Action resource takes precedence over the value that is defined in **Retention** attribute of the Client resource and the **Retention** attribute of the Pool resource.
- **Legacy Backup Rules**—Enabled for migrations only. NetWorker uses the values that are defined in the **Retention** attribute of the Pool resource to determine which the retention policy to assign to backup data from a client. The value that is defined in the **Retention** attribute of the Pool resource take precedence over the **Retention** value that is defined in the **Action** resource and the **Retention** value that is defined in the Client resource.

 **Note:** You cannot modify the legacy attributes in the migrated Pool resources.

Retention for manual backups

If you specify retention with a manual backup from the command prompt with `save -w`, the retention setting applies to all the save sets that are in the manual backup. Specify the retention setting by using the time and date formats that are accepted by the `nsr_getdate` program. The `save` and `nsr_getdate` UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about data formats.

If you do not specify retention for a manual backup, then retention is applied based on the retention setting of either the Client resource or the media pool for the backup, whichever is longer. If there are multiple Client resources for the host, then the longest retention setting applies.

Assigning a retention policy to a Client resource

You can override the retention setting specified in the data protection policies that apply to a Client resource by specifying a retention setting for the Client resource itself.

About this task

NetWorker provides one of the following default retention policies that you can assign to the Client resource. Default retention policies include:

- Day
- Week
- Month
- Quarter
- Year

- Decade

You can also create a custom retention policy.

Procedure

1. (Optional) Create or customize the retention policy that you plan to assign to the Client resource.
 - a. In the **NetWorker Administration** window, click **Server**.
 - b. In the expanded left pane, select **Time Policies**.
 - c. Create a policy or modify a retention Policy resource:
 - To create a policy, from the **File** menu, select **New**.
 - To modify a policy, right-click the retention policy and select **Properties**.
 - d. For a new policy only, in the **Name** box, type a name for the retention policy.
 - e. Optionally, in the **Comment** box, type a description of the retention policy.
 - f. From the **Number of periods** and **Period** lists, specify the duration of the retention period.
 - g. Click **OK**.
2. In the **NetWorker Administration** window, select **View > Diagnostic Mode** to enable diagnostic mode view.

A check mark next to **Diagnostic Mode** in the **View** menu indicates that diagnostic mode view is enabled.

3. In the **NetWorker Administration** window, click **Protection**.
4. In the expanded left pane, select **Clients**.
5. In the right pane, right-click the client resource and select **Modify Client Properties**.

The **Client Properties** dialog box appears, starting with the **General** tab.

6. From the **Retention policy** list, select the retention policy to apply to all backups of the client resource, regardless of the retention setting for any data protection policies that apply to the client resource.
7. Click **OK**.

Editing retention for a save set

Use the `nsrmm` program with the `-e` option to edit the retention setting of a save set after the backup has occurred.

Specify the save set ID with the `-S` option, and specify the updated time in quotation marks with the `-e` option. The time and date format must use a format that is accepted by the `nsr_getdate` program.

Use the `mminfo` command with the `-p` option to view a report on the retention times for save sets.

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrmm`, `nsr_getdate`, and `mminfo` commands.

Example commands to edit retention for a save set

The following command updates the retention time for save set ID 3315861249 to midnight on January 1, 2016:


```
nsrmm -S 3315861249 -e "01/01/16 23:59:59"
```

The following command updates the retention time for save set ID 3315861249 to two years from the current date and time:

```
nsrmm -S 3315861249 -e "2 years"
```

General backup considerations

Before you configure Client resources to backup data on a host, review this section for information that applies to Windows, UNIX, and Mac OS-X hosts.

Renamed directories

When you rename a directory, a full backup is performed on all subdirectories and files of the renamed directory.

If you then rename the directory back to its original name, then files and subdirectories of the directory are not eligible for backup until the files or subdirectories are updated or the next full backup occurs.

You can change this default behavior by clearing the **Backup renamed directories** checkbox on the **General** tab of the **Client Properties** dialog box for a Client resource. You must select **View > Diagnostic Mode** in the Administration interface to access the **Backup renamed directories** attribute in the **Client Properties** dialog box.

When you clear the **Backup renamed directories** checkbox for a Client resource, unchanged files and folders under the renamed directory are skipped during a non-full backup. This behavior can cause unexpected results during a recovery operation. If you try to recover data under a renamed directory from a date between the time that the directory was renamed and the next full backup, it may appear that data is missing. For that recovery period, any files or folders that were unchanged do not appear under the renamed directory. Instead, they appear under the previous directory name.

You must leave the **Backup renamed directories** checkbox selected for clients that perform synthetic full backups.

Raw partitions

The NetWorker software must have exclusive access to a file system to perform a raw backup. Close as many applications as possible before doing a raw disk backup. If the raw partition contains data that are managed by an active database management system (DBMS), ensure that the partition is offline and the database manager is shut down. For greater flexibility when backing up partitions that contain DBMS data, use a NetWorker Module application.

Raw partitions on Windows

Back up raw disk partitions on Windows by specifying the raw disk partition in a save set with the `save` command. Identify the raw partition as a physical drive or logical drive. For example:

```
save -s NetWorker_server_name -o VSS:*=off \\.\e:
save -s NetWorker_server_name -o VSS:*=off \\.\PhysicalDrive0
```

Raw partitions on UNIX

Back up raw disk partitions on UNIX by using the `rawasm` directive.

Raw partitions on Linux

NetWorker can only save an unbound Linux raw device. When you back up a Linux raw disk partition, you must specify `/dev/sd` or `/dev/hd` in the **Save set** attribute on the **General** tab of the **Client Properties** dialog box for the Linux Client resource. The backup fails if you use the `/dev/raw` device.

Access control lists

The NetWorker software supports backup and restore of Access Control Lists (ACLs) and extended ACLs for Linux, HP-UX, AIX, DEC, Solaris, OS X, and Windows.

When a file with an associated ACL is backed up, the ACL is backed up along with the file data. When the file is recovered, any associated ACL is also recovered.

The **ACL passthrough** checkbox on the **Configuration** tab of the **NetWorker Server Properties** dialog box controls whether to recover files with associated ACLs. Select the checkbox to recover files with associated ACLs.

Client parallelism and parallel save streams

Client parallelism defines the number of data streams that a client can use simultaneously during backup.

Data streams include backup data streams, savefs processes, and probe jobs.

The default value is different for the NetWorker server than it is for all other client resources:

- For the NetWorker server client resource, the default value is 12. This higher default value enables the server to complete a larger number of index backups during a Server backup action.
- For all other clients, the default value is 4.

To define client parallelism, use the **Parallelism** attribute of the Client resource. You can find the parallelism attribute on the **Globals(1 of 2)** tab of the Client property dialog box, in the **NetWorker Administration** window.

The *NetWorker Network Data Management Protocol (NDMP) User Guide* provides more information about recommended parallelism settings for NDMP clients.

To avoid disk contention for clients other than the NetWorker server, specify a value that is the same as or fewer than the number of physical disks on the client that are included in the backup.

For a Windows client with the ALL keyword save set attribute, the backup includes the local disks, for example C: and D: drives as well as the System State and System DB. In this example, you can keep the default parallelism setting of 4. If you define multiple save sets on the same disk, for example, C:\users, C:\system, C:\docs and so on, a higher client parallelism results in multiple save streams attempting to access the disk at the same time.

The *NetWorker Performance Optimization Planning Guide* provides more information about recommended client parallelism values and performance benefits.

Enabling the parallel save streams (PSS) feature for a Client resource allows you to back up each save set for the client by using multiple parallel save streams to one or more destination backup devices. You can use PSS to perform the scheduled file level backup of file systems, and block based backups.

You can use PSS for clients with supported UNIX, Linux, and Windows operating systems. Supported save sets for PSS include the Save Set ALL, and individual save

points including `Disaster_Recovery`, `deduplicated`, and CSV volumes (Windows only). Checkpoint restart is not supported when you use PSS.

When you enable PSS, you can specify the maximum number of save streams that a client can send simultaneously for one or more save set backups concurrently running by using the **Parallelism** attribute in the **Client Properties** dialog. The default value for the **Parallelism** attribute is different for the NetWorker Server than it is for all other Client resources:

- For the NetWorker Server Client resource, the default value is 12. This higher default value enables the server to complete a larger number of index backups during a file system backup of the server or other index backups.
- For all other clients, the default value is 4.

Enabling PSS results in significant performance improvements due to save set aggregation, where the NetWorker Server starts a single save process per client with all client save sets that are passed to the single process for various processing optimizations, such as minimal Windows VSS snapshots and support for the following:

- Four parallel streams are started per save set, subject to any client parallelism limitations that might prevent all save sets from starting simultaneously.
- The ability to modify the number of parallel streams per save set by defining the new `PSS:streams_per_ss` environment variable save operations attribute in the properties of a Client resource. For example, setting `PSS:streams_per_ss=2,*` splits all save sets into two parallel save streams, whereas `PSS:streams_per_ss=3,/data1, 5,/data2` splits `/data1` into three parallel save streams and `/data2` into five parallel save streams.
- Automatic stream reclaiming, which dynamically increases the number of active streams for an already running save set backup to maximize utilization of limited client parallelism conditions.

Note: It is recommended that you set the client parallelism value to be a multiple of the `PSS:streams_per_ss` parameter default value 4 or its largest defined value when configured. For example, a multiple of 4 is 8, 12, or 16.

If the client parallelism is less than the `PSS:streams_per_ss` default 4 or the lowest configured value, the backup fails displaying an error message.

The `PSS:streams_per_ss` values range from 1 to 8. If you specify an invalid value, the backup proceeds with the default value 4, and a warning message displays stating that the entire `PSS:streams_per_ss` parameter is ignored.

The *NetWorker Performance Optimization Planning Guide* provides complete details on PSS requirements and performance benefits.

Configuring parallel save streams

Enable parallel save streams and specify the maximum number of save streams for a client by using the **Client Properties** dialog box. Note that the value specified for parallelism as part of an action in a policy is ignored for PSS backups.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Clients**.
3. Right-click the Client resource and select **Modify Client Properties**.

The Client Properties dialog box appears, starting with the **General** tab.

4. In the **Save set** attribute, specify **all** or a list of paths, for example, on UNIX /X and /Y or on Windows X:\ and Y:\.
5. Select the **Globals (1 of 2)** tab.
6. From the **Parallelism** list, specify the maximum number of save streams.
7. Select the **Parallel save streams per save set** checkbox.
8. Click **OK**.

Configuring parallel save streams for virtual clients

If you are backing up virtual clients, you can base the client parallelism setting on the underlying physical host. In this way, the total number of save streams for all virtual clients that reside on a physical host are limited to the value specified for the physical host.

About this task

For example, consider an environment with ten virtual machines running on the same physical host. Each virtual machine is a NetWorker client, and each client has a client parallelism setting of 4. This setting can result in a total of 40 save streams occurring on the same physical host, which would significantly slow down that system. To avoid this situation, you can specify that the client parallelism values are to be based on the underlying physical host. In this example, that would result in no more than four save streams occurring for the backup of the ten virtual clients.

Procedure

1. In the **Administration** window, select **View > Diagnostic Mode** to enable diagnostic mode view.
A check mark next to **Diagnostic Mode** in the **View** menu indicates that diagnostic mode view is enabled.
2. Click **Protection**.
3. In the expanded left pane, select **Clients**.
4. Right-click the Client resource for the virtual client and select **Modify Client Properties**.

The **Client Properties** dialog box appears, starting with the **General** tab.

5. Select the **Virtual client** checkbox.
6. Type the name of the underlying physical host in the **Physical host** box.
7. Select the **Globals (1 of 2)** tab.
8. From the **Parallelism** list, specify the maximum number of save streams.
9. Select the **Physical client parallelism** checkbox.
10. Select the **Parallel save streams per save set** checkbox.
11. Click **OK**.
12. Repeat these steps for all virtual NetWorker clients that share the same physical host.

Ensure that the value in the **Physical host** attribute is the same for all virtual NetWorker Client resources that share the same physical host.

Configuring backup retry and retry delay

This can be set in backup action.

The number of retries can be set maximum up to 5. The default retry value is 1. The more the number of retries means that NetWorker will try to complete the backup in successive attempts

The retry delay range can be set up from 1 to 720 seconds. The default retry delay value is 30 seconds.

Note: Starting with NetWorker 18.1 and later, the number of retries can be set maximum up to 24 and retry delay can be set maximum up to 3600 seconds from CLI only.

Troubleshooting PSS

About this task

It is recommended that you troubleshoot PSS with the guidance of Customer Service. The *NetWorker Performance Optimization Planning Guide* provides complete details on PSS requirements and performance benefits.

Procedure

1. Enable detailed logging for the client:
 - a. Specify the following value for the **Backup command** attribute on the **Apps & Modules** tab of the **Client Properties** box:

```
save -v -D7 (or D9 for more detailed logging)
```

- b. Type the following command at the command prompt on the client computer:

```
touch /nsr/debug/mbsdfopen
```

2. In the **Protection** window of the Administration interface, enable the `-v verbose` option for scheduled backups by selecting **Policies > policy name > workflow name**.
3. Wait for the next backup to occur, or manually start a backup by using one of the following methods:
 - In the **Protection** window of the Administration interface, right-click the workflow and select **Start**.
 - Use the `nsrpolicy` command on NetWorker server:


```
nsrpolicy start -p "policy" -w "workflow"
```

where *policy* is the name of the policy and *workflow* is the name of the workflow to start.
4. After the workflow finishes, collect the log files in the following table for Customer Service.

Table 67 Log files for PSS troubleshooting

Log file type	Log files to collect
Client	All log files in <code>/nsr/tmp/save-mbs-*</code>

Table 67 Log files for PSS troubleshooting (continued)

Log file type	Log files to collect
NetWorker server	• /nsr/logs/daemon.raw • All log files in /nsr/logs/policy/policy_name/ workflow_name/ action_name_sequence#_logs/* For example, /nsr/logs/policy/Silver/Filesystem/Backup_032334_logs/* • /nsr/tmp/savegrp.log

Maximum path and save set length

The maximum supported length in the NetWorker software for a pathname is 12 KB, and the maximum length for a save set name is 1024 bytes. The number of characters that are allowed by each of these limits depends on the locale.

All operating systems have an internal limit for path and file names. The limit depends on the operating system and file system. Typically, the pathname component size is 256.

For UNIX, only the path component length is checked against the limit. As a result, it is possible to create a path and file name that is greater than the limit supported by the operating system, but an try to access this path fails.

Open files

Open files are a problem that all data backup applications must solve. Open files that are not backed up correctly represent a potential data loss. They might be skipped, improperly backed up, or locked.

NetWorker can open files that are owned by the operating system and files that are owned by a specific application.

When you use VSS technology with NetWorker to create snapshot backups of volumes and exact copies of files, the backup includes all open files and files that change during the backup process.

Files owned by the operating system

Most open files that are owned by the operating system can be backed up. However, some applications can apply operating system locks to open files. These locks prevent other applications, such as NetWorker software, from writing to or reading from the open file.

The NetWorker software normally skips locked files and returns the following message:

```
save: filename cannot open
```

Also, the operating system might return a permission denied error.

To back up locked open files, close any open files if possible. To automate this process, create a pre- and postprocessing backup command that shuts down specific applications, backs up the open files, and then restarts any applications after the backup finishes.

You can also use Open File Manager to back up open files.

Files owned by a specific application

The NetWorker software cannot normally back up an open file that belongs to a specific application, like a database. To back up these open files, use a NetWorker Module. For example, use the NetWorker Module for SAP to back up open files in an Oracle database.

Files that change during the backup

If a file changes during a backup, the NetWorker software displays the following message in the **Monitoring** window:

```
warning: filename changed during save
```

To ensure that the changed file is backed up, either rerun the scheduled backup or perform a manual backup of the file.

NetWorker Modules can back up these types of files correctly if they are files that are related to the database that the module is backing up.

Data deduplication

Data deduplication is a type of data compression that removes duplicate information to reduce the amount of backup data sent to storage devices and reduce the bandwidth that is required for the data transport. You can implement data deduplication of NetWorker backup data by storing backups on Data Domain Boost deduplication devices.

Deduplication with DD Boost devices

The NetWorker client software includes the DD Boost library API and the distributed segment processing (DSP) component to enable deduplication on the client. The API enables the NetWorker software to communicate with the Data Domain system. The DSP component reviews the data that is already stored on the Data Domain system, and adds only unique data to storage.

DD Boost can run as many as 60 concurrent sessions (save streams) for a DD Boost device for backup and recovery. This high throughput reduces the number of necessary devices and the performance and maintenance impact on the Data Domain system. The resulting performance gain provides an advantage over conventional advanced file type device (AFTD) or virtual tape library (VTL) interfaces that do not handle these high session rates.

To perform deduplication backups with a Data Domain system, perform the following tasks:

- Configure the Data Domain system for use with NetWorker.
- Add the device in the NetWorker Administration interface.
- Select Data Domain backup options for Client resources.

The *NetWorker Data Domain Boost Integration Guide* provides details on system requirements and configuration steps.

Deduplication with Avamar

The Avamar client software only provides support to NetWorker hosts that used an Avamar system as a data protection target with a previous release of NetWorker. You cannot configure new Avamar nodes in NetWorker 19.1

Directives

Directives are resources that contain special instructions that control how the NetWorker server processes files and directories during backup and recovery. Directives enable you to customize the NetWorker software, maximize the efficiency of backups, and apply special handling to individual files or directories.

Types of directives

There are three types of directives.

- **Global directives**—Stored as resources on the NetWorker server and can be selectively applied to individual clients by using the **Directive** attribute of the Client resource.
- **NetWorker User local directive**—On Windows clients only, users with local Windows Administrator or Backup Operator privileges can create a local directive in the NetWorker User program. A file that is named `networkkr.cfg` on the client file system contains the directive configuration information. NetWorker uses the directive that is specified in the `networkkr.cfg` during a scheduled backup, a backup that is started with the NetWorker User application, and `save` operations that do not include the `-i` option.
- **Local directive files**—User-created files named `nsr.dir` (Windows) or `.nsr` (UNIX) anywhere on a client file system where they have permission to create files. These directives apply only to the immediate data within the path where the directive file is located.

If there is a conflict between directives, global directives are enforced over local directives. Also, NetWorker User program local directives are enforced over local directive files (`nsr.dir` files) on Windows hosts.

NOTICE If you use the Windows BMR feature, implement user-defined directives with caution. Using such directives in directories with system state files can lead to an incomplete BMR backup image and potentially render the BMR backup image unusable. If you create user-defined directives, test the BMR backup image to ensure that you can recover the Windows system state correctly.

Format of directive statements

Directive statements specify the files or directories and then the action to perform on the files and directories.

A directive statement specifies the following items:

- The directory for a directive statement.
- The action to perform, specified using either of the following:
 - An ASM specification with a pattern list of child file or directory names, which could include wildcards.
 - A save environment keyword specifies the action to perform.

A directive statement is written in the following format:


```
<< "directory_specification" >>
[+] ASM: pattern
save_environment_keyword
# comment
```

where:

- The directive statement does not contain blank lines.
- *directory_specification* is the absolute path to the highest-level directory for which the ASM in the directive applies.
- [+] is optional. The presence of the plus (+) sign indicates that the directive recursively applies to the directory defined by the absolute path and all subdirectories.
- *ASM* is the ASM that specifies the action to take on one or more files in the current directory.
- *pattern* is a list of file or directory names, in the current directory on which to apply the ASM.
The pattern can include multiple names that are separated by spaces, and wildcards. Wildcards can replace a single character or string of characters. Directive statements support the use of standard shell command interpreter file matching patterns. You cannot specify lower-level subdirectories in the pattern. That is, the pattern must not contain the Unix "/" or Windows "\" directory separator.



Note: File names are case-insensitive for directives that are applied to Windows clients.

- *save_environment_keyword* is a NetWorker keyword that controls how the current ASM and subsequent ASMs that apply to the current directory and subdirectories are applied in the directive statement.
- *comment* is a user-defined description of the directive statement. A hash (#) character must precede the comment.



Note: If an ASM or pattern name includes a space, enclose the name or argument in double quotation marks.

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about directives in the `nsr` and `nsr_directive` commands.

Defining directory specifications

A directory specification is the absolute path to the highest-level directory for which the ASM in the directive applies.

Consider the following information before you define a directory specification.

Defining directory specifications using wildcards

You can use wildcards in the directory specification to avoid maintaining multiple directives for specific directory paths in both the NetWorker server and the NetWorker client directive files.

To use wildcards in a directory specification, you must type an asterisk before the directory path. For example: This asterisk does not do any matching and only enables the use of any following wildcards in the directory specification.

- UNIX: << **/directory_specification/* >>
- Windows: << ***\directory_specification* >>

Consider the following examples for using wildcards:

- The following UNIX directive skips backing up the `tmp` folder data for each user whose name starts with letters A, B, C, or D:
`<< */*/users/[A-D]*/tmp/ >> +skip: *`
- The following Windows directive compresses the `tmp` folder data for each user whose name starts with letters A, B, C, or D:
`<< "C:\*\users\[A-D]*\tmp\" >> +compressasm: *`

Note:

- Networker does not support symbolic links in skip directives. Symbolic links are considered as regular files.
 For example, the symbolic link for the file `example.dbf` is `.example.dbf`. To skip the file `.example.dbf`, you must use `+skip: *.dbf`.
- Do not use wildcards in Windows VSS source volume directory mount names.

Defining multiple directory specifications

- Directives that follow a directory specification apply the ASM action to that directory until the next directory specification.
- Directory specifications that do not contain wildcards take precedence over directory specifications that contain wildcards. For example, when a directive encounters a directory that matches a wildcard directory specification and a non-wildcard directory specification, the directive will only apply the action defined in the ASM of the non-wildcard directory specification of the directory.

Defining mount points in the directory specification

Directory mount names of Windows VSS source volumes cannot contain wildcards. For example, you cannot create a directory specification in the format `<< "?:\data\*\\" >>`. You must specify each drive letter, for example `<< "C:\data\*\\" >>`, and create a separate directory specification for it.

Defining a directory specification for a Windows client system

- File and directory names are case-insensitive.
- If there is a colon (:) in the pathname, enclose the entire path in quotation marks.

Order of execution in the directive

The command that is listed first in the directive takes precedence over the rest of the commands.

When the first command in the directive is `compressasm`, then all the files are compressed. The skip command in the directive is ignored. Example

```
+compressasm -gzip -l: * .* *.*
+skip: *.LRG
+skip: *.CAT
+skip: *.TLB
```

When the `compressasm` is positioned at the end of the directive, the commands in the directive work in the order that is mentioned. The files specified with `+skip` are skipped and the rest are compressed. Example

```
+skip: zzz_DUMMY*
+skip: *.tmp
+skip: .db2diag.log.swp
skip: tmp_mnt
```

```
+compressasm -gzip -1: *.*.*.*
```

Using wildcards in directive statements

NetWorker supports wildcards in directive statements.

In a directive statement, you can use wildcard characters in both the directory and path list specifications.

The following table describes the supported wildcard characters and their descriptions.

Table 68 Supported wildcards in directives

Wildcard	Name	Description
*	Asterisk	Matches any sequence of characters.
?	Question mark	Matches any single character.
[and]	Square brackets	Forms an expression. Represents any of the characters enclosed within the square brackets. For example, the following directory specification includes all user folders beginning with A or D: <pre><< */*/users/[AD]*/tmp/ >></pre> You can use exclamation marks and hyphens within expressions.
!	Exclamation	Must be used directly after the opening square bracket in an expression, for example, [!...]. Matches any single character except for the characters typed after the !. For example, to match any single character except a, b, and c, type: <pre>[!abc]</pre>
-	Hyphen	When a hyphen is used between two characters within square brackets, it indicates a range inclusive of those two characters. For example, the following directory specification includes all user folders beginning with A, B, C, or D: <pre><< */*/users/[A-D]*/tmp/ >></pre>

Note: Wildcards directives are applied only by the save process in the scheduled backup workflow. When you skip unwanted mount points using directory specification wildcard directives with the save set keyword **All**, save processes still run for the unwanted mount points. However, the save processes only log messages that the contents of the mount points were skipped. For example:

```
<< */TestFileSystem[5-9]?/ >> +skip: *.
```

Directive specification with ASMs and save environment keywords

A directive specification specifies the action to perform in a directive statement.

Save environment keywords

Save environment keywords control how the current ASM and subsequent ASMs applied to the directory and subdirectories are applied in the directive statement.

NetWorker supports the following *save_environment_keyword* values:

- forget**—Instructs the NetWorker Server to no longer apply inherited directives (those directives that begin with a +). The **forget** keyword works only if the corresponding directories are also explicitly specified in the NetWorker Client resource **Save Set** attribute.
- ignore**—Instructs the NetWorker Server to ignore all directives that are applied to the subdirectories below the current directory.

- **allow**—Used in subdirectories that currently have the `ignore` keyword applied to them, and overrides the `ignore`.

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about directives in the `nsr` and `nsr_directive` commands.

Review the following examples of directive specifications that include ASMs and save environment keywords.

Using the skip directive for a Windows host

The following example directive statement skips the `C:\Program Files` folder on a Windows host during a backup:

```
<< "C:\Program Files" >>
skip
```

Using the skip directive for a UNIX host

The following example directive statement skips all files in the `/tmp` directory on a UNIX host, including hidden files:

```
<< /tmp >>
+skip: * .*
```

 **Note:** A space appears after the first asterisk (*) in the pattern.

Using the skip ASM and forget save environment keyword

The following example directive statement skips all `*.o` files in the `G:\SRC` directory except those `*.o` files in the `G:\SRC\SYS` directory:

```
<< "G:\SRC" >>
+skip: *.o
<< "G:\SRC\SYS" >>
forget
```

This example uses the `skip` ASM to instruct the NetWorker server to skip all files that are named `*.o` in the `SRC` directory and all subdirectories. It then uses the `forget` keyword to instruct the server to not apply the `skip` ASM to the `SYS` subdirectory.

Both the `G:\SRC` and the `G:\SRC\SYS` directories must be explicitly specified on separate lines in the client resource **Save Set** attribute.

Using the ignore save environment keyword

The following example allows directives in the `HOMEDOC` directory to be applied to the preceding example for the `ignore` keyword:

```
<< HOME >>
ignore
<< HOMEDOC >>
allow
```

Using the allow save environment keyword

The following example directive statement overrides any local directives set in user home directories:

```
<< HOME >>
ignore
```

Order of Execution in the Directive

The command that is listed first in the directive takes precedence over the rest of the commands.

When the first command in the directive is `compressasm`, then all the files are compressed. The `skip` command in the directive is ignored. For example,

```
+compressasm -gzip -l: * .* *.*
+skip: *.LRG
+skip: *.CAT
+skip: *.TLB
```

When the `compressasm` is positioned at the end of the directive, the commands in the directive work as expected. The files that are specified with `+skip` are skipped and the rest are compressed. For example,

```
+skip: *.LRG
+skip: *.CAT
+skip: *.TLB
+compressasm -gzip -l: * .* *.*
```

Global directives

Global directives are stored as resources on the NetWorker server and can be selectively applied to individual clients by using the **Directive** attribute of the Client resource.

Global directives are listed when you select **Directives** in the expanded left pane of the **Server** window in the Administration interface. You can add, edit, copy, and delete global directives.

Preconfigured global Directive resources

The NetWorker software includes a number of preconfigured global Directive resources. All preconfigured Directive resources can be modified, but they cannot be deleted.

The following table lists the preconfigured directives and their descriptions.

Table 69 Preconfigured directives

Directive resource	Description
AES	Encrypts backup data with the aes ASM, which provides 256-bit data encryption.
Mac OS with compression	Contains the same set of directives as the Mac OS standard directive, along with applying the <code>compressasm</code> ASM to specific directories.
Mac OS standard	Contains a set of directives that are used to back up standard Mac OS clients. Applies these ASMs: The <code>skip</code> ASM is applied to these files and directories:

Table 69 Preconfigured directives (continued)

Directive resource	Description
	/Desktop DB /Desktop DF /cores /VM_Storage /TheVolumeSettingsFolder /private/var/db/netinfo /private/var/db/openldap /private/tmp /.Spotlight-V100 /.hotfiles.btree • The allow save environment keyword is applied to the <code>/nsr</code> directory to ensure that local directives in <code>/nsr</code> and subsequent subdirectories are applied. • The logasm ASM is applied to the <code>/nsr/logs</code> and <code>/var</code> directories. • The swapasm ASM is applied to the <code>/private/var/vm</code>
NT standard	Is used to back up Windows clients. By default, this resource has no directives.
NT with compression	Used to back up and compress Windows clients. It applies the compressasm ASM to all files.
UNIX standard	Contains a set of directives that are used to back up standard UNIX clients. Applies these ASMs: • The skip ASM is applied to the <code>tmp_mnt</code> directory. • The skip ASM is applied to core files on the file system. • The allow save environment keyword is applied to the <code>/nsr</code> directory to ensure that local directives in <code>/nsr</code> and subsequent subdirectories are applied. • The skip ASM is applied to the <code>/tmp</code> directory. • The swapasm ASM is applied to the <code>/export/swap</code> directory. If swap files are located in a different directory, modify

Table 69 Preconfigured directives (continued)

Directive resource	Description
	this directive to use the appropriate directory. - The logasm ASM is applied to the <code>/nsr/logs</code>, <code>/var</code>, <code>/usr/adm</code>, and <code>/usr/spool</code> directories. You can apply this ASM to other directories as well. - The mailasm ASM is applied to the <code>/usr/spool/mail</code> and <code>/usr/mail</code> directories. If email files are located in different directories, modify these directives to use the appropriate locations.
UNIX with compression	Contains the same set of directives as the UNIX standard directive, along with applying the compressasm ASM to all files. This directive is only applied to save sets that contain directories. If the save set is defined by using a file name, this directive is not applied.
VCB directives	VCB directives are valid for backing up virtual machines using the VCB methodology. This directive is supported in the following scenarios: - When file level incremental backups are performed instead of FULL image level backups. - When FULL file level or incremental file level backups are performed when the save set is ALLVMFS. The vcb directive skips the following files and folders: <code>pagefile.sys</code> <code>hiberfil.sys</code> (Hibernation file) <code>WINDOWS\system</code> folder <code>WINDOWS\System32</code> folder

Creating a global Directive resource

Procedure

1. In the **Administration** window, click **Server**.
2. In the expanded left pane, select **Directives**.

3. From the **File** menu, select **New**.

The **Create Directive** dialog box appears.

4. In the **Name** box on the **General** tab, type a name for the new directive.
5. In the **Comment** box, type a description of the directive.
6. In the **Directive** attribute, type one or more directive statements.

A directive statement specifies the files and directories for a directive statement, and then an ASM specification or a save environment keywords specifies the action to perform. You can also include comments in a directive statement by preceding text with a hash (#) character.

For example, the following directive statement skips the `C:\TEMP` folder on a Windows system during a backup:

```
<<"C:\TEMP">>
skip
```

 **NOTICE** Do not leave blank lines in the directive statement.

[Format of directive statements](#) provides more information about how to create a directive statement.

7. To specify a restricted datazone (RDZ) for the directive, click the **Restricted Data Zones** tab and then select the RDZ from the list.
8. Click **OK**.

After you finish

Apply the global directive to a Client resource by selecting the directive from the **Directive** list on the **General** tab of the **Client Properties** dialog box for the Client resource.

Editing a global Directive resource

You can edit the directive statement, description, or RDZ of a global Directive resource. To rename a global directive, delete the global directive and create a global directive with the new name.

Procedure

1. In the **Administration** window, click **Server**.
2. In the expanded left pane, select **Directives**.
3. In the right pane, perform one of the following tasks:
 - To modify multiple attributes in a single configuration resource by using the **Directive Properties** window, right-click the staging configuration and select **Properties**.
 - To modify a specific attribute that appears in the resource window, place the mouse in the cell that contains the attribute that you want to change, then right-click. The menu displays an option to edit the attribute. For example, to modify the **Comment** attribute, right-click the resource in the **Comment** cell and select **Edit Comment**.

 **Note:** To modify a specific attribute for multiple resources, press and hold the **Ctrl** key, select each resource, and then right-click in the cell that contains the attribute that you want to change. The menu displays an option to edit the attribute.

4. Edit the settings for the global directive, then click **OK**.

Copying a global Directive resource

Procedure

1. In the **Administration** window, click **Server**.
2. In the expanded left pane, select **Directives**.
3. In the right pane, right-click the directive and select **Copy**.

The **Create Directive** dialog box appears with the settings from the original directive.

4. In the **Name** box, specify a name for the directive.
5. Edit the other settings for the directive as necessary.
6. Click **OK**.

After you finish

Apply the global directive to a Client resource by selecting the directive from the **Directive** list on the **General** tab of the **Client Properties** dialog box for the Client resource.

Deleting a global Directive resource

Before you begin

- Ensure that the global Directive resource is not a default global Directive resource. You cannot delete global Directive resources that are available by default when you install the NetWorker server software.
- Ensure that the Directive resource is not selected for any Client resources.

Procedure

1. In the **Administration** window, click **Server**.
2. In the expanded left pane, select **Directives**.
3. In the right pane, right-click the directive and select **Delete**.

A confirmation message appears.

4. Click **Yes**.

NetWorker User local directives

On Windows clients, users with local Windows Administrator or Backup Operator privileges can create local directives by using the NetWorker User program. These directives are stored on the client in a file named `networkr.cfg`.

About this task

When you perform a manual backup from the NetWorker User program, only local directives that were created with the NetWorker User program are enforced. Global directives and local directive files (`nsr.dir` files) are not enforced. However, all local directives are enforced when the NetWorker `save` command without the `-i` option is run at the command prompt.

NetWorker User program local directives are also enforced during scheduled backups and archive operations.

Procedure

1. Log in to the client computer as a member of either the local Windows Administrators or Backup Operators security group.
2. Start the NetWorker User Program.
3. From the **Options** menu, select **Local Backup Directives**.
4. Set the local directive for each data item. You can clear data items to exclude them from scheduled backups, and select items for password protection, encryption, and compression. This applies for both manual and scheduled saves.

 **Note:** If password protection or encryption is selected, the password must be specified first.

5. From the **File** menu, select **Save Backup Directives** to save changes.

Depending on user privileges and the operating system version, the `networkr.cfg` file is created in one of the following locations:

- If you are logged in with local Windows Administrator or Backup Operator privileges, `networkr.cfg` is created in the root of the system volume (usually `C:\`).
- If you are not logged in with local Windows Administrator or Backup Operator privileges, `networkr.cfg` is created in `%SystemDrive%\Documents and Settings\User_name\Application Data\EMC NetWorker`.

 **Note:** The `Application Data` directories are hidden by default. To view these directories by using Windows Explorer, select **Tools > Folder Options**. On the **View** tab of the **View Options** dialog box, select the **Show hidden files and folders** option.

Creating local directives

Local directives are text files that are on the file system of the client. The directives apply only to the immediate data within the path where the directive file is saved.

Procedure

1. Use a text editor to create the directive file in the directory that contains the files to which you plan to apply the directive.
2. Create the directive statement.

A directive statement specifies the files and directories for a directive statement, and then an ASM specification or a save environment keywords specifies the action to perform. You can also include comments in a directive statement by preceding text with a hash (`#`) character.

For example, the following directive statement skips the `C:\TEMP` folder on a Windows system during a backup:

```
<<"C:\TEMP">>
skip
```

 **NOTICE** Do not leave blank lines in the directive statement.

[Format of directive statements](#) provides more information about how to create a directive statement.

3. Save the local directive file.

- On Windows, the file must be named `nsr.dir`. The user account that creates the file must have the permissions to create files either within the root of the volume or in a folder within the volume.
- On UNIX, the file must be named `.nsr`.

CHAPTER 6

Backing Up Data

This chapter contains the following topics:

• Configuring a Client resource for backups on Windows hosts	410
• Configuring a Client resource for backups on UNIX hosts	452
• Configuring a Client resource for backups on Mac OS X hosts	459
• Sending client data to AFTD or Data Domain devices only	464
• Non-ASCII files and directories	465
• Configuring checkpoint restart backups	465
• Probe-based backups	470
• Encryption	471
• Compression	474
• Configuring Client Direct backups	475
• Backup command customization	477
• Client resources	483
• Manual backups	488
• Verifying backup data	491

Configuring a Client resource for backups on Windows hosts

This section describes how to configure a Client resource to backup data on Windows hosts.

Windows backup considerations

Use the NetWorker software to back up Window file systems. The NetWorker Module for Microsoft (NMM) provides VSS-based backup and recovery of the Windows operating system and Microsoft server applications such as Microsoft Exchange Server, Microsoft SQL Server, and Microsoft SharePoint Services.

The *NetWorker Module for Microsoft Administration Guide* provides more information about the NMM product.

Configuring how NetWorker determines when to back up a file

You can configure NetWorker to back up a file that is based on the setting of the **Archive** file attribute in the properties of a Windows file or based on the modification time.

The NetWorker software saves a file when the **Archive** attribute is enabled. After NetWorker saves the file, the NetWorker software disables the **Archive** attribute. If you restore the file from a backup, then the NetWorker software enables the **Archive** attribute to ensure that the next backup includes the file.

To configure NetWorker to use the modification time of a file instead of the **Archive** attribute, perform the following steps:

1. Browse to **Control Panel > System > Advanced System Settings**.
2. On the **Advanced** tab, click **Environment Variables...**
3. In the **System Variables** section, click **New**.
4. In the **Variable name** field, type `NSR_AVOID_ARCHIVE`.
5. In the **Variable value** field, type `yes`.
6. Click **OK** to close the **Environment Variables** window, and then click **OK** to close the **System Properties** window.
7. Log off or restart the client computer, or restart the `NetWorker Remote Exec Service` to make Windows aware of the environment variable change.

Backup Operators group

The Windows Backup Operators local group provides its members the privileges necessary to back up and recover data from a Windows computer.

Users who request backups must be in the Backup Operators or Administrators group of the domain into which they are logged. The Backup Operators group is assigned on a computer-by-computer basis, rather than globally by the domain. If you are having trouble performing tasks on one NetWorker server but not another, check the Backup Operators group on the problematic computer to ensure that you are correctly assigned.

Enabling NetWorker logging operations performed by backup operator

By default, members of the Windows Backup Operators group do not have write permission to the <NetWorker_install_path>\logs directory.

About this task

NetWorker log operations are performed by members of the Windows Backup Operators group.

Enable NetWorker logging for Backup Operators by modifying the security settings on the <NetWorker_install_path>\logs directory. For example:

Procedure

1. In **Windows Explorer**, navigate to the <NetWorker_install_path>\logs directory.
2. Right-click the <NetWorker_install_path>\logs directory icon and select **Properties**.
3. On the **Security** tab of the **Properties** dialog box, add the **Backup Operators** group to the list of groups and users.
4. Select the **Backup Operators** group and click **Allow Write**.
5. Click **OK**.

Windows backup considerations

Use the NetWorker software to backup Windows file systems. NetWorker Module for Microsoft (NMM) provides VSS-based backup and recovery of the Windows operating system, and Microsoft server applications, for example, Microsoft Exchange Server, Microsoft SQL Server, and Microsoft SharePoint Services. The *NetWorker Module for Microsoft Administration Guide* provides more information about the NMM product.

Table 70 Backup considerations for Windows features

Windows Feature	Backup considerations
Event logs—Used for troubleshooting hardware problems as well as monitoring security conditions, and system and application software problems.	To back up event log files, configure a file system backup that includes the C:\Windows\system32\winevt\logs folder. The size of a recovered event log might be smaller than the backup size. This is a characteristic of Windows event logs and does not cause any data loss or change of data. You can use Microsoft Event Viewer to view the recovered, smaller log file. NetWorker backs up all event log files when more than one active event log is marked for backup (for example,

Table 70 Backup considerations for Windows features (continued)

Windows Feature	Backup considerations
	SecEvent.Evt and SysEvent.Evt). You can recover event logs to a location that differs from the location at the time of the backup. You cannot recover event logs files that were on an NTFS partition at the time of the backup to an FAT16 or FAT32 partition.
Encrypted File System (EFS)—Allows NTFS files to be stored in encrypted format. A user without the private key to the file cannot access the file.	NetWorker will not encrypt or compress a file already encrypted by Windows. Do not use AES encryption when you backup EFS encrypted files. Files can become unusable if the encryption keys change on the domain controller. For example, when you move the domain controller from one computer to another or the domain controller failures. NetWorker does not backup the encryption keys, or keep a copy of the keys to ensure a successful recovery of EFS encrypted files to an EFS that you reinstall after a disaster. When recovering encrypted files to an encrypted folder that has been removed, consider the following: • If you recover the encrypted files and the encrypted folder, the recovered folder and files are all encrypted. • If you recover only individual encrypted files (but do not recover the encrypted folder that contains them) the individual recovered files are encrypted but the re-created folder is not encrypted. Windows documentation provides instructions on encrypting the re-created folder. • Windows EFS encrypted data is backed up and recovered in its encrypted state.
Internet Information Server (IIS)— A web server that enables the publication of information on the Internet or a corporate intranet by using HTTP.	NetWorker uses the active metabase to back up IIS and can restore the backup versions to the metabase location. NetWorker supports the recover of the metabase to the default

Table 70 Backup considerations for Windows features (continued)

Windows Feature	Backup considerations
	location %SystemRoot%\system32\inetsrv\MetaBase.bin or in a location that you specify in the registry. The Microsoft documentation provides information about how to create a registry key that specifies an alternate metabase location.
Sparse files— Enables a program to create huge files without actually committing disk space for every byte.	NetWorker provides complete backup and recovery support for sparse files.
Windows Print Queues	NetWorker backs up and recovers print queues as a part of the file system backup. During a recover operation, you may have to restart the host depending on the status of the print queue at the time of the backup.
Disk quota database	The WINDOWS ROLES AND FEATURES save set contains the disk quota database. During a backup operation, NetWorker creates temporary files to store the disk quota database settings in the root directory of each drive on the client. Note: To backup the disk quota database, the local system account must have full control permissions on the local drive, otherwise a backup fails with an error message similar to the following: Failed to write to quota file, 0x80070005
POSIX compliance	NetWorker performs case sensitive backup and recovery operations. During a recovery operation on a Windows host, NetWorker may create multiple files with the same name but different cases. For example, you back up a file on a Windows host that is named <code>temp.txt</code>. The file is later deleted and created with a new file named <code>Temp.txt</code>. When you select the <code>temp.txt</code> file for recovery, NetWorker will not overwrite the file that is named <code>Temp.txt</code>. You will have two identical files in the directory, one named <code>temp.txt</code> and the other named <code>Temp.txt</code>. To configure NetWorker to ignore the case of a file, you can set the system environment variable

Table 70 Backup considerations for Windows features (continued)

Windows Feature	Backup considerations
	<i>NSR_DISABLE_POSIX_CREATE=YES</i> , which disables POSIX compliance.
Windows Dynamic Host Configuration Protocol (DHCP) and Windows Internet Naming Service (WINS) databases	The WINDOWS ROLES AND FEATURES component of the DISASTER_RECOVERY:\ save set contains the DHCP and WINS databases. Use Windows BMR recovery to perform an offline restore of these databases.
Native Virtual Hard Disk (VHD) volumes— Used as a mounted volume on designated hardware without any other parent operating system, virtual machine, or hypervisor. You can use a VHD volume as a boot volume or as a data volume.	The ALL save set does not include native VHD volumes. Configure a separate client resource to backup native VHD volumes. Do not use VHD volumes as critical volumes if the volume that contains the native VHD is also a critical volume. This situation creates a conflict during a Windows BMR backup.
Windows Content Index Server (CIS) or Windows Search Index— Index the full textual contents and property values of files and documents that are stored on the local computer. The information in the index can be queried from the Windows search function, the Indexing Server query form, or a web browser.	The WINDOWS ROLES AND FEATURES component of the DISASTER_RECOVERY:\ save set contains the CIS or Windows Search Index. The CIS or Windows Search is automatically regenerated on system restart. NetWorker performs the following actions when performing a CIS or Windows Search backup: • Pauses any CIS or Windows Search catalogs. You can still query a paused catalog, so the indexing functionality is not lost during the CIS or Windows Search backup. • Backs up all catalog files. • Turns on the catalogs when the backup completes. • CIS or Windows Search deletes the catalog folder during a backup and restores it as part of a recovery operation.

DHCP and WINS databases

The `WINDOWS ROLES AND FEATURES` component of the `DISASTER_RECOVERY:\` save set contains the Windows Dynamic Host Configuration Protocol (DHCP) and Windows Internet Naming Service (WINS) databases. Use Windows BMR recovery to perform an offline restore of these databases.

About this task

The `ALL` save set also includes the DHCP and WINS databases because the `ALL` save set automatically includes the `DISASTER_RECOVERY:\` save set.

If you do not specify the `ALL` save set or the `DISASTER_RECOVERY:\` save set in the **Save set** attribute for the client, then include the databases as part of a file system backup:

- To back up a DHCP database, include the `%SystemRoot%\System32\dhcp` directory in the **Save set** attribute of the Client resource for the DHCP server.
- To back up a WINS database, use the Microsoft WINS administrative tools to configure an automated backup of the WINS database to a local drive on the WINS server. Then specify the path to the database backup on the local drive in the **Save set** attribute of the Client resource for the WINS server.

Hard links

You can back up and recover files with hard links on a Windows client. However, the hard links of files that are created by using a Portable Operating System Interface (POSIX) application are not preserved during recovery.

Support for hard links is disabled by default to improve performance.

Backup and recovery of hard links is disabled by default to improve performance. To enable backup and recovery of hard links on a client, select the **Hard links** checkbox on the **Globals (2 of 2)** tab of the **Client Properties** dialog box for the Client resource.

Enable diagnostic mode view by selecting **View > Diagnostic Mode** in the **Administration** window to access the **Hard links** checkbox.

Microsoft DFS

You can back up and restore Microsoft Distributed File System (DFS) data.

Microsoft DFS is a Windows file system feature that enables you to create a namespace of shared directories that are physically distributed across a network. With DFS, you can organize a set of distributed directories logically, according to any scheme you choose, to provide centralized access to files that reside in a variety of locations.

DFS junctions

A DFS junction is a DFS root or link:

- A DFS root is a namespace for files and DFS links.
- A DFS link is a connection to a shared file or folder.

DFS junctions are file system objects, not files or directories. Therefore, the NetWorker software does not treat DFS junctions the same as files or directories for backup and recovery. However, DFS junctions appear as files and directories in the NetWorker User program.

DFS backups with the ALL-DFSR save set

The `ALL-DFSR` save set includes all DFS related save sets for a backup. Unlike other all-inclusive save set types, `ALL-DFSR` is not related to any particular file system. `ALL-DFSR` backs up all components that are defined by DFS\FRS writers. Backups fail if you specify `ALL-DFSR` for a system where DFS or FRS is not installed.

The syntax for this save set is `ALL-DFSR`. It is not case sensitive.

The `ALL-DFSR` save set does not support BBB. BBB only creates backups at the volume level, and DFSR replication folders can be a subfolder, which creates a conflict.

Synthetic full backup is not supported with `ALL-DFSR`.

The `ALL-DFSR` save set registers the corresponding writer and writer component nodes under `WINDOWS ROLES AND FEATURES`. All Replication folders are restored through these nodes.

Configuring a scheduled DFS backup

About this task

To avoid inconsistencies among the various save sets, configure a scheduled backup that includes the DFS topology information, junctions, and destination directories. Alternatively, you can use the `ALL-DFSR` save set.

NOTICE When a DFS client resource is run for the first time, the save set sizes should be verified to ensure that they are correct.

To configure a scheduled backup for a DFS:

Procedure

1. In the **Administration** screen, include the following clients in the NetWorker group that will back up the DFS:
 - The DFS host server
 - Any computer where remote DFS destination directories reside
 - A domain controller (domain-based DFS only)

For example, you could create a NetWorker group named DFS, then make each of the preceding clients a member of the DFS group.

2. Enter the following save sets in the **Save Set** attribute of the DFS host server's client resource:
 - The DFS root. For example, `C:\MyDfsRoot`.
 - DFS destination directories that reside on the DFS host. For example, `D:\MyLocalDir`

Note: DFS destination directories are also be backed up if you enter the entire volume (for example, `D:\`) in the Save Set attribute.

3. For clients where remote DFS destination directories reside, enter the destination directory paths in the **Save Set** attribute. For example:

```
E:\MyRemoteDir
E:\MyOtherRemoteDir
E:\
```

Windows Optimized Deduplication

NetWorker supports backup of optimized data deduplication volumes and files and can restore optimized deduplication backups to a set of eligible restore targets. However, due to recovery performance issues observed with optimized backup for Windows deduplication volumes, it is recommended that you use non-optimized backup, which is enabled by default.

When the backup is set to non-optimized, NetWorker will not deduplicate the backup. Instead, the deduplicated files get rehydrated in memory before they are backed up. This type of backup requires you to enable VSS. If you disabled VSS (for example, by specifying `vss:*=off` in the **Save Operations** attribute), the backup might back up the chunk stores unnecessarily. To back up the deduplicated volume, it is recommended to use block based backup (BBB) instead.

If you still require optimized backup, you can add `vss:NSR_DEDUP_NON_OPTIMIZED=no` to the **Save Operations** attribute to restore settings to the traditional (non-BBB)

optimized backup. However, it is not recommended that you use this setting because recovery performance issues might result in an unusable backup. Note also that even when optimized deduplication backup is enabled, NetWorker will not perform an optimized backup when the backup path is a subdirectory of the volume, or the non-optimized deduplication save option is specified in the **Save operations** field of the Client resource.

NetWorker supports the data deduplication feature on Windows Server 2012, Windows Server 2012 R2, Windows Storage Server 2012, and Windows Storage Server 2012 R2. NetWorker does not support the feature on Windows 8 client computers or computers that run the older versions of the Windows operating system. On computers that run the Windows Server operating system, NetWorker supports the feature on volumes that use the NTFS file system, which can be part of a fail over cluster, including CSV volumes.

To back up and restore Windows Server deduplication volumes or files, you must use a NetWorker 8.1 or later client. You can only restore deduplicated backups to computers that run on supported versions of Windows Server that have the data deduplication role enabled. The data deduplication role is a child role of File Services, which is a File and Storage Services role.

Detecting deduplication in a backup

When a deduplication volume is backed up, you can verify the form of the data that was backed up. This information is identified in the `mminfo` extended save set attributes output. To show all extended save set attributes, use the `mminfo` output flag `-r attrs`. Deduplication backups are indicated with

`*MSFT_OPTIMIZED_DEDUP_ENABLED:yes`.

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `mminfo` command.

Data Deduplication Backup and Restore

NetWorker supports two types of backup and four types of restores for data stored on a deduplication volume.

Optimized full-volume backup

Optimized full-volume backups are the default backup type for Windows data deduplication volumes. The backup type occurs when the non-optimized data deduplication save option is not specified and the backup path is a mount point, drive letter or full volume backup. NetWorker full, incremental, and synthetic full backups are supported with Windows data deduplicated volumes.

The optimized data deduplication files that are part of the backup include:

- Windows data deduplication reparse points
- Chunk store containers and data deduplication meta data files

NetWorker backup does not differentiate whether a volume is configured for data deduplication, except to add the media database attribute if the volume is deduplicated. The media database attribute, `*MSFT_OPTIMIZED_DEDUP_ENABLED`, is set to true and is saved as part of an optimized data deduplication volume save set.

For Windows BMR, the Windows Server 2012 and Windows Server 2012 R2 data deduplication writer is not part of the system state. Additionally, data deduplication volumes can be critical volumes and are supported with Windows BMR.

Unoptimized full and incremental backup

NetWorker creates an unoptimized data deduplication backup under the following conditions:

- When you specify in the save set attribute of the client resource, a backup path that is a subdirectory of the volume, except in the case where the subdirectory is the root of a mount point.

- When you perform a manual backup of the client that does not make up the entire volume.
- When you specify the string `VSS:NSR_DEDUP_NON_OPTIMIZED=yes` in the **Save Operations** settings of the client resource. If the save operation flag is set to yes the data deduplication backup is not optimized. If no string is present, or if the attribute is set to no, a normal volume level backup is performed.

To add this string, perform the following steps:

1. From the NetWorker **Administration** window, select the **Protection** menu.
2. In the left navigation pane, select Clients, right-click the client, and then select **Modify Client Properties**.
3. On the client **Properties** text box, select the **Apps & Modules** tab.
4. In the **Save operations** field, enter the string and attribute setting and then click **OK**.

In an unoptimized data deduplication backup, all files are rehydrated before the back up is performed. The deduplication chunk store directory is not backed up.

Windows dedup backups, either optimized or unoptimized, will be corrupt if they are backed up with VSS off.

Reasons to create an unoptimized data deduplication volume backup include:

- Support restores of a Windows Server 2012 and Windows Server 2012 R2 backups to an earlier version of Windows Server.
- Support restores of a Windows Server 2012 Windows Server 2012 R2 backups to a non-Windows computer.

Full volume restore to original path on the original computer

NetWorker supports a restore to the original volume mount path on the original server. All optimized files newer than the backup time of the restore save sets are rehydrated to prevent data loss.

When a deduplicated CSV volume is restored, CSV ownership is moved to the cluster node where the restore is being performed. This ensures that deduplication jobs and data access can be disabled during the restore process. The CSV is assigned back to original ownership when the restore is complete.

Full volume restore to original path on a different computer

NetWorker supports a restore of a data deduplication backup from one computer to the same volume mount path on another compatible computer. Part of this type of restore includes validation checks to ensure that Windows Server 2012 or Windows Server 2012 R2 is installed on the target computer and that the deduplication role is enabled.

You can manually reformat the volume, but this is not a requirement for NetWorker. The restore can only take place if the volume does not have a pre-existing chunk store. Additionally, the volume will be enabled for data deduplication after the restore is complete.

Support for save set restore of level FULL backups

A save set restore of a FULL backup is identical to a full volume restore with the following limitations:

- Limited to level Full backups in order to maintain chunk store integrity.
- Limited to volume level restores to the same path on the same computer where the backup was performed.

- No support for selective file restores due to insufficient information about the save set's restore context.
- No support for save set recovery of a Windows Optimized Deduplication volume backup. If you see the following sample error message, perform either a complete save set restore or select files to restore to an alternate location:

```
Failed to recover save sets: Unable to relocate a Windows deduplication save set, so will not recover save set 123xxxxxxx. Either restore the whole save set to its original location or select the files to restore to an alternate location using the Recover operation.
```

In the preceding sample error message, "123xxxxxxx" represents the save set ID.

File level restore

File level restore is performed if the volume to be restored is a subset of the original volume or if the restore is to a different volume. All files are restored in rehydrated form. The data deduplication meta data and chunk stores are not restored. For file level restores, the system account of the host where the restore is performed has to be a member of the NetWorker server's NetWorker Operators User Group. For example, if you are performing a dedup file level restore on host1, add system@host1 to the group.

NOTICE If an optimized deduplication restore is aborted, it is likely to have mismatched reparse point and chunk store entries. This restored volume is not a valid restore. You must restore the backup again and allow the restore process to complete.

Windows Data Deduplication Volume Best Practices

Review the following information, which describes the recommended best practices when you backup volumes that have Windows data deduplication enabled.

- A full backup should be performed immediately after deduplication has been enabled on a volume.
- Windows performs garbage collection on the chunk store of each deduplicated volume to remove no-longer-used chunks. By default, a garbage collection job is scheduled weekly for data deduplicated volumes. A full backup should be scheduled to run after garbage collection, because the garbage collection job may result in many changes in the chunk store, as a result of file deletions since the last garbage collection job.
- If there is significant chunk store container activity, control the size of incremental backups by limiting the frequency of Windows deduplication optimization jobs.
- Avoid performing extremely large file level restores. If a large percentage of a volume is restored, it is more time efficient to restore the entire volume. Because file level restores recover files in rehydrated form, a file level restore that includes many files might take up more space than is available on the volume.
- If a large file level restore is to be performed, first perform a full backup of the volume in its current state.
- When you choose to unoptimize many files at once from an optimized deduplication backup, the process can take a significant period of time. The selected files restore feature is best used to restore a moderate number of files. If most of a volume is to be restored, a full volume restore is a preferred solution. If a small amount of data needs to be skipped, that data can be moved to a temporary storage area, then back to its original location after the volume level restore is completed.

Recommended Deduplication Workloads

Based on recommendations by Microsoft, the ideal workloads for data deduplication include:

- **General file shares:** Group content publication/sharing, user home folders and profile redirection (offline files)
- **Software deployment shares:** Software binaries, images, and updates
- **VHD libraries:** VHD file storage for provisioning to hypervisors

For NetWorker, AFTD device directories are good candidates for deduplication. AFTD directories contain a large number of redundant data blocks, which in general are infrequently accessed.

Short filenames

You can back up and recover the short filenames that are automatically assigned by the Windows filename mapping feature.

Windows filename mapping is an operating system feature in which each file or folder with a name that does not conform to the MS-DOS 8.3 naming standard is automatically assigned a second name that does. For example, a directory named `Microsoft Office` might be assigned a second name of `MICROS~2`.

Backup and recovery of short filenames is disabled by default to improve performance. To enable backup and recovery of short filenames on a client, select the **Short filenames** checkbox on the **Globals (2 of 2)** tab of the **Client Properties** dialog box for the client resource.

You must enable diagnostic mode view by selecting **View > Diagnostic Mode** in the **Administration** window to access the **Short filenames** checkbox.

Volume mount points

You can back up and restore data available through a volume mount point (or mount point) on a Windows client.

Assigning a drive letter to a mount point is optional. Many disk volumes can be linked into a single directory tree, with a single drive letter assigned to the root of the host volume.

To include mount points in scheduled backups for a client, specify the host volume and each mount point in the **Save set** attribute on the **General** tab of the **Client Properties** dialog box for the Client resource. For example, to back up a single mount point on drive `D:` and all its data, type `D:\mount_point_name` in the **Save set** attribute.

To include nested mount points in scheduled backups, either use the `ALL` save set or specify the host volume and the full path to each mount point. For example, to back up three nested mount points and their data on drive `D:`, type the following values in the **Save set** attribute:

```
D:\mount_point_name1
D:\mount_point_name1\mount_point_name2
D:\mount_point_name1\mount_point_name2\ mount_point_name3
```

To include mount points in a manual backup with the NetWorker User program, select the checkbox next to the mount point name within the host volume entry in the **Backup** window.

To perform a manual backup of nested mount points and their data, perform a separate backup for each mount point. When you select a mount point in the **Backup** window, all files, directories, and nested mount points beneath the mount point are

selected by default. Before you start the backup, clear the checkboxes next to any nested mount points. Then perform separate backups for the nested mount points.

Windows file system backups

You can configure NetWorker to use VSS technology to backup file systems on a Windows host. You can recover individual file system objects from a VSS backup.

Overview of VSS

If the NetWorker Module for Microsoft is installed on the client computer, information in this chapter might be superseded by information in the NetWorker Module for Microsoft documentation. The *NetWorker Module for Microsoft Administration Guide* provides more information about the NetWorker Module for Microsoft.

Volume Shadow Copy Service (VSS) is a Microsoft technology that acts as a coordinator among all the components that create, archive, modify, back up, and restore data, including:

- The operating system
- Storage hardware
- Applications
- Utility or backup programs, such as NetWorker software

VSS allows for the creation of a point-in-time snapshot, or temporary copy, of a volume. Instead of backing up data directly from the physical file system, data is backed up from the snapshot. In addition, VSS allows for a single, point-in-time capture of the system state.

NetWorker uses VSS technology to create snapshot backups of volumes and exact copies of files, including all open files. Databases and files that are open due to operator or system activity are backed up during a volume shadow copy. In this way, files that have changed during the backup process are copied correctly.

Shadow copy (snapshot) backups ensure that:

- Applications can continue to write data to the volume during a backup.
- Open files are not omitted during a backup.
- Backups can be performed at any time, without locking out users.

 **Note:** VSS backups do not use snapshot policies, which are required to perform snapshot backups. The *NetWorker Snapshot Management Integration Guide* documentation provides more information.

VSS and the backup process

In VSS terms, NetWorker software is a requestor — an application that needs data from other applications or services. When a requestor needs data from an application or service, this process occurs:

1. The requestor asks for this information from VSS.
2. VSS reviews the request for validity.
3. If the request is valid and the specified application has the requested data, the request goes to the application-specific writer, which prepares the requested data.


Each application and service that supports VSS has its own writer, which understands how the application or service works:

1. After the writer signals that it has prepared the data, VSS directs the writer to freeze I/O to the selected volumes, queuing it for later processing.
2. VSS then calls a *provider* to capture the requested data.
3. The provider, which is either software-based or associated with particular hardware (for example, a disk array), captures the prepared data, creating a snapshot (or shadow copy) that exists side-by-side with the live volume. [Provider support](#) on page 423 contains more information.

The process of creating a snapshot involves interaction with the operating system. The amount of time it takes to create a snapshot depends on a number of factors, including the writer activity taking place at the time. Once the snapshot is created, the provider signals VSS, which tells the writer to resume activity. I/O is released to the selected volumes and any queued writes that arrived during the provider's work are processed.

The following figure provides a graphical representation of the VSS backup process.

Figure 46 VSS backup process



This figure provides a graphical representation of the VSS backup process:

1. NetWorker software (the requestor) asks VSS to enumerate writers and gather their metadata.
2. Writers provide an XML description of backup components and define the recover method.
3. VSS asks which providers can support a snapshot for each of the required volumes.
4. Requestor asks VSS to create snapshot.

Figure 46 VSS backup process (continued)

5. VSS tells the writers to freeze activity.
6. VSS tells the providers to create the snapshot of the current state on disk.

VSS tells the writers to resume activity.

NetWorker software backs up data from the point-in-time snapshot that is created during this process. Any subsequent data access is performed on the snapshot, *not* the live (in-use) file system. The requestor has no direct contact with the provider; the process of taking a snapshot is seamlessly handled by VSS. Once the backup is complete, VSS deletes the snapshot.

Provider support

By default, the NetWorker client always chooses the Windows VSS system provider for backups. If you want to use a hardware provider or a specific software provider for a particular NetWorker client, enter the following command in the NetWorker client resource Save Operations attribute:

VSS:VSS_ALLOW_DEFAULT_PROVIDER=yes

When the previous command is specified for a NetWorker client, a backup provider is selected based on the following default criteria as specified by Microsoft:

1. If a hardware provider that supports the given volume on the NetWorker client is available, it is selected.
2. If no hardware provider is available, then if any software provider specific to the given NetWorker client volume is available, it is selected.
3. If no hardware provider and no software provider specific to the volumes is available, the Microsoft VSS system provider is selected.

[Controlling VSS from NetWorker software](#) on page 424 provides more information about specifying VSS commands for a NetWorker client. [VSS commands](#) on page 426 provides information about other VSS commands.

NOTICE Windows Bare Metal Recovery backups always use the Windows VSS system provider even if the VSS:VSS_ALLOW_DEFAULT_PROVIDER=yes command is specified for the NetWorker client resource.

Troubleshooting hardware providers

If you have specified the VSS:VSS_ALLOW_DEFAULT_PROVIDER=yes command as described in [Provider support](#) on page 423 and the hardware provider and NetWorker are incompatible, try one of the following workarounds:

- Uninstall the hardware provider.
- Migrate any data that is backed up by the NetWorker client to a disk LUN (Logical Unit Number), such as C:\, that is not controlled by a hardware provider. In this way, the NetWorker client will backup all data using the software provider.

Be aware that if the NetWorker Module for Microsoft is installed on the client host, then the previously mentioned workarounds may not be required. Refer to the NetWorker Module for Microsoft documentation for details.

The importance of writers

Writers play an important role in correctly backing up data. They provide metadata information about what data to back up, and specific methods for correctly handling components and applications during backup and restore. They also identify the type of

application or service that is being backed up. Writers do *not* play a role in backing up the file system.

Writers are currently only available for active services or applications. If a service or application is present on a system but is not active, information from its writer is not available. Consequently, a writer can appear or disappear from backup to backup.

Also, NetWorker software maintains a list of supported writers in the NSRLA database of the client computer. When backing up data, the software checks to ensure that these conditions exist:

- The writer that is associated with the application is present on the system and active.
- The writer appears on the list of supported writers in the NSRLA database.
- A user has not disabled the writer.

If these conditions are all true for a particular writer, NetWorker software defaults to backing up data by using VSS technology. If any of the conditions are false for a particular writer, the data that is served by that writer is excluded from the backup operation.

List of supported writers

During a VSS backup operation, NetWorker software validates each writer against a list of supported writers. As part of a software release, or between releases, there may be updates to the list of supported writers. The *NetWorker E-LAB Navigator* provides a list of the currently supported writers.

Controlling VSS from NetWorker software

By default, NetWorker uses VSS technology to back up a client. For VSS SYSTEM save sets, this means NetWorker software uses VSS for most save sets and writers. For the file system, this means the software tries to take a snapshot of each drive, but if it fails, then it saves the file system by using the legacy method (that is, no snapshot is taken). During a particular backup for an individual client, either the VSS method or the legacy method is used, but not both.

There may be times when you need finer control over how NetWorker software uses VSS. For example, if you must disable VSS. You can control VSS from the Administration window, the NetWorker User program, or the command prompt.

Controlling VSS from the Administration window Procedure

1. From the **Administration** window, click **Protection**.
2. Click **Clients**.
3. Right-click the client for which you want to control VSS, then select **Properties**. The **Properties** dialog box appears, with the **General** tab displayed.
4. Click the **Apps & Modules** tab.
5. In the **Save Operations** attribute, type the command, then click **OK**.
 - Separate multiple commands with a semicolon (;).
 - If the **Save Operations** attribute is left blank, NetWorker software backs up data by using VSS.

Notes:

- The **Save Operations** attribute does not support NetWorker Module save sets. If a NetWorker Module save set name is entered in the window, the backup fails.

- If you enter a VSS command in the **Save Operations** attribute of the **Administration** window, the command runs when the client backup is started as part of a save set.
- Use the **Save Operations** attribute only for clients running NetWorker software release 7.2 or later. If anything is entered in this attribute for a client that is running an earlier NetWorker software release, the backup will fail.

Control VSS from the command-prompt

You can control VSS from the command-prompt on a NetWorker client or the NMC server by using the `-o` option and the **Save Operations** commands, but only while performing a `save`, `savefs`, or `nsrchive` operation.

For example, to completely disable VSS while backing up `C:\myfile` to the server `jupiter`, type:

```
save -s jupiter -o "vss:*=off" "C:\myfile"
```

Although the server name is not required in the preceding command example, include the name to ensure that the `save` command finds the correct server. Separate multiple **Save Operations** commands with a semicolon (;).

The *NetWorker Command Reference Guide* provides more information about the `save`, `savefs`, and `nsrchive` commands.

Note: If you change the VSS setting on a client by using the **Local Save Operations** dialog box or the command prompt, it does not affect that client's VSS setting on the server. Likewise, if you change a client's VSS setting on the server, it does not affect the **Local Save Operations** setting or the command-prompt VSS setting on the client.

Globally disabling VSS

Use the `nsradmin` program to disable VSS for all clients globally or only for clients with a certain Windows operating system.

About this task

To disable VSS for all NetWorker clients, perform the following steps:

Procedure

1. Log in as root or as Windows Administrator on the NetWorker server.
2. Create an input file for the `nsradmin` command. The input file eliminates interactive prompting as each client gets updated.
3. Run the `nsradmin` command and specify the input file.
4. Create an input text file. For example, create a file that is named `disable-vss-nt.txt` and type the following into the file:
5. Type the following at the command prompt:

```
nsradmin -i <path>\disable-vss.txt nsradmin -i <path>\disable-  
vss-nt.txt  
where <path> is the directory location of the input file.
```

Example 8 Example: Disable VSS for all NetWorker clients
Example: Disable VSS for all Windows NetWorker clients

Example 8 Example: Disable VSS for all NetWorker clients
Example: Disable VSS for all Windows NetWorker clients (continued)

1. Create a text file that is named `disable-vss.txt` , and then type the following into the file:

```
show name; client OS type; Save operations
print type: NSR client
update Save operations: "VSS\*=off"
print
```

2. Type the following command at the command prompt:

```
nsradmin -i <path>\disable-vss.txt
```

where `<path>` is the directory location of the input file.

1. Create a text file that is named `disable-vss-nt.txt` , and then type the following into the file:

```
show name; client OS type; Save operations
print type: NSR client; client OS type: "Windows NT Server on
Intel"
update Save operations: "VSS\*=off"
print
```

2. Type the following command at the command prompt:

```
nsradmin -i <path>\disable-vss-nt.txt
```

where `<path>` is the directory location of the input file.

VSS commands

This section lists the commands and syntax that are used to control VSS.

Table 71 VSS Save operation attribute values

Task	Save operations attribute	Behavior
To enable VSS.	Blank	Leaving the attribute empty results in NetWorker software automatically using VSS.
To completely disable VSS.	VSS*=off	VSS backups will not occur and backing up the following save sets for a NetWorker client resource yields these results: • DISASTER_RECOVERY:\ save set Backup fails at the beginning of backup operation. • All save set

Table 71 VSS Save operation attribute values (continued)

Task	Save operations attribute	Behavior
		Backups fail.
To use a hardware provider or a specific software provider for a NetWorker client backup.	VSS:VSS_ALLOW_DEFAULT_PROVIDER=yes	A backup provider is selected based on the following default Microsoft criteria: If a hardware provider that supports the particular volume on the NetWorker client is available, it is selected. If no hardware provider is available, then if any software provider specific to the particular NetWorker client volume is available, it is selected. If no hardware provider and no software provider specific to the volumes is available, the Microsoft VSS system provider is selected. Windows Bare Metal recovery backups always use the Windows VSS system provider even if the VSS:VSS_ALLOW_DEFAULT_PROVIDER=yes command is specified for the NetWorker client resource. Windows Bare Metal Recovery on page 427 provides more information about Windows Bare Metal recovery backups.

Windows Bare Metal Recovery

You can configure a Windows Bare Metal Recovery (BMR) backup on a Windows host. NetWorker Windows BMR is an automated recovery solution that uses the Windows ASR writer and other Microsoft VSS writers to identify critical volumes and perform a

full recovery on a target host. You cannot recover individual file system objects from a Windows BMR backup.

Terminology

The following list provides a description of typical Windows BMR backup and recovery terminology. The road map indicates which steps you must perform before you try a Windows BMR recovery.

This chapter uses the following terms to describe NetWorker support for Windows BMR technology:

Bare Metal Recovery (BMR)

The operation that restores the operating system and data on a host after a catastrophic failure, such as a hard disk failure or the corruption of critical operating system components. A BMR is an automated process that does not require the manual installation of an operating system. NetWorker provides an automated BMR solution for Windows that uses the Windows ASR writer and other Microsoft VSS writers to identify critical volumes and perform a full recovery on a disabled computer.

Offline recovery

A restore operation that is performed from the NetWorker Windows BMR boot image. A BMR recovery is an offline recovery. You cannot select specific files or save sets to recover during an offline recovery. You must perform an offline recover to the same or similar hardware.

Online recovery

A restore operation that is performed from the NetWorker User interface or recover command. An online recovery requires you to start the computer from an installed operating system and enables you to recover only specific files or save sets. The topic Recovering file system data provides more information about online recoveries.

Application data

User data that an application creates, such as log files or a database. For example, the application data of a SQL server includes databases and log files. You cannot use Windows BMR to recover the application data. You must back up and recover application data with NetWorker Module for Microsoft (NMM).

ASR writer

The Volume Shadow Copy Service (VSS) writer that identifies the critical data that NetWorker must back up to perform an offline recovery.

Boot Configuration Data (BCD)

A data store that contains a description of the boot applications and boot application settings that start the Windows operating system. To perform an offline recovery, you must back up this ASR writer component.

Critical volume

One of the following:

- Any volume that contains files for an installed service. The volume can be mounted as an NTFS directory. Exchange 2010 is an example of an installed service, but the Exchange database and log files are not considered critical.
- Any parent volume with a mounted critical volume.

 **NOTICE** NetWorker considers all volumes on all dynamic disks critical if at least one of the volumes is critical.

A Windows BMR recovery requires a current backup of all critical volumes.

Recovery

The restoration of the operating system and data for a host after a catastrophic failure, such as a hard disk failure or the corruption of critical operating system components. The recovery operation might be an offline recovery (Windows BMR) or an online recovery.

NetWorker Windows BMR image

A bootable image that contains the NetWorker binaries and a wizard to control the Windows BMR recovery process.

Non-critical volume

A volume that contains user data and does not contain installed applications that run as a service.

System State data

All the files that belong to VSS writers with a usage type of BootableSystemState or SystemService. You require these files to perform an offline recovery.

User data

Data that users generate, typically for the purposes of a business function. For example, a Microsoft Word document or an Excel spreadsheet. Windows BMR does not back up or recover user data unless the data resides on a critical volume. The simplest way to back up all user data is to specify the keyword All in the backup save set of the client resource. You can recover user data online at any time (on demand) or after a Windows BMR recovery operation.

WinPE

A bootable stripped-down version of the Windows operating system. The NetWorker Windows BMR image contains a customized WinPE with NetWorker binaries and a wizard to control the offline recovery process. WinPE does not support writers, except for the ASR writer. Therefore, VSS writers are not available with a NetWorker Windows BMR.

Overview of Windows Bare Metal Recovery (BMR)

Bare Metal Recovery (BMR) is data recovery and restoration where the backed up data is available in a form that allows you to restore a system from bare metal, that is, without any requirements as to previously installed software or operating system. Typically, the backed up data includes the necessary operating system, applications, and data components to rebuild or restore the backed up system to an entirely separate piece of hardware. The hardware receiving the restore should have a similar configuration as that of the hardware that was the source of the backup. The basic BMR is the process of bringing up a server after a disaster and ensuring that the system recovers with the operating system, the applications, and the data as they were at the time of the failure.

NetWorker Windows BMR is an automated recovery solution that uses the Windows ASR writer and other Microsoft VSS writers to identify critical volumes and perform a full recovery on a target host.

NetWorker Windows BMR supports file system backup and recovery of critical volumes. NetWorker Module for Microsoft (NMM) supports application data backup

and recovery. Additional backup and recovery procedures are required to backup and restore application data. The NMM documentation provides specific instructions on how to backup and recover applications.

You can use Windows BMR to recover a backup from a physical host. You can also use Windows BMR to recover a VMware virtual machine or VMware CD to a physical host, VMware virtual machine, or a VMware CD.

NetWorker uses a special save set called `DISASTER_RECOVERY:\`, a subset of the `ALL` save set, to backup all the data that is required to perform a Windows BMR. NetWorker performs the BMR backup while the Windows operating system is active. You can recover an offline BMR backup without first reinstalling the Windows operating system. This action prevents problems that can occur when you restore operating system files to a running version of Windows.

To support a NetWorker Windows BMR recovery, download the Windows BMR image from Online Support website. This image enables you to create a bootable Windows BMR ISO that contains NetWorker binaries and a wizard, which controls the recovery process.

 **Note:** The *NetWorker E-LAB Navigator* provides more information about operating systems support for Windows BMR.

Components of the `DISASTER_RECOVERY:\` save set

The `DISASTER_RECOVERY:\` save set contains a group of component save sets that are required to perform a Windows BMR recovery. A full backup of the `DISASTER_RECOVERY:\` save set contains the following components:

- All critical volumes.
- `WINDOWS_ROLES_AND_FEATURES:\` (a subset of the `DISASTER_RECOVERY:\` and `ALL` save sets).
- System Reserved partition.
- UEFI partition (if available).

NetWorker supports full and incremental backup levels of the `DISASTER_RECOVERY:\` save set. Also, when the Windows BMR recovery operation recovers data from an incremental backup, the recovery operation recovers all incremental backups.

The first time NetWorker performs a backup of the `DISASTER_RECOVERY:\` save set, NetWorker performs a level Full backup, regardless of the level that is defined for the backup.

When you configure a level Incremental backup of the `DISASTER_RECOVERY:\` save set, NetWorker backs up some components of the save set at a level Full, and other components at an Incremental level.

The following table summarizes the backup level of each save set component of the `DISASTER_RECOVERY:\` save set, when you perform an incremental backup:

Table 72 `DISASTER_RECOVERY:\` components in an incremental backup

Save set	Backup level
Critical volumes	Incremental
<code>WINDOWS_ROLES_AND_FEATURES:\</code>	Incremental
UEFI partitions	Full

Table 72 `DISASTER_RECOVERY:\` components in an incremental backup (continued)

Save set	Backup level
System reserved partition	Full

During an incremental backup, the backup operation checks both the modification time and the archive bit to determine if a file must be backed up. The backup operation ignores the archive bit when you assign the `nsr_avoid_archive` variable a value of **Yes** on the client host. As a result, NetWorker only uses the modification time to determine which files to back up.

Use the environment variable `nsr_avoid_archive` with caution. If you use the environment variable `nsr_avoid_archive`, test the BMR backup image to ensure that you can recover the Windows system state correctly. [Performing a BMR recovery to a physical computer](#) provides more information on validating the BMR backup image.

A Windows BMR recovery requires a successful backup of each component save set in the `DISASTER_RECOVERY:\` save set. If one component of the save set fails, then the backup operation fails. For a scheduled backup, NetWorker retries the `DISASTER_RECOVERY:\` backup. The number of retries that NetWorker performs is based on the value that is defined in the client retries attribute of the protection group that the Client resource is assigned to.

Note: In NMC Administration GUI, the **Log** tab of the **Monitoring** window, or the **Save Set** tab of the **Media** window displays each component save set of a `DISASTER_RECOVERY:\` backup.

Critical volumes

This topic describes critical volumes and the associated management tools.

NetWorker considers a volume as critical when it contains files for an installed Windows service. NetWorker also considers the following volumes as critical and will include the volumes in a `DISASTER_RECOVERY:\` backup:

- A non-critical volume that has a critical volume mounted on it, or a non-critical volume that serves as a parent to a critical volume.
- All volumes on a dynamic disk when one of the volumes critical. If one disk in a dynamic disk pack is critical, then NetWorker must treat all disks in that pack as critical. This can substantially increase the number of disks that NetWorker includes in the BMR backup. It is recommended that you do not install services on a dynamic disk.

Note: By default, the Windows 2012 System Writer does not report Win32 Service Files as a part of systems components. As a result, the volumes that contain Win32 Service Files are not considered critical and the `DISASTER_RECOVERY:\` save set will not include a volume that contains files for an installed service. To configure the Windows 2012 server to report Win32 Service Files as a part of system components, set the `ReportWin32ServicesNonSystemState` registry sub key to 0. Microsoft KB article 2792088 provides more information.

A Windows BMR backup does not back up the following files on a critical volume:

- Files listed in the `FilesNotToBackup` registry key.
- Files excluded by system writers.
- Files that an application VSS writer backs up. For example, Exchange databases. Use NetWorker Module for Microsoft Applications (NMM) to backup these files.

Excluded critical volumes during a Windows BMR backup

A NetWorker Windows BMR backup excludes critical volumes based on the operating system, disk types, configuration and installation of your computer.

Install applications with third-party services on the system disk, or a disk that already has other services installed. To identify the disks that contain third-party services, use the utility, `list writers detailed` command.

For Windows Server 2008 R2 SP1, set the `ExcludedBinaryPaths` registry key to exclude third-party services from the System Writer. This prevents the disk where the service is installed from being classified as critical.

NetWorker excludes a volume from a backup when one of the following Windows application service is installed on the host:

- Storage Spaces volume
- Cluster volume
- Cluster Shared Volume



NOTICE To ensure that you can recover all required files, perform a file system backup of any excluded disk.

Displaying a list of the critical volumes

To view a list of the critical volumes for a NetWorker client, type the NetWorker command `save -o VSS:LCV=yes` from the command line on the client host.

For example:

```
NetWorker_install_path\bin>save -o VSS:LCV=yes
```

Output similar to the following appears:

```
The following volumes are determined as critical by the system
state writers:
C:\ (disk num 0)
i:\mount\ (disk num 7)
The following volumes are critical because they are parents for
one or more mounted critical volumes:
i:\ (disk num 2)
The following volumes are critical because they are in the
same dynamic disk pack with one or more critical volumes:
H:\ (disk num 4,5)
i:\ (disk num 2)
```

WINDOWS ROLES AND FEATURES save set

The `WINDOWS ROLES AND FEATURES` save set was introduced in NetWorker 8.1 and replaces the `VSS SYSTEM BOOT`, `VSS SYSTEM FILESET` and `VSS SYSTEM SERVICES` save sets. The `DISASTER_RECOVERY:\` save set contains the `WINDOWS ROLES AND FEATURES` save set as a component save set.

The `WINDOWS ROLES AND FEATURES` save set contains:

- Data that are associated with the roles and features that are installed on the Windows server.
- Metadata that represents the volume data which the `ALL` or `DISASTER_RECOVERY:\` save set backs up.

Before backing up the `WINDOWS ROLES AND FEATURES` save set, consider the following:

- **Block Based Backups (BBB)** do not support the `WINDOWS ROLES AND FEATURES` save set.
- You cannot restore the `WINDOWS ROLES AND FEATURES` save set simultaneously with data from a file system backup. If you must recover data from both the `WINDOWS ROLES AND FEATURES` backup and a file system backup, restore the file system data first, and then restore the `WINDOWS ROLES AND FEATURES` data.
- The NetWorker software automatically backs up AD as a component of the `WINDOWS ROLES AND FEATURES` save sets. An AD backup or restore includes the AD log files, database, patch files, and expiry token.
- You can perform an online recovery of the `WINDOWS ROLES AND FEATURES` save set to recover the Active Directory, DFSR, or Windows Server Failover Cluster services. The topic [Online recovery of Active Directory, DFSR, or Cluster service](#) provides more information.
- If you cancel a deduplication recovery, the state of the recovered data is not reliable and may contain corrupted data. To ensure that the recovery is correct, restart the deduplication recovery process.
- The backup operation only confirms that the VSS System Writer exists on the target host. If the backup operation does not detect the writer, the backup of the `DISASTER_RECOVERY:\` or `ALL` save set fails. The backup operation does not track and report any other missing VSS writers.
- You can perform a component level granular restore of the `WINDOWS ROLES AND FEATURES` save set with a command line recover or the NetWorker User application. For example, you can recover the system state and replication folders separately. You cannot use the NMC Recovery UI to perform a component level restore.
- Do not restore the `WINDOWS ROLES AND FEATURES` system state multiple times in succession without restarting the computer as required. If you do not restart the computer, you can put the system in an unreliable operational state.

Note: The NetWorker 8.2 and later clients can only recover `WINDOWS ROLES AND FEATURES` save sets. If you try to recover a VSS System State save set that was created with a NetWorker 8.0 SP1 client or earlier, then the Windows host will not function correctly. To recover VSS system state save sets that are created with a NetWorker 8.0 SP1 or earlier backup, use the NetWorker 8.0 SP1 or earlier client to create a backup. It is recommended that you restore the `WINDOWS ROLES AND FEATURES` save set from a NetWorker 8.1 or later backup.

The `DISASTER_RECOVERY:\` save set

The `DISASTER_RECOVERY:\` save set is available for Windows clients.

The `DISASTER_RECOVERY:\` save set backs up critical volumes, UEFI, the system reserved partition, and `WINDOWS ROLES AND FEATURES`.

The `DISASTER_RECOVERY:\` save set does not include data for clusters, Active Directory, DFS-R, and Windows Server Failover Cluster.

Checkpoint restart is not supported for backups of the `DISASTER_RECOVERY:\` save set. If you enable checkpoint restart for a client with the `DISASTER_RECOVERY:\` save set, then the setting is quietly ignored for the `DISASTER_RECOVERY:\` save set. The save set is marked with a `cb` flag instead of a `k` flag, indicating that the checkpoint is not considered for `DISASTER_RECOVERY:\`.

The `DISASTER_RECOVERY:\` save set is also in the `ALL` save set.

UEFI Partition Support

NetWorker supports a backup and recovery of unmounted Unified Extensible Firmware Interface (UEFI) partitions on hosts that use a supported operating system. The *NetWorker E-LAB Navigator* provides more information about supported operating systems.

The topic [Performing a Windows BMR recovery to a physical computer](#) describes how to perform a Windows BMR of a computer that has UEFI partitions.

The following list summarizes the properties of a UEFI partition backup:

- NetWorker can backup an unmounted partition.
- NetWorker uses the following path pattern to backup the UEFI partitions:
`\\<root>\Device\HarddiskVolume#`
 where # is the number of the volume.
- The `DISASTER_RECOVERY \:` save set contains a backup of the UEFI partitions.
- NetWorker always performs a level Full backup of UEFI partitions, regardless of the backup level of the `DISASTER_RECOVERY \:` save set.
- NetWorker does not index the UEFI partitions or make the UEFI partitions available for online recoveries.

After a successful BMR restore, a host that uses UEFI might fail to start. This can occur when the UEFI boot manager does not have a valid Boot Order entry, for example, when you delete the Boot Order entry or restore the Windows BMR backup to different hardware. In these situations, the operating system recreates the Boot Order entry during a restart operation but may not use the same path.

To resolve this issue, load **Boot Manager** and select **Boot** from the **File** menu to correct the Boot Order entry.

Boot Configuration Data

In earlier versions of the Windows operating system, the `BOOT` directory was present in the system drive. In Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2 SP1, Windows Server 2012, and Windows Server 2012 R2, a hidden, unmounted system-reserved partition can be present, and the Boot Configuration Data (BCD) store is on this partition. The BCD store contains the boot configuration parameters and controls the computer boot environment.

The NetWorker Windows client backs up the system reserved partition and the BCD store only for Windows offline Bare Metal Recovery (BMR). During a Windows offline BMR backup, NetWorker checks the type of operating system. If it is Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2 SP1, Windows Server 2012 or Windows Server 2012 R2, NetWorker assigns a GUID to the partition and performs the backup of the BCD. The BCD partition does not need to be mounted for the backup to occur. If the BCD partition is not mounted, the backup is not indexed. The save set name is `GLOBALROOT/xxxxxx/`.

The BCD can only be restored as part of offline BMR. Online recovery of the BCD is not available. Consult Microsoft documentation for using the `BCDEdit` tool to save copies of BCD before making Boot Configuration Data changes.

Windows Server 2012 Cluster Shared Volumes (CSV)

NetWorker does not support Windows Server 2012 Cluster Shared Volumes (CSV) as a critical volume. If a CSV disk is marked as a NetWorker critical disk, then the Windows BMR backup reports a warning, and continues to perform the backup operation as if

the CSV is not on the critical list. NetWorker does not backup the CSV because a CSV cannot reside in the same shadow copy set with a local volumes.

Applications such as SQL Server and Hyper-V in a Windows Continuous Availability scenario using CSV are not supported.

The *NetWorker Cluster Integration Guide* provides more details.

Windows Server 2012 Storage Spaces

NetWorker Windows BMR does not support the backup and recovery of critical System State data that are on virtual disks. A NetWorker BMR backup skips all critical volume data that are on Storage Spaces and does not add the volume to the BMR critical volume list.

A BMR recovery cannot recover critical volume data on Storage Spaces. If the Storage Pool disks that compose a Storage Spaces virtual disk are not damaged, a recovery operation to the original computer will mount the Storage Pool virtual disks after the critical volume recovery operation completes.

NOTICE It is recommended that you detach the physical disks that Storage Spaces use when you recover critical volumes, and then reattach the physical disks after recovery. A Windows BMR recovery operation can overwrite data on attached Storage Spaces disks.

The topic [Windows Storage Pools considerations](#) describes how to perform a Windows BMR recovery of Storage Spaces to a new computer.

NOTICE To backup and recover data on virtual hard disks and volumes that are created by Storage Spaces, use NetWorker file system backup and recovery operations.

A Windows BMR backup of a Windows 2012 host creates a file that is named `OSSR_sysinfo.xml`. The file is located in the `[root]\EMC NetWorker\nsr\tmp` directory. This file captures pertinent information about the configuration of the backed up host. For example:

- Host information (name, boot drive, BIOS, or EFI).
- NIC cards and their parameters.
- Disk information.
- Storage Spaces information.

The purpose of this file is to support the manual recreation of the Storage Spaces configuration following a BMR recovery.

Synthetic full backups

A synthetic full backup uses the most recent full and incremental backups to create a full backup without transferring any data from the client. NetWorker performs all the work to synthesize a full backup on the NetWorker server. A synthetic full backup gives you the benefits of a full backup, such as a faster restore, without having to perform a full backup.

The topic [Synthetic full backups](#) describes the synthetic full backup feature.

When a client backup includes the `DISASTER_RECOVERY:\ save set`, NetWorker will always backup volumes that are identified as critical, at a level full. NetWorker will not create a synthetic full backup for critical volumes. The `DISASTER_RECOVERY:\ save set` is included during full backups when either the `ALL` or `DISASTER_RECOVERY:\ save set` is specified in the NetWorker Client resource.

Example 9 Synthetic full backups with save set ALL

Example 9 Synthetic full backups with save set ALL

The save set attribute of the Client resource contains the `ALL` save set and the backup schedule includes a synthetic full backup on Sundays. The NetWorker client host has four volumes: two are critical, and two are non-critical.

- `C:\` and `E:\` are critical volumes.
- `F:\` and `G:\` are non-critical volumes.

On Sundays, NetWorker performs a backup of the following save sets:

- `C:\` — At a true level full backup level.
- `E:\` — At a true level full backup level.
- `F:\` — At a synthetic full backup level.
- `G:\` — At a synthetic full backup level.
- `DISASTER_RECOVERY:\` — At a true level full backup level.

Example 10 Synthetic full backups with file system save sets

The save set attribute of the Client resource contains a list of all volumes and the backup schedule includes a synthetic full backup on Sundays. The save set attribute does not contain the `DISASTER_RECOVERY:\` save set. The NetWorker client host has four volumes: two are critical, and two are non-critical.

- `C:\` and `E:\` are critical volumes.
- `F:\` and `G:\` are non-critical volumes.

On Sundays, NetWorker performs a backup of the following save sets:

- `C:\` — At a synthetic full backup level.
- `E:\` — At a synthetic full backup level.
- `F:\` — At a synthetic full backup level.
- `G:\` — At a synthetic full backup level.

Online recovery of Active Directory, DFSR, or Cluster services

The `DISASTER_RECOVERY:\` save set includes the `WINDOWS ROLES AND FEATURES` component save set. You can recover the `WINDOWS ROLES AND FEATURES` backup in an online recovery operation, to a host that uses the same Windows operating system instance. NetWorker 8.2 and higher support the online recovery of the following Windows services, which the `WINDOWS ROLES AND FEATURES` component contains:

Active Directory

Solve Desktop provides procedures that describe how to recover this service.

Distributed File System Replication (DFSR)

The topic, Backing Up and Restoring a Microsoft DFS, provides more information.

Cluster

Solve Desktop provides procedures that describe how to recover this service.

NetWorker does not support the online recovery of any other Windows service that the `WINDOWS ROLES AND FEATURES save set` contains. Unsupported online recovery of `WINDOWS ROLES AND FEATURES` components results in an inconsistent state of the Windows server.

NOTICE When you perform an online recovery, you cannot mark the `WINDOWS ROLES AND FEATURES save set` and use the Required Volumes option. To determine the volume that contains the `WINDOWS ROLES AND FEATURES save set` that you want to restore, mark the `DISASTER RECOVERY:\ save set`, then use the Required Volumes option. After you determine the required volumes, unmark the `DISASTER RECOVERY:\ save set` and mark the `WINDOWS ROLES AND FEATURES save set`.

Windows BMR Planning

This section provides guidelines on how to plan your Windows BMR backups.

Requirements for Windows BMR backup and restore

The BMR recovery process restores the operating system that was installed on the source host. If you perform a BMR recovery to a different host with different hardware, after the recovery operation and restart completes, Windows prompts you to install the required drivers.

Before you perform a BMR recovery to a different host, ensure that you meet the following requirements:

- The source and target hosts use the same processor architecture.
 - The hardware on the target host is operational.
 - The target host has a minimum of 512 MB of RAM.
 - The target host startup hard disk capacity must be larger or the same size as on the source host, regardless of the amount of space actually in use. If the disk is smaller by a single byte, BMR fails.
- Note:** Verify whether the source critical volumes are part of a larger physical disk. If critical volumes are on a larger physical disk, all target critical volumes must be large enough to accommodate the entire underlying physical disk. Use the Windows Disk Management utility to verify disk configuration and size.
- The number of disks on the target host is greater than or equal to the number of disks there were on the source host. The disk LUN numbering on the target host must match the disk LUN numbering on the source host.
 - The RAID configuration on the target host should match the disk order of the hard disks.
 - The disk or RAID drivers that are used on the source system must be compatible with the disk or RAID controllers in the target system. The recovery process restores the backup to the same logical disk number that was used by the source host. You cannot restore the operating system to another hard disk.
 - Windows BMR supports IDE, SATA, or SCSI hard disks. You can make the backup on one type of hard disk and recover on another type of hard disk. For example, SAS to SATA is supported.

- The target system can access the Windows BMR image as a bootable CD/DVD volume or from a network start location.
- The target system has the NIC or storage device drivers installed that match the NIC.

Note: All NIC or storage device drivers must not require a restart to complete the driver installation process. If the drivers require a restart, then the BMR recovery process fails and prompts you to install the drivers again.

Save set configuration by host type

This section describes the attributes of save sets that are used by Windows BMR. This information helps you select the correct save set configuration for the computer and operating system.

The following table lists the save sets to back up, depending on the Windows host to be protected.

Table 73 Save set configuration for a specific host

To back up this host	Specify these save sets in the client resource Save Set attribute	Considerations
A host or file server that is not a Microsoft Application server	• Specify the save set <code>All</code> in the NetWorker Client resource. • By default, the save set <code>All</code> includes the <code>DISASTER RECOVERY:\</code> save set and all of the local physical drives.	• <code>WINDOWS ROLES AND FEATURES</code> must be backed up. • <code>WINDOWS ROLES AND FEATURES</code> save sets are recovered in a Windows BMR operation and are also available for online recovery. <code>WINDOWS ROLES AND FEATURES</code> save sets should only be recovered online as part of an Active Directory, DFSR, or Windows Server Failover Cluster online recovery.
A host that is a Microsoft Application server. For example, a Microsoft Exchange Server, Microsoft SQL Server, Hyper-V, or Microsoft SharePoint Server	• Specify the <code>ALL</code> save set in the Save set attribute in the NetWorker Client resource. • Use NMM to back up the application databases. The NMM provides details.	• Use the Windows BMR Wizard to recover the data contained in the <code>DISASTER RECOVERY:\</code> save set. • Use NMM to recover the application databases.

Best Practices for Windows BMR

The following sections outline best practices for Windows BMR.

Perform regular backups

Perform a full backup that contains the `DISASTER_RECOVERY:\ save` set regularly and after any you install, remove or update any system components. For example, when you add, change, or remove Windows roles and features, or install Windows updates and service packs.

NetWorker will automatically back up the `DISASTER_RECOVERY:\ save` set when you specify the `ALL` save set in the Save Set attribute of the NetWorker Client resource.

Capture disk configuration changes for Windows BMR

The NetWorker BMR recovery operation uses the Microsoft ASR writer to reconstruct a disk configuration. The ASR writer is sensitive to the disk numbers and disk configuration of the original host. NetWorker saves this disk information during a Windows BMR backup and uses the disk configuration information to perform the recovery. After you reconfigure any disk on a host, reboot the host and then perform a Windows BMR backup to ensure that NetWorker captures the new disk configuration. Examples of a disk reconfiguration include the addition or removal of a disk or partition.

Mixing critical and non-critical volumes on a physical disk

Windows allows you to partition a physical disk into multiple volumes. These volumes can be either critical or non-critical, depending on the type of data that they contain. During a Windows BMR recovery operation, the ASR writer can re-create and format a partition, including non-critical partitions. If the ASR writer formats a non-critical partition, the use of an online recovery is required to recover data on the non-critical partitions. Recovering the Data describes how to perform an online recovery.

 **NOTICE** Do not mix critical and non-critical volumes on the same physical disk.

Considerations for NetWorker user defined directives

Use user defined directives, such as `nsr.dir`, with caution. When you use directives in directories where system state and installed services data resides, the backup creates an incomplete BMR backup image and potentially render the BMR backup image unusable. If you create user defined directives, test the BMR backup image to ensure that you can successfully perform a BMR Recovery. [Performing a Windows BMR recovery to a physical computer](#) provides more information about testing the BMR backup image.

Critical volume recommendations

Use the following practices to minimize the size of Windows BMR backups.

- Do not store non-critical data, such as MPEG files, on critical volumes.
- Consolidate critical volumes. For example, install services on the same disk.
- Do not mount critical volumes on a non-critical volume.

Windows BMR limitations and considerations

Review the following Windows BMR limitations and special considerations before you perform Windows BMR backup, clone and recovery operations.

Disk configuration limitations

This sections describes disk configuration limitations in Windows BMR.

Dynamic disks

A BMR recovery does not bring dynamic disk volumes online. After the BMR recovery completes, use Windows Disk Manager to bring the dynamic disks back online.

NTFS and ReFS

Only NTFS and ReFS file systems are recognized as critical volumes

Although the backup of the `DISASTER_RECOVERY:\ save set` fails, NetWorker will backup, the contents of the partition and the data is available for an online recovery only.

To ensure a successful backup of the `DISASTER_RECOVERY:\ save set`, install all services or application on an NTFS or ReFS volume.

Critical volumes

Windows BMR only supports critical volumes on NTFS and ReFS partitions. This is a Microsoft ASR limitation. If a critical volume is on a partition other than NTFS or ReFS, the backup of the `DISASTER_RECOVERY:\ save set` fails. A message similar to the following appears in the `policy.log` file:

```
Disaster Recovery: critical volume volumename identified for
disaster recovery backup has a non-NTFS file system,
filesystemname. Backups of non-NTFS critical volumes are not
supported.
```

Note: Windows BMR does not support FAT and FAT32 file systems as critical volumes.

HP ProLiant system considerations

You cannot recover from a Windows BMR backup on an HP ProLiant system when the HP i Provisioning Tool (IPT) 1.4 or 1.5 was used to configure an entire disk as a critical volume, such as the system partition.

To resolve this issue, shrink the logical volume before you perform the Windows BMR restore. The HP website contains a customer advisory that describes the issue and the impact to Windows Bare Metal Recovery with Windows Server Backup. This advisory and the resolution also applies to NetWorker Windows BMR critical volumes.

Note: It is recommended that you test your BMR solution before a disaster recovery is required.

Optimized deduplication backup considerations

Review this section before you configure backups that use optimized deduplication.

- You can recover a complete volume backup recovery to the original volume only if the backup was performed at a level Full.
- You cannot recover specific files from a level FULL or INCREMENTAL save set.
- You cannot perform a full volume recovery of a non-full level save set.
- You cannot recover data from an optimized and unoptimized deduplication backup when VSS is disabled. The backups that NetWorker created are corrupt.
- You cannot cancel the recovery of an optimized deduplication backup to a deduplication volume. If the recovery process is interrupted or fails, the destination volume becomes unusable. You must repeat the recovery process and the recovery operation must complete successfully to prevent volume corruption.
- If the optimized deduplication recovery cannot successfully complete, you can perform a selected files restore of directories from the optimized deduplication backup. This restores the directories' files to a rehydrated state, but will take significantly more time.

Save set considerations

This topic describes limitations and considerations that relate to save sets.

Checkpoint restart backup for Windows DISASTER_RECOVERY:\ save set is not supported

The NetWorker software does not support a checkpoint restart backup for the Windows `DISASTER_RECOVERY:\` save set. When you enable the Checkpoint restart option for a Client resource that you configure to back up the `DISASTER_RECOVERY:\` save set, the backup fails.

Including DISASTER_RECOVERY:\ in multiple save sets

When you use specify multiple save sets with the save command, you must use the `-N` option to specify the symbolic name of `DISASTER_RECOVERY:\` save set, and specify the `DISASTER_RECOVERY:\` as the last save set in the save set list.

For example:

```
save.exe -s server -N "DISASTER_RECOVERY:\"" save_set1 save_set2 ...
"DISASTER_RECOVERY:\""
```

where:

save_set1 or *save_set2* are unique save set names, such as a drive letter (`f:\`) or mount point (`n:\mountpoint`).

Monitoring save operations

When you monitor Windows BMR save operations, for example, by using the **NetWorker Administration > Monitoring > Sessions** window, you might notice that the number of save sessions differ from the number of save sets that appear in the **Save set** attribute of the Client resource. This is because NetWorker optimizes Windows BMR backups to generate the correct number of Windows BMR backup sessions and save sets.

Cloning considerations

To clone a Windows BMR backup, ensure that you clone all of the critical volumes, `DISASTER_RECOVERY:\`, and `WINDOWS ROLES AND FEATURES` save sets that were created during the backup operation. While you can clone individual save sets, you cannot perform a successful BMR recovery unless you recover each save set that the backup operation created.

To ensure that you clone all of the BMR save sets, review the following information before you start a clone operation:

- When you use the automatic clone, you enable the Clone attribute on the group resource that contains the BMR client. The automatic clone operation will clone all of the required save sets after the scheduled backup operation completes.
-  **Note:** Synchronize the NetWorker server and client host clocks before the backup operation to ensure that all of the save sets are cloned.
- When you use the `nsrclone` command to perform a manual clone, ensure that you include the *ssid/cloneid* for each save set. Use the `mminfo` or `nsrinfo -v` command to report all save set backups that occurred for the Windows client during the save session. The *Command Reference Guide* provides detailed information about using the `mminfo` and `nsrinfo` commands.
- When you use the schedule clone function, do not filter on other attributes such as save set name. Filter only by client name. When you enable automatic cloning for a backup group that contains the `DISASTER_RECOVERY:\` save set, synchronize

the clocks on the NetWorker server and client host clocks across the network to ensure that NetWorker clones all save sets.

Security considerations

This section describes security issues related to planning Windows BMR backup and recovery.

NetWorker support for Windows Encrypting File System (EFS)

This topic describes the behavior of EFA and BitLocker after you complete a BMR with NetWorker.

Windows BMR supports backup and recovery of files and folders encrypted with Windows Encrypting File System (EFS), and volumes encrypted with BitLocker. After BMR, the EFS or BitLocker services might be running but the EFS encryption attributes on files or folders must be re-enabled and BitLocker volumes must be re-encrypted. Consult Microsoft documentation for steps to encrypt with EFS and BitLocker.

If a folder is encrypted in Windows, for example, by selecting **Folder Properties > Advanced > Encrypt** contents to secure data, it is recovered as encrypted. However, the encryption attribute is not be set on the folder. You can manually reset the encryption attribute after the recovery operation. This is a Microsoft limitation.

Windows BMR and third-party encryption tools

This topic provides information on how to correctly validate Windows BMR when you use a third-party encryption tool.

NetWorker Windows BMR has not been thoroughly tested with third-party drive encryption products other than Microsoft's BitLocker. If you use a third-party drive encryption product, then validate the backup and recovery procedures by performing a Windows BMR backup and recovery to verify that the restored computer is fully functional. Perform the test against the original hardware and new hardware to confirm both scenarios. You must learn if any additional steps are required to reencrypt the drivers after a successful restore.

NetWorker Strong Authentication and Windows BMR recoveries

This topic describes how to use NetWorker strong authentication.

 **Note:** Starting with NetWorker 19.1, BMR can be performed for clients that uses nsrauth as authentication method.

When you recover a Windows client that uses NetWorker strong authentication (`nsrauth`) to communicate with other NetWorker hosts, communications with the NetWorker server may fail during a Windows BMR recovery. When you perform a Windows BMR recovery for a host that uses nsrauth authentication only, the Windows PE image does not have the nsrauth credentials file that the original client used, and it generates the new credentials. NetWorker server refuses to allow the recovery operation to complete as it is not aware of the new credentials.

To resolve this issue, before you perform the BMR recovery perform the following task:

- Delete the NSR Peer Information resource for the NetWorker Client from the NSRLA database on the NetWorker Server. This will cause the NetWorker Server to create a new NSR Peer Information resource for the client.

Server role considerations

This section describes considerations for Windows Server Roles in Windows BMR.

Protecting Windows server roles

Several server role components of Windows host store the data in a database.

Examples of Windows server roles with databases include:

- Active Directory Rights Management Services (AD RMS).
- Windows System Resource Manager (WSRM).
- Universal Description, Discovery, and Integrations (UDDI) Services.
- Windows Server Update Services (WSUS).

When you install the Windows server role on a host, the installation process prompts you to store data on either an existing SQL Server installation or in a Windows Internal Database (WID).

NetWorker uses the VSS SQL Server writer to back up the role databases that are stored in WID but does not protect role databases, which the server role component stores in a SQL Server. Use NMM or a third-party SQL backup product to backup and recovery the roles databases.

Backup and recovery workflows for server roles that use WID

These are the backup and recovery workflows are as follows:

- Perform a NetWorker Windows BMR backup, which includes all the SQL writer components for WID. If required, backup user data on the client.
- Perform a NetWorker Windows BMR recovery operation, which recovers all the WID components.

After the NetWorker Windows BMR system restart, the WID service is available and Windows server roles have access to their databases.

Saving and recovering SQL Server components with Windows BMR and NMM:

1. Perform a NetWorker Windows BMR backup. If required, backup user data on the SQL client.
2. Use NMM or a third-party backup application to back up the SQL Server application.
3. Perform a NetWorker Windows BMR recovery operation.
After the recovery and restart operations complete, you cannot start the SQL Server service. Also, any server roles that store data in SQL databases outside WID will not work.
4. For non-clustered SQL servers only, ensure that the SQL group is offline.
5. Run the following **setup.exe** command from a command prompt with elevated privileges, to rebuild the SQL Server:

```
C:\> setup /QUIET /ACTION=REBUILDDATABASE /
INSTANCENAME=Instance_name /SQLSYSADMINACCOUNTS=domain_name
\administrator
```

 **Note:** The SQL Server installation media contains the **Setup** tool.

6. Bring the SQL server services online.
7. Use NMM or a third-party backup application to recover the SQL system databases (master, model, msdb).

8. Use NMM or a third-party backup application to recover the role databases.
9. Restart the services that require the role databases that you recovered.

 **NOTICE** The *NetWorkerModule for Microsoft Applications Application Guide* provides more information about using NMM to recover SQL databases.

Microsoft server application considerations

Use both the NMM and the NetWorker software to protect Microsoft server applications, such as Microsoft Exchange Server, Microsoft SQL Server, Hyper-V, and Microsoft SharePoint. The NMM software protects the application data, such as databases and log files and the NetWorker client software protects the user data and critical disks on the host, for the purposes of Windows BMR.

Below is a high level overview of NetWorker and NMM backup and recovery workflow for Microsoft server applications:

1. Use NetWorker to back up critical and non-critical disks as part of a regular file system backup.
2. Use NMM to back up application data, such as Microsoft SQL Server.
3. Use NetWorker to perform a Windows BMR backup of the critical volumes on the host.
4. Use the Windows BMR boot image to perform a BMR recovery.
5. Use the NetWorker User application to recover any non-critical disks.
6. Use NMM to recover the application data.

The NetWorker Module documentation provides more information about recovering application data.

Online recovery of Windows services considerations

This section describes limitations and considerations that are related to Windows services.

Active Directory considerations

A Windows BMR recovery of a Domain Controller is non-authoritative by default. If you must perform an authoritative recovery, then you must start into DSRM mode directly from the Windows BMR wizard. The topic *Performing post-recovery tasks for Active Directory services*, provides more information.

DFSR considerations

DFSR namespaces are junction mount points. The `DISASTER _RECOVERY:\` and `ALL` save sets do not backup DFSR namespaces, even if the DFSR shares reside on a critical volume. To backup DFSR Shares, either use the new save set `ALL-DFSR` or provide the full DFSR Share path as the save set name. The `ALL-DFSR` save set applies to all supported platforms. Unlike the `ALL` save set, which skips the DFSR namespace because it is a junction point, the `ALL-DFSR` save set backs up every namespace, along with the associated replication folders.

The topic *Recovering Windows volume mount points*, provides more information about recovering volume mount points.

MSCS considerations

Review these considerations before you perform a Windows BMR recovery on a clustered host.

- Before you start the Windows BMR recovery operation, ensure that you detach the shared disks. After the Windows BMR recovery operation and the restart completes, attach the shared disks before you perform the online recovery.
- After an authoritative restore completes, the recovery operation does not bring the cluster services online on the remote nodes. You must bring the services online manually.

Windows Storage Pools considerations

When a system failure occurs which damages Storage Pools, perform the following steps as recommended by Microsoft to perform a BMR recovery to a new host. In the case of a complete system failure, a Storage Pool may not exist on the target host. There can only be physical disks. Some of these disks are required to create Storage Pools.

Before beginning Windows BMR wizard, physically remove from the target recovery computer any physical disks reserved for storage pools. This manual step is required because the Windows BMR wizard does not have any option to exclude the disks.

To recover Storage Spaces to a new host, perform the following steps:

1. Boot the host with the Windows BMR image.
2. Recover only the critical volumes.
3. Reboot the host.
4. Attach physical disks that are reserved for Storage Pools.
5. Use Windows Server Manager or Powershell Cmdlets to configure the Storage Pools.
6. Perform a volume or file recovery of the Storage Spaces volumes.
7. Perform a volume or file recovery of other volumes on physical disks.

WinPE considerations for SAN boot devices

When you recover to a host that uses a SAN boot device, the WinPE environment requires that you temporarily disable all but one path to the boot device. After the BMR recovery and reboot completes you can re-enable the remaining paths.

VMware network interface card driver limitations

The Windows BMR image does not contain a driver for any of the VMware VMXNET, VMXNET3, or the VMware Paravirtual SCSI NIC models. The Windows BMR image does contain a driver for the e1000 NIC. When you perform a Windows BMR recovery, ensure that the VM has at least one configured e1000 NIC, or add custom NIC drivers when you run the NetWorker BMR wizard.

The VMware Tools installation media in the `\Program Files\VMware\VMware Tools\Drivers` folder on the system drive of the VM contains the VMware NIC drivers.

BCD partition limitations

NetWorker requires that the BCD partitions are online during a Windows BMR backup. If a BCD partition is offline during a Windows BMR backup, the backup fails with an messages similar to the following:

```
save: Unable to get volume information of file system.The device is
not ready. (Win32 error 0x15) with the volume offline
```

Creating a Client resource with the Client Backup Configuration wizard

The Client Backup Configuration wizard enables you to quickly configure a client resource with a limited set of key backup options. Follow these steps to configure a file system backup and a BMR backup for a Windows host.

Before you begin

- Install the NetWorker client software on the client computer.
- Ensure that the NetWorker server host is listed in the `servers` file on the client computer.
- Ensure that the communication between the NMC server, NetWorker client, and NetWorker server uses `nsrauth` strong authentication.
- Ensure that the user who runs the wizard meets the following requirements:
 - Root (UNIX) or Administrator (Windows) privileges.
 - A member of a User Group on the NetWorker server that has Configure NetWorker privileges.
- Ensure that multiple wizard hosts are not trying to access the same client computer simultaneously.
- (Optional) Check for reverse entries in the DNS server. If the reverse entries are not present, then do not use the IP address for creating the clients.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left panel, right-click **Clients**, and then select **New Client Wizard**.

The **Client Backup Configuration** wizard appears.

3. In the **Client Name** box, type either the hostname or the Fully Qualified Domain Name (FQDN) of the client.

It is recommended that you specify the FQDN of the host. For OS cluster hosts, type the FQDN of the virtual host.

For application cluster hosts, type the FQDN of the application cluster host. For example:

- For an Oracle cluster, type the RAC hostname.
- For an Exchange IP DAG, type the DAG name.

The application module administrator guides provide more information.

 **Note:** If the Client Configuration wizard cannot resolve the specified hostname, an error message appears after you click **Next**.

4. Optionally, in the **Comment** box, type a description of the client.
If you are creating multiple client resources for the same NetWorker client host, then use this attribute to differentiate the purpose of each resource.
5. In the **Tag** box, type one or more tags to identify this Client resource for the creation of dynamic client groups for data protection policies.

Place each entry on a separate line.

6. In the **Type** box, select **Traditional NetWorker client**.
7. Optionally, from the **Group** list, select a group for the Client resource.

The group to which the client belongs determines the workflow that is used to back up the client.

 **Note:** You can also assign the client to one or more groups after you create the Client resource.

8. Click **Next**.
9. On the **Specify the Backup Configuration Type** window, select **Filesystem**, and then click **Next**.
10. On the **Select the NetWorker Client Properties** window, configure the following options:

Option	Description
Priority	Enables you to control the order in which the NetWorker server contacts clients for backup. During a backup operation, the NetWorker server contacts the client with the lowest priority value first. If you do not specify a priority for the client resources, then the backup order is random. The default value is 500. While the Priority attribute specifies the order of client contact, many variables affect the order in which clients complete their backups. For example: • The backup operation on a client does not begin until the worklists for each of the save sets on the client are complete. • The amount of work can vary greatly from one client to the next. • If a client stops responding and times out, then the backup operation puts the client backup at the end of the backup order list. The only way to guarantee that the backup of one client occurs before the backup of another client is to configure the workflows for the clients to start at different times.
Parallelism	Specifies the maximum number of data streams that a client can send simultaneously during a backup action. Data streams include backup data streams, savefs processes, and probe jobs. The default value is different for the NetWorker server than it is for all other client resources: • For the NetWorker server client resource, the default value is 12. This higher default value enables the server to complete a larger number of index backups during a Server backup action. • For all other clients, the default value is 4. To avoid disk contention for clients other than the NetWorker server, specify a value that is the same as or fewer than the

Option	Description
	number of physical disks on the client that are included in the backup. The <i>NetWorker Performance Optimization Planning Guide</i> provides more information about recommended client parallelism values and performance benefits.
Remote Access	Specifies a list of the users that have access to perform remote access operations. For example, users that can perform a directed recovery of backup data that originated on this host.
Data Domain Interface	Specifies the protocol to use if you send the backup data to a Data Domain Device. Available selections are IP, Fibre Channel, or Both.  Note: Mac OS X clients only support the IP protocol.
Block Based Backup (BBB)	Enables Block Based Backups for the host. When you select this option, you must also select the Client Direct. This option applies to Linux only.  Note: The Block Based Backup chapter provides complete information about how to configure a host for BBB backups.
Client Direct	Allows the client to try to directly connect to the backup storage device, instead of connecting to a NetWorker storage node. If a direct connection is not possible, then the backup operation connects to the NetWorker storage node that you configure to accept data from the client.
Parallel Save Streams (PSS)	Enables NetWorker to use multiple parallel save streams to backup each save set defined for the client, to one or more destination devices. PSS does not support Checkpoint Restart backups.

11. Click **Next**.
12. On the **Select the File System Objects** window, select the file system objects to backup.

 **Note:** To avoid the over consumption of memory, NetWorker limits the number of files that you can view when you browse a directory that contain a large number of files, for example, 200,000 files. When NetWorker determines that displaying the number of files will exhaust memory resources, NetWorker will display a partial list of the files and a message similar to the following appears: Expanding this directory has stopped because the result has too many entries

CIFS, DFS, and msdos file systems do not appear as selectable file system objects. [Modifying the save_sets defined for a Windows client](#) describes how to modify the save set attribute to define backup a remote file system.

 **Note:** When you select all file system objects and the `DISASTER_RECOVERY: \` save set, the **ALL** value appears in the **Save set** attribute for the client resource. When you select file system objects, enables you to perform granular recoveries of files and directories. The

`DISASTER_RECOVERY:\ save set` enables you to perform a BMR restore of the Windows host. To backup Active Directory, DFSR, or Cluster Services, ensure that you perform `DISASTER_RECOVERY:\ backup`.

13. On the **Backup Configuration Summary** window, click **Create**.
14. On the **Client Configuration Results** window, review the results of the client configuration process, then click **Finish**.

Results

The Client resource appears in the **Clients** window pane.

Verifying a valid Windows BMR backup

After you perform a Windows BMR backup, verify that the backup exists. NetWorker creates one save set for each critical volume backed up by the `DISASTER_RECOVERY:\ save set`.

About this task

You can verify that the backup exists by using the NMC console, the NetWorker User program, or the `nsrinfo` program.

NOTICE If any of the components of the Windows BMR backup fail, then NetWorker does not create a `DISASTER_RECOVERY:\ save set` and you cannot perform an offline recovery. The backup process may backup the `WINDOWS ROLES AND FEATURES` save sets or critical volumes, which NetWorker makes available for an online recovery.

Verifying that a valid backup exists by using the NMC console

Procedure

1. Use NMC to connect to the NetWorker server.
2. In the **NetWorker Administration** window, click **Media**.
3. In the left pane, click **Save Sets**.
4. On the right pane, on the **Query Save Set** tab, specify the search criteria such as the NetWorker **Client Name** and a date range for the **Save Time**.
5. Select the **Save Set List** tab in the right pane to generate and display a list of save sets that meet the search criteria.

Verifying that a valid `DISASTER_RECOVERY:\ save set` exists by using the NetWorker User Program

By default, the **Recovery** window displays the most recent backup. To verify an older backup select the **View > Change Browse Time** menu option, and then specify a different backup date and time.

Procedure

1. Start the NetWorker **User** program by using the `winworkr` command with the `-s` option to connect to the NetWorker server to which the source client data is backed up:

```
winworkr -s server_name
```

If the `-s` option is not entered and there is only one server detected, that server is connected automatically. If there are no servers detected, or if there is more than one server available, the **Change Server** dialog box appears, enabling you to choose the server.

2. Click **Recover**.

The **Source Client** dialog box appears.

3. Select the source client of the `DISASTER_RECOVERY:\` save set, and then click **OK**.
4. Select a destination client, and then click **OK**.
5. In the **Recover** window, browse and locate the save set named `DISASTER_RECOVERY:\`.

Verifying that a valid `DISASTER_RECOVERY:\` save set exists by using the `nsrinfo` program

To query the client file index of the Windows host and display information about the `DISASTER_RECOVERY:\` save set, type the following command from a command prompt.

About this task

```
nsrinfo -v -s server_name -N "DISASTER_RECOVERY:\\\" client_name
```

where:

- *server_name* is the name of the NetWorker server.
- *client_name* is the name of the client that performed the Windows BMR backup.

Performing a NetWorker Bare Metal Recovery wizard test

Before you need to perform a Windows BMR, test the wizard to ensure that you can complete a recovery and that you have the required drivers. This task is especially important for 64-bit hosts that might require additional drivers. For both 64-bit and 32-bit hosts, the wizard must use drivers that do not require a reboot.

About this task

 **NOTICE** After you test the wizard, you can safely exit the wizard before completing the entire recovery process.

Procedure

1. Follow the procedures in [Performing a Windows BMR to physical or virtual computers](#) on page 622.

Verify the following as you step through the BMR recovery wizard screens:

- If DNS is not available, that the host can resolve the NetWorker server name by some method, such as a local hosts file.
- You can see the network interface that is required to communicate with the NetWorker server. If you cannot see the network interface, use the wizard to load the required NIC driver.
- You can see the critical and non-critical disks for the host that is to be recovered. If you cannot see all of the disks, use the wizard to load the required disk drivers.

2. Click **Exit** to safely exit the wizard.
3. Exit the command window.

The system automatically reboots.

Modifying the save sets defined for a Windows client

You can modify an existing client to change the file system objects to backup on the client.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left panel, select **Clients**.
3. Right-click the client resource and select **Modify Client Properties**.

The **Client Properties** dialog box appears.

4. On the **General** tab, in the **Save set** attribute, specify the file system, directory or path to a file. Specify one file system object on each line. You can also modify specify a special **ALL** save set to backup a specific type of file system only. The following table summarizes the available **ALL** save sets.

Table 74 Special ALL save sets

Special ALL save set syntax	Backup behavior
<code>all-file_system</code>	Only back up locally mounted file systems of a particular type, where <i>file_system</i> is the name of the file system, for example <code>ntfs</code> . The <i>NetWorker E-LAB Navigator</i> provides a list of the supported file system for each operating system.
<code>all-mounts</code>	On Windows clients, the <code>all-mounts</code> save set is equivalent to the ALL save set. File systems that are normally skipped are still skipped.

Mapped drives

To back up mapped or CIFS drives on a Windows client for either a scheduled or a manual backup, you must perform additional configuration steps in the Client resource.

Before you begin

- Create a dedicated client resource for the backups of mapped drives. A common user account must have access to each mapped drive.
- Create a separate Client resource for backups of local drives.
- Ensure that the **Administration** window is in Diagnostic Mode. To enable Diagnostic Mode, from the **View** menu, select **Diagnostic Mode**.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left panel, select **Clients**.
3. Right-click the client resource and select **Modify Client Properties**.

The **Client Properties** dialog box appears.

4. On the **General** tab, in the **Save set** attribute, specify the Universal Naming Convention (UNC) path of the drive.

Do not specify the drive letter. For example, to specify the `accounts` directory on the `jupiter` server, type `\\jupiter\accounts`.

5. On the **Apps & Modules** tab, configure the following attributes:
 - a. In the **Remote user** and **Password** fields, specify a username and the associated password for an account that has access to the UNC path.
 - b. In the **Backup command** box, type `save -xL`.
 - c. In the **Save operations** box, type `vss:*=off`

Configuring a Client resource for backups on UNIX hosts

This section describes how to configure a Client resource to backup data on UNIX hosts.

UNIX/Linux backup considerations

The following topics provide details on considerations for backing up client data on Solaris, Linux, HP-UX, and AIX computers.

Linux

You can install the NetWorker client, server, storage node, and NetWorker Management Console (NMC) server software on Linux.

Backup and recovery operations are supported on the following Linux journaled file systems:

- `ext3`
- `reiserfs`
- `jfs`
- `xfs`

For `ext3` file systems with the journal set to `visible`, do not back up or recover the journal. Recovering the journal may cause the file system to become unstable. Use a directive to ensure that the file system is excluded from a backup. [Directives](#) on page 396 provides information on directives.

Solaris

You can install NetWorker client, server, storage node, and NetWorker Management Console (NMC) server software on the Solaris platform.

The NetWorker software supports local and global zones for a NetWorker client, server, and a dedicated storage node. You can install and back up a NetWorker client, server, or storage node on a computer running in a local zone. The NMC and NetWorker License Manager can only be installed in a global zone.

 **Note:** Extended file attribute data is in the calculation of the save set file size for Solaris clients. As a result, the save set file size in NetWorker appears to slightly larger than expected.

NetWorker executables not found for Solaris client

On Solaris client computers, NetWorker executables are installed by default in `/usr/sbin`. The search path for root on the NetWorker server must include `/usr/sbin`. Otherwise, scheduled backups fail on a client with NetWorker executables in `/usr/sbin` because the `savefs` command is not in the search path.

To solve this issue, edit the search path for root on the NetWorker server to include `/usr/sbin`, even if the directory does not exist locally.

Alternatively, specify `/usr/sbin` in the **Executable path** attribute on the **Globals (2 of 2)** tab of the **Client Properties** dialog box for the Client resource.

HP-UX

You can install NetWorker client, server, storage node, and NetWorker Management Console (NMC) server software on the HP-UX platform.

Customized backup scripts

On HP-UX, do not use the `posix` shell (`/bin/sh`) for customized backup scripts that are meant to be automatically started by the backup. Use the `korn` shell (`/bin/ksh`) instead.

Symbolic link entries in the `fstab` file

For HP-UX operating systems, do not use symbolic link entries in the `/etc/fstab` file. If you use symbolic links in the `fstab` file, the backup does not include the file system to which the symbolic link points.

AIX

You can install the NetWorker client, server, storage node, and NetWorker Management Console (NMC) server software on the AIX platform.

 **Note:** On AIX, non-root users who are performing a recovery cannot restore group ownership (the `set-group-id-on-execution` or `setuid` permission bit) on binaries or files. This behavior is to be expected.

Creating a Client resource with the Client Backup Configuration wizard

The **Client Backup Configuration** wizard enables you to quickly configure a Client resource with a limited set of key backup options. Follow these steps to configure a file system backup and a UNIX host.

Before you begin

- Install the NetWorker client software on the client computer.
- Ensure that the NetWorker server host is listed in the `servers` file on the client computer.
- Ensure that the communication between the NMC server, NetWorker client, and NetWorker server uses `nsrauth` strong authentication.
- Ensure that the user who runs the wizard meets the following requirements:
 - Root (UNIX) or Administrator (Windows) privileges.
 - A member of a User Group on the NetWorker server that has Configure NetWorker privileges.
- Ensure that multiple wizard hosts are not trying to access the same client computer simultaneously.
- (Optional) Check for reverse entries in the DNS server. If the reverse entries are not present, then do not use the IP address for creating the clients.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left panel, right-click **Clients**, and then select **New Client Wizard**.

The **Client Backup Configuration** wizard appears.

3. In the **Client Name** box, type either the hostname or the Fully Qualified Domain Name (FQDN) of the client.

It is recommended that you specify the FQDN of the host. For OS cluster hosts, type the FQDN of the virtual host.

For application cluster hosts, type the FQDN of the application cluster host. For example:

- For an Oracle cluster, type the RAC hostname.
- For an Exchange IP DAG, type the DAG name.

The application module administrator guides provide more information.

 **Note:** If the Client Configuration wizard cannot resolve the specified hostname, an error message appears after you click **Next**.

4. Optionally, in the **Comment** box, type a description of the client.

If you are creating multiple client resources for the same NetWorker client host, then use this attribute to differentiate the purpose of each resource.

5. In the **Tag** box, type one or more tags to identify this Client resource for the creation of dynamic client groups for data protection policies.

Place each entry on a separate line.

6. In the **Type** box, select **Traditional NetWorker client**.
7. Optionally, from the **Group** list, select a group for the Client resource.

The group to which the client belongs determines the workflow that is used to back up the client.

 **Note:** You can also assign the client to one or more groups after you create the Client resource.

8. Click **Next**.
9. On the **Specify the Backup Configuration Type** window, select **Filesystem**, and then click **Next**.
10. On the **Select the NetWorker Client Properties** window, configure the following options:

Option	Description
Priority	Enables you to control the order in which the NetWorker server contacts clients for backup. During a backup operation, the NetWorker server contacts the client with the lowest priority value first. If you do not specify a priority for the client resources, then the backup order is random. The default value is 500. While the Priority attribute specifies the order of client contact, many variables affect the order in which clients complete their backups. For example: • The backup operation on a client does not begin until the worklists for each of the save sets on the client are complete.

Option	Description
	- The amount of work can vary greatly from one client to the next. - If a client stops responding and times out, then the backup operation puts the client backup at the end of the backup order list. The only way to guarantee that the backup of one client occurs before the backup of another client is to configure the workflows for the clients to start at different times.
Parallelism	Specifies the maximum number of data streams that a client can send simultaneously during a backup action. Data streams include backup data streams, savefs processes, and probe jobs. The default value is different for the NetWorker server than it is for all other client resources: - For the NetWorker server client resource, the default value is 12. This higher default value enables the server to complete a larger number of index backups during a Server backup action. - For all other clients, the default value is 4. To avoid disk contention for clients other than the NetWorker server, specify a value that is the same as or fewer than the number of physical disks on the client that are included in the backup. The <i>NetWorker Performance Optimization Planning Guide</i> provides more information about recommended client parallelism values and performance benefits.
Remote Access	Specifies a list of the users that have access to perform remote access operations. For example, users that can perform a directed recovery of backup data that originated on this host.
Data Domain Interface	Specifies the protocol to use if you send the backup data to a Data Domain Device. Available selections are IP, Fibre Channel, or Both.  Note: Mac OS X clients only support the IP protocol.
Block Based Backup (BBB)	Enables Block Based Backups for the host. When you select this option, you must also select the Client Direct. This option applies to Linux only.  Note: The Block Based Backup chapter provides complete information about how to configure a host for BBB backups.
Client Direct	Allows the client to try to directly connect to the backup storage device, instead of connecting to a NetWorker storage node. If a direct connection is not possible, then the backup operation connects to the NetWorker storage node that you configure to accept data from the client.

Option	Description
Parallel Save Streams (PSS)	Enables NetWorker to use multiple parallel save streams to backup each save set defined for the client, to one or more destination devices. PSS does not support Checkpoint Restart backups.

11. Click **Next**.
12. On the **Select the File System Objects** window, select the file system objects to backup.

To avoid the over consumption of memory, NetWorker limits the number of files that you can view when you browse a directory that contain a large number of files, for example, 200,000 files. When NetWorker determines that displaying the number of files will exhaust memory resources, NetWorker will display a partial list of the files and a message similar to the following appears:

```
Expanding this directory has stopped because the result
has too many entries
```

Note: When you select all file system objects, the ALL value appears in the **Save set** attribute for the client resource. When the backup starts, the `savefs` process reads the contents of the `/etc/vfstab` file on Solaris clients, the `/etc/fstab` file on HP-UX and Linux clients, or the `/etc/filesystems` file on AIX clients. The contents of the file are compared to the currently mounted file systems and BTRFS sub-volumes. Only currently mounted file systems and BTRFS sub-volumes that are configured in these files are backed up. When NetWorker encounters a sub-directory that has a sub-volume ID that differs from the parent sub-volume ID, NetWorker will not backup the contents of the subdirectory, unless you specify the `save -x` in the Backup command field in the properties of the Client resource. After you create the client configuration wizard, you can modify the client resource or create a new client resource to include the excluded file systems. [Supported save set configurations for UNIX hosts](#) provides more information.

When you specify the ALL save set:

- For a Solaris sparse or whole root zone client, all mounted file systems in the sparse or whole root zone that are not normally skipped, such as NFS, are backed up.
- ZFS file systems are backed up.
- If the save set name includes a symbolic link, a save set recovery is not supported.

13. On the **Backup Configuration Summary** window, click **Create**.
14. On the **Client Configuration Results** window, review the results of the client configuration process, then click **Finish**.

Results

The Client resource appears in the **Clients** window pane.

Supported save set configurations for UNIX hosts

The **Client Configuration** wizard does not display some types of file systems on UNIX hosts and these save sets are not in the **ALL** save set.

When the backup starts, the `savefs` process reads the contents of the `/etc/vfstab` file on Solaris clients, the `/etc/fstab` file on HP-UX and Linux clients, or the `/etc/filesystems` file on AIX clients. The contents of the file are compared to the currently mounted file systems and BTRFS sub-volumes. Only currently mounted file systems and BTRFS sub-volumes that are configured in these files are backed up. When NetWorker encounters a sub-directory that has a sub-volume ID that differs from the parent sub-volume ID, NetWorker will not backup the contents of the subdirectory, unless you specify the `save -x` in the Backup command field in the properties of the Client resource.

If you edit a client resource and modify the **Save set** attribute to include file system objects for file systems that are not in the OS file system file, NetWorker will not back up the file system objects.

The following file systems are excluded from the **ALL** save set. If you manually define the file system or directories and files for one of these file systems in the **Save set** attribute of the Client resource, the backup operation excludes the object:

Table 75 File systems excluded from the **ALL** save set

• hsfs	• sharefs	• dfs	• binfmt_mi	• nucam
• proc	• nfs2	• autofs	sc	• fdfs
• fd	• nfs3	• iso9060	• usbfs	• xx
• cachefs	• nfs3perf	• udf	• devpts	• none
• lofs	• profs	• sysfs	• smbfs	
• mntfs	• nfs4	• debugfs	• swap	
• ctfs	• nfs	• subfs	• tmp	
• objfs	• brfs	• usbdevfs	• tmpfs	
			• nucfs	

When you specify the **ALL** save set:

- For a Solaris sparse or whole root zone client, all mounted file systems in the sparse or whole root zone that are not normally skipped, such as NFS, are backed up.
- ZFS file systems are backed up.
- If the save set name includes a symbolic link, a save set recovery is not supported.

Use a customized **ALL** save set to backup files.

Modifying the save sets defined for a UNIX client

You can modify a client to change the file system objects to backup on the client.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left panel, select **Clients**.
3. Right-click the client resource and select **Modify Client Properties**.

The **Client Properties** dialog box appears.

4. On the **General** tab, in the **Save set** attribute, specify the file system, directory or path to a file. Specify one file system object on each line. You can also modify specify a special ALL save set to backup a specific type of file system only. The following table summarizes the available ALL save sets.

Table 76 Special ALL save sets

Special ALL save set syntax	Backup behavior
<code>all-file_system</code>	Only back up locally mounted file systems of a particular type, where <i>file_system</i> is the name of the file system. For example, the <code>all-zfs</code> save set backs up all locally mounted <code>zfs</code> file systems on a Solaris host. File systems such as NFS that are normally skipped are still skipped. When the backup starts, the <code>savefs</code> process reads the contents of the <code>/etc/vfstab</code> file on Solaris clients, the <code>/etc/fstab</code> file on HP-UX and Linux clients, or the <code>/etc/filesystems</code> file on AIX clients. The contents of the file are compared to the currently mounted file systems and BTRFS sub-volumes. Only currently mounted file systems and BTRFS sub-volumes that are configured in these files are backed up. When NetWorker encounters a sub-directory that has a sub-volume ID that differs from the parent sub-volume ID, NetWorker will not backup the contents of the subdirectory, unless you specify the <code>save -x</code> in the Backup command field in the properties of the Client resource. The <i>NetWorker E-LAB Navigator</i> provides a list of the supported file system for each operating system.
<code>all-mounts</code>	Back up all the currently mounted file systems. File systems such as NFS that are normally skipped are still skipped.
<code>all-local</code>	For a global zone client, the file systems in the sparse

Table 76 Special ALL save sets (continued)

Special ALL save set syntax	Backup behavior
	or whole root zone on the physical host are backed up. File systems in the global zone are skipped. For a sparse or whole root zone client, the <code>all-local</code> save set is equivalent to the <code>ALL</code> save set.
<code>all-global</code>	For a global zone client, all file systems in the global zone are backed up. All sparse and whole root zone file systems on the physical host are skipped. For a Solaris sparse or whole root zone client, the <code>all-global</code> save set is equivalent to the <code>ALL</code> save set.

Note: If you explicitly list a BTRFS sub-volume in the **Save set** field, NetWorker will back up the files in the sub-volume, even if the sub-volume does not appear in the `/etc/fstab` file. When NetWorker encounters a sub-directory that has a sub-volume ID that differs from the parent sub-volume ID, NetWorker will not backup the contents of the subdirectory, unless you specify the `save -x` in the Backup command. To back up data in the subdirectories, perform one of the following tasks:

- Specify `save -x` in the **Backup command** field in the client properties window.
- Explicitly list the path of each sub-volume in the **Save set** field.
- Mount each sub-volume, include the mount point in the `/etc/fstab` file, and then specify `ALL` or `all-btrfs` in the **Save set** field.

5. Click **OK**.

Configuring a Client resource for backups on Mac OS X hosts

This section describes how to configure a Client resource to backup data on Mac OS X hosts.

Mac OS X backup considerations

You can configure a Mac OS X host as a NetWorker client. You can use any supported NetWorker server on UNIX, Linux, or Windows to back up and restore an OS X host. You cannot configure an OS X host as a NetWorker server or an NMC server.

The NetWorker client for OS X supports the following file systems:

- HFS+ (including journaled)

- HFS
- UFS

The NetWorker client for OS X also backs up and recovers all file system metadata, including:

- Finder information
- Resource forks
- Extended attributes
- Access Control Lists (ACLs)

Creating a Client resource with the Client Backup Configuration wizard

The **Client Backup Configuration** wizard enables you to quickly configure a Client resource with a limited set of key backup options. Follow these steps to configure a file system backup and an OS-X host.

Before you begin

- Install the NetWorker client software on the client computer.
- Ensure that the NetWorker server host is listed in the `servers` file on the client computer.
- Ensure that the communication between the NMC server, NetWorker client, and NetWorker server uses `nsrauth` strong authentication.
- Ensure that the user who runs the wizard meets the following requirements:
 - Root (UNIX) or Administrator (Windows) privileges.
 - A member of a User Group on the NetWorker server that has Configure NetWorker privileges.
- Ensure that multiple wizard hosts are not trying to access the same client computer simultaneously.
- (Optional) Check for reverse entries in the DNS server. If the reverse entries are not present, then do not use the IP address for creating the clients.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left panel, right-click **Clients**, and then select **New Client Wizard**.

The **Client Backup Configuration** wizard appears.

3. In the **Client Name** box, type either the hostname or the Fully Qualified Domain Name (FQDN) of the client.

It is recommended that you specify the FQDN of the host. For OS cluster hosts, type the FDQN of the virtual host.

For application cluster hosts, type the FQDN of the application cluster host. For example:

- For an Oracle cluster, type the RAC hostname.
- For an Exchange IP DAG, type the DAG name.

The application module administrator guides provide more information.

 **Note:** If the Client Configuration wizard cannot resolve the specified hostname, an error message appears after you click **Next**.

4. Optionally, in the **Comment** box, type a description of the client.

If you are creating multiple client resources for the same NetWorker client host, then use this attribute to differentiate the purpose of each resource.

5. In the **Tag** box, type one or more tags to identify this Client resource for the creation of dynamic client groups for data protection policies.

Place each entry on a separate line.

6. In the **Type** box, select **Traditional NetWorker client**.

7. Optionally, from the **Group** list, select a group for the Client resource.

The group to which the client belongs determines the workflow that is used to back up the client.

 **Note:** You can also assign the client to one or more groups after you create the Client resource.

8. Click **Next**.

9. On the **Specify the Backup Configuration Type** window, select **Filesystem**, and then click **Next**.

10. On the **Select the NetWorker Client Properties** window, configure the following options:

Option	Description
Priority	Enables you to control the order in which the NetWorker server contacts clients for backup. During a backup operation, the NetWorker server contacts the client with the lowest priority value first. If you do not specify a priority for the client resources, then the backup order is random. The default value is 500. While the Priority attribute specifies the order of client contact, many variables affect the order in which clients complete their backups. For example: • The backup operation on a client does not begin until the worklists for each of the save sets on the client are complete. • The amount of work can vary greatly from one client to the next. • If a client stops responding and times out, then the backup operation puts the client backup at the end of the backup order list. The only way to guarantee that the backup of one client occurs before the backup of another client is to configure the workflows for the clients to start at different times.
Parallelism	Specifies the maximum number of data streams that a client can send simultaneously during a backup action. Data streams include backup data streams, savefs processes, and probe jobs.

Option	Description
	The default value is different for the NetWorker server than it is for all other client resources: - For the NetWorker server client resource, the default value is 12. This higher default value enables the server to complete a larger number of index backups during a Server backup action. - For all other clients, the default value is 4. To avoid disk contention for clients other than the NetWorker server, specify a value that is the same as or fewer than the number of physical disks on the client that are included in the backup. The <i>NetWorker Performance Optimization Planning Guide</i> provides more information about recommended client parallelism values and performance benefits.
Remote Access	Specifies a list of the users that have access to perform remote access operations. For example, users that can perform a directed recovery of backup data that originated on this host.
Data Domain Interface	Specifies the protocol to use if you send the backup data to a Data Domain Device. Available selections are IP, Fibre Channel, or Both.  Note: Mac OS X clients only support the IP protocol.
Block Based Backup (BBB)	Enables Block Based Backups for the host. When you select this option, you must also select the Client Direct. This option applies to Linux only.  Note: The Block Based Backup chapter provides complete information about how to configure a host for BBB backups.
Client Direct	Allows the client to try to directly connect to the backup storage device, instead of connecting to a NetWorker storage node. If a direct connection is not possible, then the backup operation connects to the NetWorker storage node that you configure to accept data from the client.
Parallel Save Streams (PSS)	Enables NetWorker to use multiple parallel save streams to backup each save set defined for the client, to one or more destination devices. PSS does not support Checkpoint Restart backups.

11. Click **Next**.
12. On the **Select the File System Objects** window, select the file system objects to backup.

To avoid the over consumption of memory, NetWorker limits the number of files that you can view when you browse a directory that contain a large number of files, for example, 200,000 files. When NetWorker determines that displaying the number of files will exhaust memory resources, NetWorker will display a partial list of the files and a message similar to the following appears:

Expanding this directory has stopped because the result has too many entries

 **Note:** When you select all file system objects, the `ALL` value appears in the **Save set** attribute for the Client resource. The `ALL` save set includes local and mounted volumes.

13. On the **Backup Configuration Summary** window, click **Create**.
14. On the **Client Configuration Results** window, review the results of the client configuration process, then click **Finish**.

Results

The Client resource appears in the **Clients** window pane.

Assigning directives to Mac OS X clients

After you create a client resource for an OS X client, select one of the Mac OS directives to exclude certain files and directories from the backup, and ensure a consistent state after a recovery operation.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left panel, select **Clients**.
3. Right-click the client resource and select **Modify Client Properties**.
The **Client Properties** dialog box appears.
4. On the **General** tab, in the **Directive** box, select one of the following directives:
 - Mac OS Standard Directives
 - Mac OS with Compression Directives

[Preconfigured global Directive resources](#) on page 401 provides more information about the Mac OS directives.
5. Click **OK**.

Configuring Open Directory database backups

The Mac OS directive does not back up Open Directory database files, which contain system configuration information that is essential for disaster recovery. To ensure complete protection of a Mac OS X computer if a catastrophic failure occurs, create a script file and then modify the client resource for the Mac OS X host to include the Open Directory database files.

About this task

[Customizing backups with the pre and post commands](#) on page 483 provides more information about the script file and the how to modify the client resource to use the command.

Procedure

1. On the OS X host, create the script file as an executable text file.
The name of the script file must start with `nsr` or `save`. For example `nsr_opendir_backup.sh`
2. Add the commands to backup open files to the script file.

 **Note:** Open Directory database files remain available during the backup.

- To back up LDAP directory domain for the Open Directory, type:

```
#slapcat -l /var/backups/networker.ldif
```

- To back up Password Server database for the Open Directory when the OS-X host uses LDAP over SSL, type:

```
# mkdir -p /var/backups/networker.odpdb
# mkpassdb -backupdb /var/backups/networker.odpdb
```

- To back up the local NetInfo directory domain, type:

```
# nidump -r / . > /var/backups/networker.nidump
```

The following script file provides an example of how to back up the LDAP directory, Password Server, and NetInfo databases before each scheduled save:

```
"/usr/sbin/slapcat -l /var/backups/networker.ldif;
/bin/mkdir -p /var/backups/networker.odpdb;
/usr/sbin/mkpassdb -backupdb /var/backups/networker.odpdb;
/usr/bin/nidump -r / . > /var/backups/networker.nidump"
```

3. Connect to the NetWorker server by using NMC.
4. In the **NetWorker Administration** window, click **Protection**.
5. In the expanded left panel, select **Clients**.
6. Right-click the client resource and select **Modify Client Properties**.

The **Client Properties** dialog box appears.

7. On the **Apps & Modules** tab, in the **Pre command** attribute, specify the name of the script file that you require NetWorker to run before a backup.

 **Note:** Do not specify the path to the file.

8. Click **OK**.

Sending client data to AFTD or Data Domain devices only

Use the **Backup target disks** attribute of the client resource to define an ordered list of AFTD and Data Domain disk devices that will receive data for this client. When you specify a value in this attribute, NetWorker ignores the values that you specify in the **Storage nodes** attribute. This attribute does not apply to the client resource of the NetWorker server, and applies to each instance of the client resource. You can specify devices that are local or remote to the NetWorker server.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left panel, select **Clients**.
3. Right-click the client resource and select **Modify Client Properties**.

The **Client Properties** dialog box appears.

4. On the **Globals (2 of 2)** tab, in the **Backup target disks** attribute, specify the name of the AFTD or Data Domain devices that NetWorker uses to store data for this client.

Specify each device name on a separate line.

5. Click **OK**.

Results

NetWorker does not use the values in the **Storage nodes** attribute of the client resource when selecting the device to receive data for the client.

Non-ASCII files and directories

If you create a client resource by using the **Client Properties** dialog box and the **Save set** field contains non-ASCII characters, you must edit the **Save operations** field on the **Apps & Modules** tab for the client resource.

To access the **Save operations** field, in the **NetWorker Administration** window, click **View > Diagnostic Mode**.

In the **Client Properties** dialog box, on the **Apps & Modules** tab, in the **Save operations** field, specify `I18N:mode=utf8path`

Configuring checkpoint restart backups

The checkpoint restart feature allows a failed backup operation to restart at a known good point, before the point of failure during the backup.

Note: Checkpoint restart is only supported on Linux and UNIX environments when performing standard save operations; you cannot use checkpoint restart with block-based backup or parallel save streams enabled. Checkpoint restart is not supported on Windows platforms.

A known good point is defined as a point in the backup data stream where the data is successfully written to the save set and that data can be located and accessed by subsequent recovery operations. This feature allows client backups that are part of a scheduled backup to be restarted, if they fail while running. This prevents the files and directories that have already been backed up from being backed up again.

Backup failures occur for various reasons. The most common reasons include hardware failures, loss of network connectivity, and primary storage software failures. If a backup fails and checkpoint restart is enabled, then failed save sets are marked as partial instead of as aborted. Partial save sets remain in the index, the media databases, and media such as AFTD.

You can manually restart a failed backup, or you can configure the backup to restart automatically. A restarted save set has a new SSID and savetime.

The NetWorker server and storage node components must remain running to manage the client failure and to create a partial save set. If the NetWorker server or storage node components fail during a backup, then partial save sets are not created. In this case, the backup for the checkpoint-enabled client starts from the beginning.

If the checkpoint restart feature is not enabled, a failure that is encountered during a scheduled backup operation might require a rerun of an entire backup tape set. This can be costly when a limited backup window of time is available, as a significant portion of the backup data might have been successfully transferred to tape, and the NetWorker software cannot resume a save set from the point of interruption.

For example, when performing an 800 GB backup that requires approximately 10 hours to complete and spans six tapes, if a failure occurs while writing to the last tape, the

previous five tapes representing 9 hours of backup time may need to be rerun. As datasets continue to increase in size, so does the impact of backup failures.

About partial save sets

The backup sequence of partial save sets is not the same as the backup sequence for complete backups. Each partial save set provides protection for part of the file system, but the completeness and consistency of the coverage of the whole file system cannot be guaranteed.

The checkpoint restart window is user-defined and can be large. If restarted hours apart, the partial backups might provide an image of the file system that is different from the state of the file system at any fixed point in time. The resulting file system backup is not guaranteed to be consistent.

NetWorker performs file and directory backups in alphabetical order. If a failure occurs, and you restart the backup, the backup operation starts alphabetically with the next file or folder that was not previously backed up. NetWorker does not review files or folder that were previously backed up for changes. If a previously backed up file or folder was edited or added after the backup failure, NetWorker does not back up the file or directory again.

Consider the following example in which a backup is interrupted while it is saving a directory and is restarted after the directory contents have changed:

1. A save set contains `/disk1/dir` with files `file_a`, `file_c` and `file_d`.
2. The backup of the save set is interrupted while `file_d` is backed up.

As a result, the first partial save set includes only `file_a` and `file_c`.

3. A user adds `file_b` to the file system.
4. The checkpoint restart is initiated for the save set.

The second partial save set contains `file_d` and `/disk1/dir`, which includes `file_a`, `file_b`, `file_c` and `file_d`. However, `file_b` is not in the save set.

Partial saveset cloning and scanning

Partial save sets can be cloned and scanned individually. These operations must be performed on every partial save set.

If legacy automatic cloning is enabled, all partial save sets are cloned because automatic cloning is run as part of the scheduled backup.

Checkpoint restart requirements

Ensure that the environment meets the following requirements to support checkpoint restart.

Server and client software requirements

Checkpoint restart requires the server and client software listed in the following table.

Table 77 NetWorker software requirements for checkpoint restart

Client	NetWorker server and client software requirements
Non-NDMP clients	NetWorker 8.0 or later
NDMP NetApp clients	NetWorker 8.0 or later

Table 77 NetWorker software requirements for checkpoint restart (continued)

Client	NetWorker server and client software requirements
NDMP Isilon clients	NetWorker 8.1 SP1 or later

Platform requirements

Checkpoint restart is only supported on Linux and UNIX environments when performing standard save operations. You cannot use checkpoint restart with block-based backup or parallel save streams enabled.

Checkpoint restart is not supported on Windows platforms.

Client hostname requirements

Use a consistent convention for all NetWorker client hostnames. Do not configure client resources with both short and fully qualified domain names (FQDN).

Save set requirements

Backup of the Windows `DISASTER_RECOVERY:\` save set is not supported. If a client with a `DISASTER_RECOVERY:\` save set is enabled for checkpoint restart, the backup fails.

The checkpoint restart option is ignored for index and bootstrap save sets.

Client Direct requirements

Checkpoint restart supports Client Direct backups only to AFTD devices, and not to DD Boost devices. If a client is enabled for checkpoint restart and a Client Direct backup is attempted to a DD Boost device, then the backup reverts to a traditional storage node backup instead.

For Client Direct backups to AFTDs, checkpoints are made at least 15 seconds apart. Checkpoints are always made after larger files that require more than 15 seconds to back up.

Performance requirements

Enabling checkpoint restart might impact backup speed, depending on the datazone environment and configuration.

Checkpoint restart also might increase the size of the index because additional index records are created for the valid recoverable data. These partial save sets should not be manually removed from the index.

Configuring checkpoint restart

To allow a failed backup for a client to restart from a known good point, you must enable checkpoint restart for the NetWorker Client resource and configure the number of automatic retries for the backup action in the data protection policy.

About this task

When you enable checkpoint restart, you define whether to restart the backup at the directory or file level from the point of failure.

Procedure

1. In the **Administration** window, click **Protection**.
2. From the **View** menu, select **Diagnostic Mode**.

3. In the expanded left pane, select **Clients**.
4. Right-click the client resource and select **Properties**.
The **Client Properties** dialog box appears.
5. On the **General** tab, select the **Checkpoint enabled** checkbox.
6. From the **Checkpoint granularity** list, select whether to restart the backup from the point of failure at the directory or file level:
 - Select **Directory** to restart the backup at the directory level. After each directory is saved, the data is committed to the media and index database. If a directory contains a large number of entries, intermediate checkpoints are created.
 - Select **File** to restart the backup at the file level. Use this option only for save sets with a few large files. Committing every file to the index and the media database is time consuming. Performance degradation may occur for backups that contain many small files.
7. Click **OK** on the **Client Properties** dialog box.
8. Configure the number of times to retry a failed backup:
 - a. In the expanded left pane of the **NetWorker Administration** window, select **Policies**.
 - b. Select the policy.
 - c. In the right pane, select the **Actions** tab.
 - d. Right-click the action and select **Properties**.
The **Policy Action** wizard appears.
 - e. On the **Advanced Options** page, perform the following tasks:
 - a. In the **Retries** box, specify the number of retries that should occur if the backup fails.
 - b. In the **Retry Delay** box, specify a delay in seconds before a failed backup is retried.
 - c. Click **Next**.
 - f. On the **Action Wizard Summary** page, review the settings for the backup action, and then click **Configure**.

Restarting checkpoint-enabled backups

You can configure automatic restarts of checkpoint-enabled backups by specifying the number of retries for the backup action in the data protection policy. You can also manually restart a checkpoint-enabled backup.

NOTICE If you rename a save set, the checkpoint restart fails to find a match against a previous run and the restart reverts to a complete backup. Also, do not edit retention in between checkpoint restarts, as an expired partial save set may leave gaps in the backup set.

Automatically restarting a checkpoint-enabled backup

If the NetWorker server fails to connect to a client for a backup, the **Retries** attribute for the backup specifies the number of times that the server tries the connection to the client before the backup is considered a failure.

The **Retries** attribute applies to a backup regardless of whether the checkpoint restart is enabled for the client. However, a partial save set is created when there is a failure

for a checkpoint-enabled client, and the backup is automatically restarted from the checkpoint until the specified number of retries has been exceeded.

The automatic restart must occur within the restart window that you specify for the workflow for the data protection policy.

Example 1

There are six clients in a group, each with three save sets. The **Retries** attribute for the backup is 1. One save set fails and is checkpoint restarted immediately. The remaining save sets in the group continue to back up. The save set fails a second time. A checkpoint restart for the save set does not occur because the retry attempt would exceed the value for the **Retries** attribute.

When all the save set backup attempts in the group complete, the backup completion report:

- Provides a list of the successful save sets.
- Reports that the failed partial save set is unsuccessful.
- Reports that the backup failed.

Example 2

There are six clients in a group, each with three save sets. The **Retries** attribute for the backup is 2. One save set fails and is checkpoint restarted immediately. The remaining save sets continue to back up. The partial save set fails a second time and is checkpoint restarted immediately. This time, the partial save set succeeds.

When all the save set backup attempts in the group are complete, the backup completion report:

- Provides a list of the successful save sets.
- Reports that the two partial save sets are successful.
- Reports that the backup completed successfully.

Manually restarting a checkpoint-enabled backup

You can manually restart the data protection policy or workflow for a failed backup. For checkpoint-enabled clients, the backup continues from the checkpoint. For other clients, the incomplete save sets are backed up again in full.

Procedure

1. In the **Administration** window, click **Monitoring**.
2. Right-click the policy or workflow for the failed backup, and select **Restart**.
A confirmation message appears.
3. Click **Yes**.

Recovering data from partial save sets

If there is a complete sequence of partial save sets that span the original save set, then you can browse to and recover individual files and directories. If the sequence of partial save sets is incomplete and does not make up the original save set, then you must perform a save set recovery to recover the data from the partial save set.

To recover data from partial save sets that span the original save sets, perform a query for all partial save sets, and then use either the NetWorker User program on Windows or the `recover` program on UNIX to restore the data.

The steps to recover data from a single partial save set are the same as save set recovery from a complete save set. The partial save set contains only files that were successfully backed up. You cannot browse partial save sets.

When you perform a save set recovery of a partial NDMP save set, the recovery process recovers all partial save sets in the checkpoint sequence. You cannot recover data in a partial save set separately from other partial save sets in the checkpoint sequence.

Use the `nsrinfo` command to display the contents of a partial save set.

Probe-based backups

You can configure the NetWorker server to search or probe a NetWorker client for a user-defined script before the start of a scheduled backup operation. A user-defined script is any program that passes a return code.

About this task

When the NetWorker server detects the script, the NetWorker server runs the script and interprets two return codes:

- Return code 0 indicates that a client backup is required.
- Return code 1 indicates that a client backup is not required.

NetWorker interprets all other return codes as an error and does not perform a backup.

Procedure

1. Create the Probe resource script, and save the script in the same directory as the NetWorker binaries on each client that uses the client probe.

The name of the probe script must begin with `save` or `nsr`.

 **Note:** Users are responsible for creating and supporting user-defined scripts.

2. Create the Probe resource on the NetWorker server:
 - a. In the Administration interface, click **Protection**.
 - b. In the expanded left pane, right-click **Probes** and select **New**.

The **Create NSR probe** dialog box appears.

- c. In the **Name** box, specify the name of the probe.
- d. (Optional) In the **Comment** box, specify details for the probe script.
- e. In the **Command** box, type the name and path of the probe script.

 **Note:** The **Command options** box applies to NetWorker Module probes only.

- f. Click **OK**.

3. Associate the probe with a Client resource:
 - a. In the expanded left pane of the **Protection** window, select **Clients**.
 - b. In the right pane, right-click the Client resource, and select **Modify Client Properties**.

The **Client Properties** dialog box appears.

- c. Click the **Apps & Modules** tab.
 - d. Select the probe resource from the **Probe resource name** list.
 - e. Click **OK**.
4. Configure a data protection policy with a workflow that includes a probe action:
 - a. Create a group that includes the client with the assigned probe resource.
 - b. Create a policy.
 - c. Create a workflow.
 - d. Create a probe action and a backup action for the workflow.

Encryption

You can use either AES encryption or in-flight encryption to encrypt data.

The Advanced Encryption Standard feature (AES encryption) encrypts data both in transit and at rest on the backup volume.

In-flight encryption secures data that is in transit.

Note: Do not use the in-flight encryption feature and the AES encryption feature together. Combining the encryption types is redundant and could significantly increase the duration of the backup.

AES Encryption

You can apply password protection and 256-bit data Advanced Encryption Standard (AES) encryption to backup and archive data on UNIX and Windows hosts for additional security.

Note: You can apply password protection alone, AES encryption alone, password protection and encryption together, or compression alone. You cannot apply password protection and compression together or encryption and compression together. Do not apply AES encryption and in-flight encryption together.

When NetWorker uses `aes` to encrypt the backup data, backup times increase. The process of encrypting the data increases CPU and memory usage on the backup client. The impact to CPU and memory resources depends on a number of factors including the load on the host, network speed, and the number of backup files. A backup of a single large file requires less resources than a backup of a dense file system, where NetWorker must access a large number of small-sized files.

Do not use the `aes` ASM for data encryption when backing up files that are encrypted by using the Microsoft Windows Encrypting File System (EFS). The backup is reported as successful, but recovery of the file fails and the following message is written to the NetWorker log file:

```
recover: Error recovering
filename. The RPC call completed before all pipes were
processed.
```

When a backup includes EFS encrypted files, the files are transmitted and stored on backup volumes in their encrypted format. When the files are recovered, they are also recovered in their encrypted format.

Password protection

AES Encryption is supported through the use of the `aes` Application Specific Module (ASM) based on the password that is defined on the UNIX or Windows host. If a

password is not defined on the host, then data is encrypted with the default password that is configured for the NetWorker server.

NOTICE You must specify the password to recover password-protected files. If the password was configured or changed after the backup occurred, then you must provide the password that was in effect when the file was originally backed up. Keep password changes to a minimum.

Configuring encryption for scheduled backups

Procedure

1. Configure a password on the host.
To configure the password on a Windows host:
 - a. Select **Options > Password** in the NetWorker User program.
 - b. Type a password.
2. Configure the default password on the NetWorker server:
 - a. In the **Administration** window, click **Protection**.
 - b. In the left pane, right-click the NetWorker server, and select **Properties**.
The **Server Properties** dialog box appears, starting with the **Setup** tab.
 - c. Click the **Configuration** tab.
 - d. Type the password in the **Datazone pass phrase** attribute.
 - e. Click **OK**.
3. Configure a directive for the Client resource with the `aes` ASM for encryption.
You can use the Encryption global directive to apply encryption. You can also configure a local directive on the client computer. [Directives](#) on page 396 provides more information.

Configuring AES encryption or password protection for manual backups

When you perform a manual backup on Windows with the NetWorker User program, you can specify AES encryption or password protection.

Procedure

1. Configure a password on the Windows host:
 - a. Open the NetWorker User program.
 - b. Select **Options > Password**.
 - c. Type the password in the **Password** dialog box and click **OK**.
2. Open the NetWorker User program and click **Backup**.
3. Select the data to back up.
4. From the **File** menu, select **Special Handling**.
The **Special Handling** dialog box appears.
5. Select the handling method for the backup data:
 - **Password Protect**
 - **Password Protect and Encrypt**
6. Click **OK**.

7. Click **Start** to start the backup.

In-flight encryption

In-flight encryption secures data that is in transit. By default, the in-flight encryption feature is not enabled in NetWorker.

About this task

Backup times might be longer with in-flight encryption than with AES encryption. Because in-flight decryption occurs on the target NetWorker storage node, CPU and memory usage could significantly increase. The impact to CPU and memory resources depends on a number of factors including the load on the host, network speed, and the number of backup files.

You can enable in-flight encryption from the NMC or from `nsradmin`.

Note:

- Do not use the in-flight encryption feature and the AES encryption feature together. Combining the encryption types is redundant and could significantly increase the duration of the backup.
- Do not use in-flight encryption to backup and recover to Data Domain devices (DDBoost). Refer to the Data Domain documentation set to configure DDBoost encryption.
- In-flight encryption is not supported for a client direct backup and recovery operation from a NetWorker client host over a network to a remote host's AFTD. If in-flight encryption is enabled, data is not encrypted in-flight over the network.
Use AES encryption for a client direct save operation from a NetWorker client host over a network to a remote host's AFTD.

Using NMC to configure in-flight encryption for the NetWorker server

Procedure

1. Use NMC to connect to the NetWorker server.
2. On the **NetWorker Administration** window, select **Hosts**.
3. Right-click the hostname of the NetWorker server.
4. Select **Configure Local Agent**. The **Local Agent Properties** window appears.
5. Go to the **Advanced** tab and select **Connection encrypted**.
6. Click **OK**.

Using nsradmin to configure in-flight encryption for a NetWorker client

About this task

Use the following procedure to enable in-flight encryption for a NetWorker client's save set data.

Procedure

1. Log in as root or as Windows Administrator on the NetWorker client.
2. Type the following at the command prompt:

```
nsradmin -p nsrexec
```

The `nsradmin` prompt appears.

3. Edit the NSRLA resource by typing the following command:

```
print type:NSRLA
```

4. Change the value of the *connection encrypted* attribute in the NSRLA resource to enabled.

Type the following line at the `nsradmin` prompt:

```
update connection encrypted:enabled;
```

5. Type **Yes** when prompted to confirm the change.
6. If the *auth method* attribute is not set, ensure that the peer certificate for the NetWorker client matches the storage node.

Results

NOTICE When you modify an attribute with the `nsradmin` program, you must specify the attribute name and value correctly. If you do not specify the attribute name and value correctly, the `nsradmin` program does not update the attribute and `nsradmin` does not provide an error message.

The *EMC NetWorker Security Configuration Guide* provides more information about the `nsrexec` database and how to modify attributes in the `nsrexec` database.

Compression

You can compress backup data to reduce network traffic and backup storage requirements.

Compressing data for a backup generates less network traffic. However, compression uses computing resources, so its benefits may be limited on low-powered systems. If the storage device also compresses data, the result may be that more data is actually written to tape.

Note: You can apply password protection alone, encryption alone, password protection and encryption together, or compression alone. You cannot apply password protection and compression together or encryption and compression together.

Configuring compression for scheduled backups

About this task

Configure a directive for the Client resource with the `compressasm` ASM for compression.

You can use one of the global directives with compression or configure a local directive on the client computer.

Configuring compression for manual backups

About this task

The methods of configuring compression for UNIX and Windows differ.

To compress data for a manual backup on UNIX, you must use the `compressasm` ASM in a local directive file.

To configure data for a manual backup on Windows, use either the `compressasm` ASM in a local directive file, or use the following procedure.

Procedure

1. Configure a password on the Windows host.
 - a. Open the NetWorker User program.
 - b. Select **Options > Password**.
 - c. Type the password in the **Password** dialog box and click **OK**.
2. Open the NetWorker User program and click **Backup**.
3. Select the data to back up.
4. From the **File** menu, select **Special Handling**.
The **Special Handling** dialog box appears.
5. Select **Compress** as the handling method for the backup data.
6. Click **OK**.
7. Click **Start** to start the backup.

Configuring Client Direct backups

NetWorker clients with network access to AFTD or DD Boost storage devices can bypass the NetWorker storage node and send backup data directly to the devices. This type of backup is called a Client Direct backup.

The storage node manages the devices for the NetWorker clients, but does not handle the backup data.

A Client Direct backup reduces bandwidth usage and bottlenecks at the storage node, and provides highly efficient backup data transmission.

If a Client Direct backup is not available, a traditional storage node backup occurs instead.

Requirements for Client Direct backups

Ensure that the environment meets the following requirements to perform Client Direct backups:

- NetWorker clients on UNIX/Linux or Microsoft Windows can perform non-root and cross-platform Client Direct backups to AFTDs. The AFTD can be managed by either a UNIX/Linux or a Windows storage node, and can be either local or mountable on the storage node.
To perform non-root and cross-platform Client Direct backups to AFTDs, the NetWorker server and the storage node software must be version 8.1 or later.
- If an NFS server provides the AFTD storage for Client Direct backups, then the NFS server must permit access by using the NFSv3 protocol with AUTH_SYS (AUTH_UNIX) authentication. The NFS server also must not restrict access to clients by using only privileged ports.
- If you enable checkpoint restart for a client, then Client Direct backups are supported only to AFTDs, and not to DD Boost devices. If a client is enabled for checkpoint restart and a Client Direct backup is tried to a DD Boost device, then the backup reverts to a traditional storage node backup instead.

For Client Direct backups to AFTDs, checkpoint restart points are made at least 15 seconds apart. Checkpoints are always made after larger files that require more than 15 seconds to back up.

- Archive operations are not currently supported for Client Direct backups.

Configuring Client Direct backups

Procedure

1. Ensure that the clients that perform Client Direct backups have a network connection and a remote network protocol to reach the storage device.

Windows clients can use a CIFS or NFS path, although a CIFS path generally yields better performance. UNIX clients must use an NFS path.

2. Specify the complete path for the destination device in the **Device access information** attribute on the **General** tab of the **Device Properties** dialog box for the destination device.

Keep in mind the following points when you specify the path:


- If the storage device is directly attached to a Windows storage node, then the storage node uses a different path than the Client Direct clients. If the storage device is not directly attached to any storage node, then the path is the same for all storage nodes and Client Direct clients.
- The device access information path should include multiple access paths to cover local and remote use cases.
- To specify an NFS path, use the `NFS_host:/path` format regardless of whether the AFTD is local to the storage node or mountable on the storage node. Non-root UNIX/Linux NetWorker clients require this NFS format for Client Direct access.
- For Windows Client Direct backups, specify a CIFS path instead of an NFS path. A CIFS path generally yields better performance.
- If you are setting up an AFTD on a Windows storage node, specify the CIFS path first. For example:

```
\\fileserver\aftd1
fileserver:/aftd1
```

- If you are setting up a UNIX/Linux storage node, specify the NFS path first. For example:

```
fileserver:/aftd1
\\fileserver\aftd1
```

The following figure shows an example set of paths for a CIFS AFTD.

Figure 47 Paths for CIFS AFTD

3. If an NFS server provides the AFTD storage for Client Direct backups, then specify the username and password that is required to access the NFS server for the AFTD in the **Remote user** and **Password** attributes on the **Configuration** tab of the **Device Properties** dialog box for the device.
4. Ensure that the **Client direct** attribute is enabled on the **General** tab of the **Client Properties** dialog box for each Client Direct client.

Client Direct backups are enabled by default.

Select **View > Diagnostic Mode** in the Administration interface to access the **Client direct** attribute in the **Client Properties** dialog box.

Backup command customization

You can customize client backups by creating additional programs (scripts) that affect the way the NetWorker server will back up client file system data.

NetWorker provides you with the following features, which enable you to customize scheduled backups for a client:

- Create a custom backup script that starts the `save` command.
- Create a script file that performs operations before the start of a backup.
- Create a script file that performs operations after the backup of all save sets for a client completes.

For example, you can create a custom backup script that performs the following tasks:

1. Shuts down a mail server or database before the NetWorker server performs a backup.
2. Prints a message such as `Backup started at 3:33 A.M.`
3. Starts the `save` command and performs a backup.
4. Prints a message such as `Backup completed at 6:30 A.M.`
5. Restarts the mail server or database after the backup completes.

Creating a custom backup script

Create a script that runs the `save` program as part of its instructions to customize behavior of scheduled backups of a client. When NetWorker performs a back up of the

client, NetWorker runs the customized program for each save set instead of the standard `save` program.

Procedure

1. Use a text editor to create a script in the `networker_installation_dir\bin` directory on Windows clients or the `networker_installation_dir/bin` on LINUX or UNIX clients.

The script file must meet the following requirements:

- The name starts with `save` or `nsr`.
- The name contains a maximum of 64 characters.
- For Windows, the script file must end with a `.bat` extension.
- For UNIX, the script file must have executable file permissions.

For example, script file names that meet these criteria include `save_custom_script.bat` and `nsr_backup_script.bat` for windows, and `save_custom_script.sh` and `nsr_backup_script.sh` for Linux and UNIX.

2. Add commands to the script in the following order:
 - a. Declare all required environment variables, for example the `PATH` variable.
 - b. (Optional) Run a preprocessing command before each save set backup.
 - c. (Required) Back up the data by using the NetWorker `save` command. Always specify the full path of the `save` command in the script.

On UNIX and Linux hosts, run the NetWorker `save` command with the arguments `save "$@"` to enable the `save` command to accept the arguments that the NetWorker `savefs` program would run during a regular backup.
 - d. (Optional) Run a postprocessing command after each save set backup.

 **Note:** All commands within the script must complete successfully. Otherwise, the NetWorker server cannot complete the remaining instructions.

3. Save and close the script file.
4. Specify the name of the backup script in the **Backup command** attribute for the Client resource:
 - a. In the **Administration** window, click **Protection**.
 - b. In the expanded left pane, select **Clients**.
 - c. Right-click the Client resource, and select **Modify Client Properties**.
The **Client Properties** dialog box appears, starting with the **General** tab.
 - d. Select the **Apps & Modules** tab.
 - e. Type the name of the backup script in the **Backup command** box.
 - f. Click **OK**.
5. Back up the client to ensure that the new backup command works.

Results

NetWorker logs information about the backup status in separate log files, and not in the `save` output.

[Reporting policy status and backup job status](#) on page 685 provides more information about how to review backup job status.

Example backup script on Windows

In this example backup script for a Windows client computer, the customized backup program runs pre-backup commands, the NetWorker `save` command, and then post-backup commands.

Description of the example script

The following table provides details on each type of command in the example backup script.

Table 78 Example backup script on Windows

Command type	Description
Pre-backup	Redirects the output of the <code>net start</code> DOS command to create a <code>netstart.txt</code> file at the root of the <code>C:\</code> drive, and sends all information about started services for the current computer to this file.
<code>save</code>	Runs NetWorker commands that are required to start the backup process.
Post-backup	Redirects the output of the <code>set</code> DOS command to a <code>set.txt</code> file at the root of the <code>C:\</code> drive, and sends all computer system environment information to this file.

The `netstart.txt` and `set.txt` files are placed in the `C:\` directory. New information is appended to these files each time a backup is run.

Example script

```
@ECHO OFF
SETLOCAL
ECHO =====START BATCH FILE=====
ECHO =====NetWorker PRE_BACKUP COMMAND=====
ECHO =====NET START - creates netstart.txt file and
ECHO =====sends all Started Services information
ECHO =====to the file c:\netstart.txt

NET START >>C:\NETSTART.TXT

REM This command takes incoming arguments from
REM the savegrp command and handle them
REM to overcome batch file limitations:

REM PARSE ALL INCOMING ARGUMENTS
REM and pass single argument in case
REM more than 10 arguments are passed to this file
REM (ie %0-%9 is not enough).

ECHO =====NetWorker SAVE SET COMMAND=====
SHIFT
SET arg=%0

:loop
SHIFT
```

```

IF %0.==. GOTO save
SET arg=%arg% %0
GOTO loop

REM These are the save commands that run the required
REM NetWorker backup commands.

:save

REM Note: Enter correct path to your NetWorker bin
REM directory (line below is default path)
C:\PROGRA~1\nsr\bin\save.exe %arg%

ECHO =====NetWorker POST_BACKUP COMMAND=====
ECHO ====="SET" - creates set.txt file and sends all
ECHO =====computer system environment information to
ECHO =====C:\set.txt file=====

SET >>C:\SET.TXT

ECHO =====END OF BATCH FILE=====

ENDLOCAL

```

Monitoring details for the script

The following information appears in the **Monitoring** window of the Administration interface and the backup action log file. After the backup process completes, review the log output to verify the execution of the commands in the script.

```

--- Successful Save Sets ---
:* jupiter:c:\inetpub =====START BATCH FILE=====
* jupiter:c:\inetpub ===NetWorker PRE_BACKUP COMMAND===
* jupiter:c:\inetpub=====NET START
* creates netstart.txt file and sends all started
* jupiter:c:\inetpub =====services information to
* that file c:\netstart.txt==

* jupiter:c:\inetpub ===NetWorker SAVE SET COMMAND=====
* jupiter:c:\inetpub save: using 'C:\Inetpub' for
* 'c:\inetpub'
jupiter: c:\inetpub level=full,194 KB 00:00:02 37 files
* jupiter:c:\inetpub =====NetWorker POST_BACKUP COMMAND
* jupiter:c:\inetpub ====="SET" - creates set.txt
* file and sends all computer system
* jupiter:c:\inetpub ==== environment information
* to C:\set.txt file
* jupiter:c:\inetpub =====END OF BATCH FILE=====

```

Example backup script on UNIX

This example script on UNIX locks a ClearCase version object base (VOB), performs the backup, and then unlocks the VOB.

```

#!/bin/sh
# export the SHELL that we are going to use
SHELL=/bin/sh
export SHELL
# export the correct PATH so that all the required binaries can be
found

```

```

case $0 in
/* ) PATH=/usr/atria/bin:/bin:/usr/bin:`/bin/dirname $0`
c=`/bin/basename $0`
;;
* ) PATH=/usr/atria/bin:/bin:/usr/bin:/usr/sbin
c=$0
;;
esac
export PATH

# These are the valid statuses that save reports upon completion of
the backup
statuses="
failed.
abandoned.
succeeded.
completed savetime=
"

# Perform the PRECMD (Lock VOB)
/usr/atria/bin/cleartool setview -exec "/usr/atria/bin/
cleartoollock -c \
'VOB backups in progress' -vob /cm_data/mis_dev" magic_view >
/tmp/voblock.log 2>&1
# Perform backup on client
save "$@" > /tmp/saveout$$ 2>&1
# cat out the save output
cat /tmp/saveout$$
# search for backup status in output reported by save
for i in ${statuses}; do
    result=`grep "${i}" /tmp/saveout$$`
    if [ $? != 0 ]; then
        echo ${result}
    fi
done
# Perform the POSTCMD (Unlock VOB)
/usr/atria/bin/cleartool setview -exec "/usr/atria/bin/
cleartoolunlock -vob
/cm_data/mis_dev" \
magic_view > /tmp/vobunlock.log 2>&
# exit gracefully out of the shell script
exit 0

```

Table 79 NetWorker Server Versions

NetWorker Server Versions	NetWorker client version configures with the NetWorker Server	Client properties need to be updated
8.2.x	8.2.x	backup command <code>savepnpc</code>
9.x	8.2.x	backup command <code>savepnpc</code>
9.x	9.x	pre command and post command

Controlling exit status reporting for a custom backup script

Use the **Job control** attribute on the **Apps & Modules** tab of the **Client Properties** dialog box for a Client resource to control how end of job and exit status messages are determined for a custom backup script.

To access the **Job control** attribute, select **View > Diagnostic Mode** in the Administration interface to enable diagnostic mode view. A checkmark next to **Diagnostic Mode** in the **View** menu indicates that diagnostic mode view is enabled.

There are three check boxes for the **Job control** attribute:

- **end on job end**
- **end on process exit**
- **use process exit code**

The following table provides details on exit status reporting depending on the selection of one or more of the checkboxes.

Table 80 Job control attribute selections

Selections	Description
No selections (default behavior)	The <code>nsrpolicy</code> and <code>nsrjobd</code> programs determine the success or failure of a custom script based on the completion of the <code>save</code> program (end of job). The following criteria apply: • If the <code>save</code> job completion status is success, then <code>nsrpolicy</code> and <code>nsrjobd</code> report that the custom backup job succeeded. • If the <code>save</code> job completion status is failure, then <code>nsrpolicy</code> and <code>nsrjobd</code> report that the custom backup job failed. • If no completion status is received, the custom job output is examined for completed <code>savetime=savetime</code> lines. If found and the <code>savetime</code> is a value other than 0 (zero), then the custom backup job is considered to have succeeded. If the value is 0, then the custom backup job is considered to have failed. The exit code of the custom script process is not taken into consideration.
end on job end only	A backup job is considered to be ended as soon as an end job message is received from the <code>save</code> command. Select this option when you do not want to wait for the postprocessing commands of the script to end.
end on process exit only	A backup job is considered to be ended as soon as the started process exits. Background processes started by the backup command could still be running on the client. Use this option when you want the custom script to start background processes and you do not want <code>savegrp</code> or <code>nsrjobd</code> to wait for the processes to complete.
use process exit code only	Only the process exit code is used to determine the success or failure of the job. An exit code of 0 indicates success. Otherwise, the job is reported as failed. Use this option when you want the script postprocessing command status to have an impact on the status of the <code>save</code> backup command without having to unset the <code>NSR_STD_MSG_FD</code> environment variable. If the script invokes more than one NetWorker backup command such as <code>save</code>, then you must still unset the <code>NSR_STD_MSG_FD</code> environment variable.
Both end on job end and end on process exit	Either event can trigger the end of a job.

Table 80 Job control attribute selections (continued)

Selections	Description
Both end on job end and use process exit code	If an end job message is received before the process exits, then the exit status provided by the end job message is used to determine the success or failure of the job.

Customizing backups with the pre and post commands

Customize backup behavior by running preprocessing and postprocessing commands only once during the client backup, instead of once for each save set.

About this task

Preprocessing and postprocessing scripts can be useful if the client is running a database or another program that should be stopped before the client is backed up, and then restarted after the backup has completed.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left panel, select **Clients**.
3. Right-click the client resource and select **Modify Client Properties**.
The **Client Properties** dialog box appears.
4. On the **Apps & Modules** tab, in the **Pre command** attribute, specify the name of the script file that you require NetWorker to run before a backup.
 **Note:** Do not specify the path to the file.
5. Optionally, in the **Post command** attribute, specify the name of the script file that you require NetWorker to run after a backup of all the save sets for the client completes.
 **Note:** Do not specify the path to the file.
6. Click **OK**.

Results

The customized instructions are applied the next time that the client is backed up.

Client resources

A client is both a physical computer with NetWorker client software installed on it and a NetWorker *resource* that specifies a set of files and directories to be in a scheduled backup. A Client resource also controls backup settings for the client, such as the save sets to back up for the client, the groups to which the client belongs, and whether to automatically restart failed backups for the client.

You can configure multiple Client resources for a single NetWorker client computer, although clients with the same save set cannot be in the same group. You might want to create multiple Client resources for a single client computer in the following scenarios:

- To segregate different types of backup data, such as application data and operating system files. For instance, to back up the accounting data on a computer on a different schedule than the operating system files, create two client resources for the computer: one for accounting data and another for operating system data.
- To back up large client file systems more efficiently. For instance, you could create separate client resources for each file system on a computer and back them up on different schedules.

You can create a Client resource either by using the **Client Backup Configuration** wizard or the **Client Properties** dialog box.

You can configure NetWorker clients to use a unique network interface on the NetWorker server and storage node for backup and recovery operations. [Using multihomed systems](#) on page 871 provides more information.

Create a Client resource with the Client Properties dialog box

The following procedure provides the basic steps to create a client resource for scheduled backups. Additional configuration of the Client resource may be necessary for clients such as VMware or NAS device clients, or to take advantage of product features such as probe-based backups or archiving.

Before you begin

- Install the NetWorker Client software on the client computer.
- (Optional) Configure directives to control how the NetWorker Server processes files and directories during backup and recovery. For example, you can create a directive to skip certain directories or file types, to compress backup data, or to encrypt backup data. [Directives](#) on page 396 provides more information.
- (Optional) To view advanced options in the **Client Properties** dialog box, select **View > Diagnostic Mode** in the **Administration** window. Advanced options are not discussed in this procedure.
- (Optional) Check for reverse entries in the DNS server. If the reverse entries are not present, then do not use the IP address for creating the clients.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Clients**.
3. From the **File** menu, select **New**.

The **Client Properties** dialog box appears, starting with the **General** tab.

4. In the **Name** box, type the hostname of the client computer.

 **Note:** In NetWorker 19.1, users can configure NSM with NetApp ONTAP. The *NetWorker Snapshot Management for NAS Devices Integration Guide* provides more information on how to configure NSM with NetApp ONTAP 7-Mode and ONTAP Cluster Mode.

5. (Optional) In the **Comment** box, type a description of the client.

If multiple Client resources are being set up for the same host, type a comment that distinguishes the Client resources.

6. In the **Tag** box, type one or more tags to identify this Client resource for the creation of dynamic client groups for data protection policies.

Place each entry on a separate line.

7. To allow a failed backup operation to restart at a known good point before the point of failure during the backup, select the **Checkpoint enabled** checkbox.
[Configuring checkpoint restart backups](#) on page 465 provides more information on the requirements for checkpoint restart.
8. From the **Directive** list, select a directive to control how the NetWorker Server processes files and directories during backup and recovery.
9. In the **Save set** box, type the name of the files or directories to back up, or click the **Browse** button to browse and select file system objects.

Note: To avoid the over consumption of memory, NetWorker limits the number of files that you can view when you browse a directory that contain a large number of files, for example, 200,000 files. When NetWorker determines that displaying the number of files will exhaust memory resources, NetWorker will display a partial list of the files and a message similar to the following appears: Expanding this directory has stopped because the result has too many entries

When you manually specify the save set value, place multiple entries on separate lines. For example, to back up a log file directory that is named `C:\log` and all the data under the directory that is named `D:\accounting`, type the following entries:

```
C:\log
D:\accounting
```

Follow the guidelines in the section "Mapped drives" to back up mapped drives on Windows systems.

To back up all client data, type `ALL`. For Windows operating systems, the `ALL` save set includes the `DISASTER_RECOVERY:\` save set, which includes the `WINDOWS ROLES AND FEATURES` save set.

NOTICE Some operating systems contain files and directories that should not be backed up. Use directives to ensure that these files and directories are not backed up.

[Save sets](#) on page 356 provides more information on defining the save sets for a Client resource.

10. Select the other tabs in the **Client Properties** dialog box and configure options as necessary.
11. Click **OK**.

Results

Verify that the client is enabled for scheduled backups by ensuring that a check mark appears next to the client in the **Scheduled backup** column in the right pane for the client.

Editing a Client resource

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Clients**.
3. In the right pane, perform one of the following tasks:

- To modify multiple attributes in a single configuration resource by using the **Client Properties** window, right-click the staging configuration and select **Modify Client Properties**.
- To modify a specific attribute that appears in the resource window, place the mouse in the cell that contains the attribute that you want to change, then right-click. The menu displays an option to edit the attribute. For example, to modify the **Comment** attribute, right-click the resource in the **Comment** cell and select **Edit Comment**.

 **Note:** To modify a specific attribute for multiple resources, press and hold the **Ctrl** key, select each resource, and then right-click in the cell that contains the attribute that you want to change. The menu displays an option to edit the attribute.

4. Edit the attributes of the Client resource.
5. Click **OK**.

Client priority

The **Priority** attribute on the **Globals (1 of 2)** tab of the **Client Properties** dialog box for a Client resource enables you to control the order in which the NetWorker server contacts clients for backup.

The attribute can contain a value between 1 and 1,000. The lower the value, the higher the priority.

You must select **View > Diagnostic Mode** in the Administration interface to access the **Priority** attribute in the **Client Properties** dialog box.

During a backup operation, the NetWorker server contacts the client with the lowest priority value first. If you do not specify a priority for the Client resources, then the backup order is random.

While the **Priority** attribute specifies the order of client contact, many variables affect the order in which clients complete their backups. For example:

- The backup operation on a client does not begin until the worklists for each of the save sets on the client are complete.
- The amount of work can vary greatly from one client to the next.
- If a client stops responding and times out, then the backup operation puts the client backup at the end of the backup order list.

The only way to guarantee that the backup of one client occurs before the backup of another client is to configure the data protection policies for the clients to start at different times.

Copying a Client resource

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Clients**.
3. In the right pane, right-click the Client resource and select **Copy**.

The **Create Client** dialog box appears with the same attributes as the original client except for the client name.

4. Type the hostname of the client in the **Name** box.
5. (Optional) Edit other attributes for the Client resource.

6. Click **OK**.

Changing the hostname of a client

To change the hostname of a client, you must delete the Client resource, rename the directory with the client file index for the client, and then create a Client resource with the new hostname and the original client ID.

About this task

If you create the new Client resource but do not use the client ID of the original NetWorker host:

- The NetWorker server considers the new hostname to be a new NetWorker host.
- The NetWorker server assigns the new hostname a new client ID.
- To recover data, you must perform a directed recovery from the original hostname to the new hostname.
- You cannot perform a browsable recovery, only a save set recovery.

Use the `nsrclientfix` command to analyze the media database and identify client ID inconsistencies. To resolve client ID issues, use the `nsrclientfix` command to merge information about multiple clients in the media database and resource database into one client resource with the original client ID. The following KB articles on the Online Support website provide more information about using the `nsrclientfix` command:

- For NetWorker Server client ID issues: 000185727
- For NetWorker Client client ID issues: 000193911

Procedure

1. Record the client ID of the original Client resource:
 - a. Enable diagnostic mode view by selecting **View > Diagnostic Mode** in the **Administration** window.
 - b. In the **Administration** window, click **Protection**.
 - c. In the expanded left pane, select **Clients**.
 - d. In the right pane, right-click the Client resource and select **Modify Client Properties**.
The **Client Properties** dialog box appears.
 - e. Select the **Globals (1 of 2)** tab.
 - f. Record the value in the **Client ID** attribute.
 - g. Click **Cancel**.
2. Delete the Client resource:
 - a. Right-click the resource, and select **Delete**.
A confirmation message appears.
 - b. Click **Yes**.
3. Stop all the NetWorker services on the NetWorker server.
4. On the NetWorker server, rename the client file index directory for this client from `old_client_name.domain.com` to `new_client_name.domain.com`.

The default location for the client file index is `NetWorker_install_path\index\client_name.domain.com` on Windows and `/nsr/index/client_name.domain.com` on UNIX/Linux.

5. Restart the NetWorker services on the NetWorker server.
6. Create a Client resource with the new hostname and the original client ID.

Deleting a Client resource

When you delete a Client resource, the NetWorker server can no longer back up the client computer. The backup history for the client remains in the client file index and media database until the entries are removed. You can still access and recover backup data for the client directly from the volume that contains the data by using the `scanner` command.

About this task

If you create a Client resource to re-create the deleted client, specify the same hostname for the client. The NetWorker server recalls and uses the original client ID for the hostname.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, click **Clients**.
3. In the right pane, right-click the Client resource and select **Delete**.

A confirmation message appears.

4. Click **Yes**.

Manual backups

Manual backups enable users to make quick backups of a few files from the client host.

When you perform a client-initiated or manual backup, by default NetWorker backs up the data to a volume assigned to the Default pool on the NetWorker server. The retention policy that is assigned to the data is one year, and the level is manual.

Perform manual backups on Windows by using the NetWorker User program. Perform manual backups on UNIX and Linux only from the command line.

Performing a manual backup on Windows

Before you begin

Create a local directive on the client computer to exclude local file type devices from manual backups with the NetWorker User program:

1. Start the NetWorker User program.
2. From the **Options** menu, select **Local Backup Directives**.
3. Clear the checkbox for the local file type device.
4. From the **File** menu, select **Save Directive**.

[NetWorker User local directives](#) on page 405 provides more information on local directives.

Note: You cannot perform data deduplication during backups with the NetWorker User program. You must perform scheduled backups or manual backups from the command line to perform data deduplication during the backup.

Procedure

1. In the NetWorker User program, click **Backup**.

The **Backup** window appears.

2. Select the data to back up.

To back up critical volumes, UEFI, the system reserved partition, and `WINDOWS ROLES AND FEATURES` for disaster recovery purposes, select the `DISASTER_RECOVERY` save set.

3. Click **Start**.

The **Backup Status** dialog box displays the progress of the backup. When the backup finishes, a `Backup completion time` message appears.

If the backup fails due to a problem with VSS or a writer, an error message appears. Use the Windows Event Viewer to examine the event logs for more information. VSS backup error messages are also written to the NetWorker log file.

The NetWorker log file in `\install_path\logs\networkr.raw` contains a record of every file that was part of an attempted manual backup from the NetWorker User program. This file is overwritten with the next manual backup. To save the information in the file, rename the file or export the information by using the `nsr_render_log` program.

NOTICE Certain types of corrupt files or errors on computer disk volumes are not detected. NetWorker might back up this corrupt data. To avoid this situation, run diagnostic programs regularly to correct disk volume errors.

Including Windows BMR in manual backups

When you use the NetWorker User program to back up a host, to ensure the backup operation will backup all of the data on the host, select **Computer** in the **Backup** window.

About this task

If you only select the `DISASTER_RECOVERY:\` save set, then the NetWorker User program automatically selects the critical volumes and `WINDOWS ROLES AND FEATURES` save sets.

Note: When you use the NetWorker User program or the `save` command to perform a manual backup, NetWorker performs the backup operation as a single backup stream. To multi-stream the backup operation, run a scheduled group backup.

[Backing Up Data](#) on page 409 provides more information about manual backups.

Performing a manual backup from the command prompt

Perform a manual backup from the command prompt by using the `save` command.

For example, to back up `myfile` to the `jupiterserver`, type:

```
save -s jupiter myfile
```

If you do not specify the `-s` option with the `save` command, the files are backed up to the NetWorker server that is alphabetically listed first in the `/nsr/res/servers` file on the client computer.

UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `save` command.

BTRFS backups

NetWorker support BTRFS volume backups. When you specify a BTRFS volume or sub-volume save set, NetWorker performs a recursive back up of the directory tree that you specified with the `save` command. When NetWorker encounters a sub-directory that has a sub-volume ID that differs from the parent sub-volume ID, NetWorker will not back up the contents of the subdirectory, unless you specify the `-x` option with the `save` command.

Performing a manual backup on Mac OS X

To perform a manual backup on a Mac OS X client, use the `save` command in a Terminal session.

About this task

For example:

```
$ save "file_or_directory_to_back_up" -s NetWorker_server
```

If you do not specify the `-s NetWorker_server` option, the `save` command contacts the NetWorker server that is defined in the `/nsr/res/servers` file. The *NetWorker Command Reference Guide* provides more information about the `save` command.

Troubleshooting manual backups

This section describes how to troubleshoot error messages that might appear during a manual or client-initiated backup.

Could not create log file: Permission denied

This message appears when a non-root user performs a manual client direct-enabled backup to a CloudBoost device but the user account does not have write access to the `/nsr/logs/cloudboost` directory. To resolve this issue, configure the following environment variables to define an alternate location for the log files, where the non-root user has write access:

```
export CB_CACHE_LOCATION=cache_dir
```

```
export CB_LOG_DIR_LOCATION=log_dir
```

where:

- `cache_dir` is the directory that stores backup cache files.
- `log_dir` is the directory that stores for the backup log files.

Verifying backup data

You can use the NetWorker User program on Windows clients to ensure that backup data on the NetWorker server matches the data on the local disk. This verification process enables you to test whether you can successfully recover the data.

About this task

During the verification, the file types, file change times, file sizes, and file contents are compared. Other system attributes, such as read-only, archive, hidden, system, compressed, and file access control list (ACL), are not part of the verification.

The NetWorker server alerts you to any changes that have occurred to the data since the backup. Verification also determines whether a hardware failure kept the NetWorker server from completing a successful backup.

 **NOTICE** This feature is not available on UNIX clients.

Procedure

1. Log in as an administrator on the Windows client computer.
2. Open the NetWorker User program.
3. From the **Operation** menu, select **Verify Files**.
4. Select the data items to verify.
5. Click **Start**.
6. Monitor the data verification progress in the **Verify Files Status** dialog box.

After the verification is complete, the **Verify Status** dialog box shows any data discrepancies.

CHAPTER 7

Cloning, Staging, and Archiving

This chapter contains the following topics:

- [Cloning, staging, and archiving](#).....494
- [Benefits of cloning and staging](#)..... 494
- [Cloning save sets and volumes](#).....495
- [Staging save sets](#).....509
- [Archiving data](#).....516

Cloning, staging, and archiving

The storage device that you use for the initial backup is often a compromise between a number of factors, including location, availability, capacity, speed, and cost. As a result, the backup data on the initial storage device is unlikely to be on the ideal or best storage for the entire duration of the retention period.

NetWorker provides you with three ways to manage data for long term storage.

- Cloning—The clone process allows you to perform the following tasks:
 - Create a duplicate copy of backup data securely offsite.
 - Transfer data from one location to another.
 - Verify backups.

You can clone volumes and save sets. The clone process copies existing save sets from a volume in one device to a volume in a different device. The target volume can be the same media type or a different media type than the original.

- Staging—The stage process uses the clone process to transfer backup data from an AFTD or file type device to another medium, then removes the data from the original location.
- Archiving—The archive process captures files or directories as they exist at a specific time, and writes the data to archive storage volumes. NetWorker does not automatically recycle the archive volumes. After the archive process completes, you can delete or groom the original files from the disk to conserve space.

 **Note:** Cloning of APP Consistent TLOG backup will fail when you use CloudBoost 19.1 embedded storage node. For cloning APP consistent full and Tlog backups to CloudBoost 19.1, use the external storage node with NetWorker 19.1.

Benefits of cloning and staging

Cloning and staging enables you to use storage devices more effectively by moving data between different types of devices. You can copy the data that is stored on local tape devices to other devices in remote locations without an impact to the initial backup performance. You can copy backups from disk devices to tape device to facilitate offsite or long term storage. When you move data from disk to tape, you can use the storage capacity more effectively. When you make use of a deduplicated disk, NetWorker can reclaim the initial storage space for new backups.

NetWorker can only perform a clone operation after a successful backup, which provides the following benefits:

- Allows the backup process to complete at maximum performance without any impact to speed due to multiple write acknowledgments, delays, or retries on one or more devices. A clone operation limits the performance impact on a client, while providing data protection as quickly as possible.
- Ensures that a successful backup, that the data is valid, and that the clone operation completes successfully.
- Ensures that the storage requirements have been determined, and that the storage is made available.
- Allows you to schedule and rank the clone operation outside of the backup window, when resources are less constrained.

- Reduces the load on the backup infrastructure.
- Allows you to easily start recoveries because the backup operation has already completed.

Note: You cannot use the NetWorker software to create an instant clone by writing to two devices simultaneously. This operation is also referred to as parallel cloning, twinning, or inline copy. Where parallel cloning or twinning is required, consider using the NetWorker cloning feature. Using cloning helps ensure that the initial backup completes successfully. Additional data protection can also be implemented by using the best devices and bandwidth available for the backup environment.

Cloning save sets and volumes

The cloning operation reads save sets from a volume within a backup or archive pool and writes the data to a volume in a clone pool. You can clone save sets multiple times, but NetWorker must write each clone to a separate volume.

When you clone backup data, the clone operation validate that NetWorker can read the original backup data successfully in the media database and on the media volume, which provides additional assurance that you can recover the data.

To schedule save set cloning, configure Data Protection Group, Data Protection Policy, followed by a workflow having a clone action. The Data Protection Policies chapter provides detailed information about creating a clone action. To manually clone backup save sets or to clone the backup volume itself from the command prompt, use the `nsrclone` command.

Deciding when to clone

The need to clone data is normally driven by a requirement for additional protection, or the need to move data to a specific media type or location. In both cases, the priority is to secure the data as quickly as possible.

There is a high probability that any restore request within the first 48 hours is due to local failure or corruption and that the original backup copy is the most likely source for that recovery. If there is a local disaster recovery or site loss, the recovery actions and objectives are likely to be very different. Selected systems and services are assigned specific priorities, recovery point objective (RPO) values, and recovery time objective (RTO) values.

Clone retention

NetWorker supports the ability to define a retention time for a clone save set that differs from the original save set.

The following attributes determine the retention time that NetWorker assigns to the original save set and clone save set.

- **Retention policy** attribute that is defined for the Client resource.
- **Retention policy** attribute that is defined for the Action resource that created the save set.
- **Retention policy** attribute that is defined for the Pool resource that contains the save set.

Note: This read-only attribute appears on the **Configuration** tab of the Pool resource, when **Diagnostic mode** is enabled in the **NetWorker**

Administration window. This is a 8.2.x and earlier attribute, which you cannot modify.

It is recommended that you define the retention policy for data in the Action resource. If you define the retention policy for save sets in multiple resources, you might experience unexpected save set expirations.

Cloning requirements and considerations

Review this section before you configure a clone action or perform a manual clone operation.

Note: The Clone Data Domain must be running on an operating system version that is similar to or later than that of Back up Data Domain operating system.

Device requirements

NetWorker requires two or more storage devices to perform a clone operation. One device contains the volume with the original data and one device contains the volume to which NetWorker writes the clone data. The clone data must reside on a volume that differs from the original volume. Each clone volume can only contain one instance of a cloned save set, even if the clone operation did not complete successfully. For example, if you want to create three clone copies of a save set, NetWorker must write each clone save set to a separate volume. As a result, you would need three separate volumes.

When using a tape library with multiple devices, the NetWorker Server automatically mounts the volumes that are required to complete the clone operation. When you use standalone tape devices, you must manually mount the volumes. A message in the **Alert** tab of the **Monitoring** window indicates which volumes to mount.

Often businesses choose devices for the initial backup that is based on speed or cost requirements. NetWorker supports cloning or staging data to a device type that differs from the source data volume. A common cloning or staging scenario includes using an AFTD for the initial backup to gain speed and versatility benefits, then to clone or stage the data to tape devices or deduplication devices. This scenario allows for an extended retention period without increasing disk space requirements. The use of deduplication can also provide efficient use of storage. Cloning to or from deduplication devices can ensure that these devices are used effectively. If the clone operation includes save sets from different devices, and you want all the save sets to be written to the same volume, include only one volume in the clone target pool.

Note: It is recommended that you do not write NDMP and non-NDMP data to the same clone volume because the number of file marks and positioning on the device differs for both data types.

Cloning multiplexed backups

You can clone multiplexed save sets. NetWorker writes the clone copies of multiplexed save sets as a single contiguous data stream on the target media (demultiplexed). When you recover from a multiplexed save set, read and recovery times increase as a result of the time NetWorker spends reading and locating the data. The process of demultiplexing save sets by the clone operation allowed you to read and recover data faster from a clone save set than a backup save set.

When you clone multiplex save sets, you can only clone one save set to the same target volume simultaneously. However, if the save sets have separate target volumes, you can start multiple clone sessions simultaneously from the same source.

Save set spanning

Some devices, for example Data Domain, support save set spanning across multiple volumes. When NetWorker clones a save set, the clone copy might start on one volume but continue on one or more additional volumes.

When using devices that support save set spanning, ensure that you:

- Identify save sets that span multiple volumes.
- Keep the number of continued save sets to a minimum.
- Use separate pools and larger or alternative devices.
- Use the Data Domain backup-to-disk and optimized cloning feature with Data Domain devices.
- Plan ahead to ensure that the volumes are available and that they are read in the best sequence.

Note: You can create a custom, scripted solution that uses the `nsrclone` command to manage save set spanning.

Save set status

NetWorker does not clone save sets that are recyclable or eligible for recycling. If NetWorker encounters a save set that is not browseable, the save set is skipped and is not cloned. However, the clone status is successful.

Recovery scenarios

When you clone data, you provide the datazone with an alternative data recovery source, which helps to protect against media loss or corruption. However, if the media is located in one of the following locations, then the second copy of the data is still vulnerable to major disasters that can affect the entire site:

- On the same tape library as the original data volume.
- On a deduplication device within the same data center, in a Data Domain environment.
- In an onsite safe.

Sometimes, you may require more copies of a save set to ensure that all the recovery scenarios are accommodated while maintaining the expected return on investment. This requirement may not apply to all clients and all data, or be practical. However, consider the reasons for cloning to ensure that the cloning strategy meets requirements and expectations.

Changing the target device, or moving tapes to a second location after the cloning operation completes, can provide additional protection.

Retention considerations

A Retention policy value applies to every type of save set. The retention policy value determines the length of time that the data remains available for recovery in the NetWorker media database and the client file index. You can specify a retention policy value for the clone save set that differs from the value that is defined for the original save set. When the retention policy differs for the original and clone save set, you can expire the original save set and reclaim the space on the source volume but maintain the data on a clone volume for future recoveries.

Note: The retention setting impacts the amount of disk space that is required by the NetWorker Server. The recovery procedure is likely to be different if retention has expired. The retention setting should be equal to or greater than the client or data requirements, and allow for the expected recovery conditions.


Cloning example

In this example, a backup of a client with three data drives creates three save sets. These save sets are stored on a volume that is accessible through Storage Node A. Once a cloning action occurs, the copies of these save sets are sent to eligible devices in the clone pool on Storage Node B.

In this figure:

- A client performs a backup of three data drives to Storage Node A. NetWorker creates three save sets, one save set for each data drive.
- A clone operation reads the data from the volumes on Storage Node A, and then copies the save sets to Storage Node B.

Figure 48 Cloning example



GEN-001749

Cloning with tape devices

There are a number of reasons why tape devices are used as part of the cloning process:

- In cases where tape is used as a secondary storage tier where selected data is cloned to tape for offsite storage or for extended data retention periods. This allows disk devices to be used for the initial backup where their speed and flexibility can be most effectively used for fast backup and recovery performance.
- In cases where tape is used as the primary backup media, there are still benefits in creating clone copies, including:
 - Secondary copy at different location or for offsite storage.

- Data validation.
- Verification of the ability to read data from the media.
- Added protection of multiple copies across multiple volumes.
- De-multiplexing of multiplex backups for faster recovery.

Cloning with tape devices provides two benefits which should be considered for every clone:

- Unlike disk-based devices, tape devices read data in a serial format. This means that while multiplexing is beneficial from a backup streaming perspective, this is not the case for recovery.
- If recovery speed is important, the use of clone copies as the source is likely to result in faster recovery throughput.

Tape clone copies are often the preferred method to read data in a disaster recovery situation. The ability to acquire, install, and configure a tape unit to read data is often the first task on a disaster recovery plan.

By creating a copy of the backup on tape, you can eliminate the need for appliances such as VTLs or disk systems to be in place. This often takes longer to acquire, install, and configure. However, ensure that the tape copy is a full and complete copy, without the dependence on other backups or deduplication appliances to complete the restore operation.

Production storage node cloning of data to physical tape

This section outlines the advantages and disadvantages of cloning data to physical tapes:

- The NetWorker software can clone from virtual tape in the disk library through a production storage node to a SAN-attached tape library to produce copies of save sets. This operation is a standard NetWorker cloning procedure.
- For the disk library, a virtual tape drive works in conjunction with a SAN-attached target tape device to complete the cloning process.
- Cloning from a production storage node to a second storage node can also be performed over IP.

 **Note:** Do not use a production storage node to perform cloning operations when the embedded storage node cloning capability is present.

Advantages

The advantages of cloning data to physical tapes include the following:

- Cloning can occur with the disk libraries under NetWorker control with standard NetWorker policy support. You can use multiple retention policies for different cloned copies of the data.
 - Cloning can occur at the save set and volume level.
-  **Note:** NetWorker can clone a single save set, multiple save sets or all of the save sets on a volume.
- Copying can occur from one tape type (virtual) to another tape type (target tape library), also known as tape conversion.
 - Copying can occur from multiple virtual tapes to a single tape, also known as tape stacking.

Disadvantages

The disadvantages of cloning data to physical tapes include the following:

- Requires storage node licenses.
- Requires maintenance of front-end SAN infrastructure to a target tape library as well as the virtual tape library.
- Consumes SAN bandwidth as data must be from virtual tape over the SAN to a target device on the SAN.

Cloning with file type and AFTD devices

Disk backup devices such as file type devices and advanced file type devices (AFTD) are ideal for cloning operations because they provide high speed, random access, and flexibility.

There are differences in the cloning process for file type devices and advanced file type devices.

- For file type devices, scheduled and manual cloning begins only after all save sets in a group have been backed up.
- For AFTDs, scheduled cloning begins only after all save sets in a group have been backed up. However, you can begin manually cloning a save set when it has finished its backup. For example, if there are three save sets (A, B, and C) in a backup, you can begin manually cloning Save Set A after its backup is complete and while the backups of Save Sets B and C are in progress. You can only manually clone one save set at a time. AFTDs allow recoveries during cloning operations (Read(source) or Write(target)). This assumes that the recover operation is not from the active save set and that only one clone operation is running at a time.

Often, the disk devices are used as the initial target device for backups, especially in situations where slower clients are unable to match the speeds that are expected for modern tape devices. In these situations, the ability to clone or stage data to tape often provides extended retention and data protection, while maximizing the disk use and benefits.

Data can remain on the disk devices for short periods, typically 3 to 14 days, which allows for:

- Adequate time for immediate and urgent restore operations to occur.
- Plenty of time to create further copies to tape or other disk-based devices for longer term retention.

Cloning with Avamar

When you configure NetWorker with Avamar to deduplicate backup data, the backup data is stored on an Avamar deduplication node on the Avamar server. The metadata (hash information) is stored on a NetWorker storage node.

 **Note:** NetWorker does not support the protection of new Avamar clients. You can only protect Avamar 7.2 clients that were initially configured on a NetWorker 8.2.x and earlier host.

To clone Avamar deduplication backups:

- Configure a clone action to clone the metadata. Cloning this hash metadata is highly recommended.
- Configure replication of the backup data from the original Avamar deduplication node to another Avamar deduplication node. The NetWorker software does not

start replication. A replication host (an Avamar server) must be configured by Customer Service before a deduplication backup can be replicated.

NOTICE For disaster recovery, you must replicate the client data to another Avamar deduplication node and clone the metadata. Both the metadata and the client data are required to recover client backup data.

You can also output the backup data of Avamar deduplication nodes to tape volumes. Create a second Client resource for the client, but do not configure the second instance as a deduplication client. Configure a data protection policy to back up the second client instance as a normal NetWorker client and store the backups on tape.

Cloning with Data Domain (DD Boost)

As with other NetWorker devices, you can use Data Domain device types to perform clone operations. You can clone single save sets or the entire Data Domain volume from a Data Domain device. You can also use the Data Domain device as the target device, to receive cloned data.

Cloning works differently for deduplication devices. You can perform clone-controlled replication (CCR), or optimized cloning of data, from one Data Domain system to another. Or you can clone data from a Data Domain device to tape or to any other device type.

Note: To use Data Domain with NetWorker, the NetWorker server hostname should be in lower case. Data Domain functions with lowercase and DD Cloud tier operations fails if it is mixed case.

Controlling storage node selection for cloning

You can control the storage node from which clone data is read (read source) and the storage node to which the clone data is written (write source). If you do not specify the read and write source storage nodes, then the cloning operation uses the `nsrserverhost` or NetWorker server as the storage node.

When you use data protection policies to clone, the selections that you make from **Source Storage Node** and **Destination Storage Node** lists on the **Clone Options** page for the clone action control the read source and write source. The "Creating a clone action" topic provides more information about how to configure a clone action and configure the filters that enable you to define the criteria that NetWorker uses to create the list of eligible save sets to clone.

When you use the `nsrclone` command, Use the `-J recover storage node` option to specify the read source host for the original volume and the `-d save storage node` option to specify the write source for the clone volume. The *NetWorker Command Reference Guide* or the UNIX man pages provide more information about the `nsrclone` command.

Determining the storage node for reading clone data

When you do not specify the source storage node for a clone action in a data protection policy or for the `nsrclone` command, the storage node from which clone data is read (read source) depends on whether the source volume is mounted or unmounted, as well as environment variable settings.

To control the storage node from which clone data is read, ensure that the source volume is mounted on the device for the storage node, or list the storage node in the **Recover storage nodes** attribute of the Client resource for the NetWorker server and in the **Read Hostname** attribute for the Library resource, if the source volume is in a

media library. Select **View > Diagnostic Mode** in the Administration interface to access the **Recover storage nodes** and **Read Hostname** attributes in the **Client Properties** dialog box.

NOTICE If the clone source volume is on a remote storage node and is unmounted, a volume clone operation cannot complete successfully, even if the source volume is mounted after the clone operation tries to start. The `nsrclone` program is unavailable with a message that the server is busy. This issue does not occur when the storage node is on the NetWorker server (or, not remote) or when you perform a clone controlled replication (optimized clone) operation.

Cloning operation logic for selecting a read source storage node

The cloning operation uses the following logic to determine the read source storage node:

1. If the source volume is mounted, then the storage node of the device on which the volume is mounted is used as the read source except in the following scenarios:

- If the `FORCE_REC_AFFINITY` environment variable is set to `Yes`.
- If the volume resides in a Virtual Tape Library (VTL) environment such as a CLARiiON Disk Library (CDL).

In these scenarios, the NetWorker software ignores whether the source volume is mounted and behaves as though the volume is not mounted.

2. If the source volume is not mounted or the `FORCE_REC_AFFINITY` environment variable is set to `Yes`, then the NetWorker software creates a list of eligible storage nodes, based on the storage nodes that meet both of the following criteria:

- The storage node is listed in the **Recover storage nodes** attribute of the Client resource for the NetWorker server.

If there are no storage nodes in the list and the **Autoselect storage node** checkbox in the NetWorker server Client resource is clear, then the clone operation uses the value in the **Storage Nodes** attribute for the NetWorker server Client resource.

If there are no storage nodes in the list and the **Autoselect storage node** checkbox in the NetWorker server Client resource is selected, then the clone operation uses autoselect logic to choose the storage node.

- If the requested volume is in a media library, then the storage node is listed in the **Read Hostname** attribute for the Library resource is used.

If the **Read Hostname** attribute for the Library resource is not set, then all storage nodes on which any device in the library is configured are added to the list of eligible storage nodes.

Note: If the volume is not in a media library, then the list of storage nodes is based only on the criterion for storage node settings in the NetWorker server Client resource.

Example

Consider the following example for a volume that resides in a media library and is not mounted:

- The **Recover storage nodes** attribute in the NetWorker server Client resource lists the following storage nodes in order:
 - Storage node F

- Storage node E
- Storage node D
- The **Read Hostname** attribute for the Library resource is not set, but the following devices in the media library are configured with storage nodes:
 - Device A is configured on storage node D.
 - Device B is configured on storage node E.
 - Device C is configured on storage node B.

The list of eligible storage nodes is the intersection of the two previous lists (storage nodes E and D). The order in which the storage node is selected depends on the order of the storage nodes in the **Recover storage node** attribute list. In this example, storage node E is selected first as the read source storage node. If storage node E is not available, then storage node D is selected.

If no matching storage nodes are found in the intersecting list, then an error is written to the daemon log file that indicates that no matching devices are available for the operation. To correct the problem, ensure that at least one matching storage node appears in both lists.

Determining the storage node for writing cloned data

When you do not specify the destination storage node for a clone action in a data protection policy or for the `nsrclone` command, the storage node to which clone data is written (write source) depends on the storage nodes listed in the **Clone storage nodes** attribute for the read source storage node or the NetWorker server storage node.

To control the storage node to which clone data is written, specify the hostname of the write source storage node in the **Clone storage nodes** attribute for the read source storage node.

To clone from many read source storage nodes to a single write source storage node, specify the hostname for the write source storage node in the **Clone storage nodes** attribute for the NetWorker server storage node.

In backup-to-disk environments, a single backup volume can be shared by multiple storage devices on different storage nodes. To ensure unambiguous clone write sources in this situation, specify the same write source storage node in the **Clone storage nodes** attribute of all storage nodes that have access to the backup volume.

Regardless of where the cloned data is written, the client file index and media database entries for the cloned save sets reside on the NetWorker server.

Cloning operation logic for selecting a write source storage node

The cloning operation uses the following logic to determine the storage node that stores cloned backup data:

1. The write source storage node is listed in the **Clone storage nodes** attribute for the read source storage node.
2. If the **Clone storage nodes** attribute for the read source storage node is blank, then the write source storage node is listed in the **Clone storage nodes** attribute for the NetWorker server storage node.
3. If the **Clone storage nodes** attribute for the NetWorker server storage node is blank, then the write source storage node depends on whether the **Autoselect storage node** checkbox is selected or clear in the Client resource for the NetWorker server:

- If the checkbox is clear, then the clone operation uses the value in the **Storage Nodes** attribute of the Client resource for the NetWorker server.
- If the checkbox is selected, then the clone operation uses autoselect logic to choose the storage node.

You must select **View > Diagnostic Mode** in the Administration interface to access the **Autoselect storage node** attribute in the **Client Properties** dialog box.

Determining the storage node for recovering cloned data

The storage node from which clone data is recovered depends on whether the source volume is mounted or unmounted, as well as environment variable settings.

To control the storage node from which cloned data is recovered, ensure that the source volume is mounted on the device for the storage node. Alternatively, list the storage node in the **Recover storage nodes** attribute of the Client resource that is being recovered and in the **Read Hostname** attribute for the Library resource, if the source volume is in a media library. You must select **View > Diagnostic Mode** in the Administration interface to access the **Recover storage nodes** and **Read Hostname** attributes in the **Client Properties** dialog box.

Recovery operation logic for selecting the storage node from which to recover cloned data

The recovery operation uses the following logic to determine the storage node from which to recover cloned data:

1. If the source volume is mounted, then the storage node of the device on which the volume is mounted is used as the read source except in the following scenarios:
 - If the `FORCE_REC_AFFINITY` environment variable is set to `Yes`.
 - In a Virtual Tape Library (VTL) environment such as a CLARiiON Disk Library (CDL).

In these scenarios, the NetWorker software ignores whether the source volume is mounted and behaves as though the volume is not mounted.

2. If the source volume is not mounted, or the `FORCE_REC_AFFINITY` environment variable is set to `Yes`, then the NetWorker software creates a list of eligible storage nodes, based on the following criteria:

- The storage node is listed in the **Recover storage nodes** attribute of the NetWorker Client resource that is being recovered.

If there are no storage nodes in the list and the **Autoselect storage node** checkbox in the Client resource is clear, then the clone operation uses the value in the **Storage Nodes** attribute for the Client resource.

If there are no storage nodes in the list and the **Autoselect storage node** checkbox in the Client resource is selected, then the recovery operation uses autoselect logic to choose the storage node.

- If the requested volume is in a media library, then the storage node is listed in the **Read Hostname** attribute for the Library resource is used.

If the **Read Hostname** attribute for the Library resource is not set, then all storage nodes on which any device in the library is configured are added to the list of eligible storage nodes.

Note: If the volume is not in a media library, then the list of storage nodes is based only on the criterion for storage node settings in the NetWorker server Client resource.

Recover Pipe to Save

Recover Pipe to Save (RPS) is a feature that improves performance by allowing clone, backup, and recovery operations to access the same `nsrmmmd` process concurrently.

About this task

To clear or select the **Disable (RPS) Clone** attribute, perform the following steps.

Procedure

1. On the **Administration** window, click **Server**.
2. In the left pane of the **Server** window, right-click the NetWorker server.
3. From the **File** menu, select **Properties**.
4. Select the **Configuration** tab.
5. Clear or select the **Disable RPS Clone** attribute, and then click **OK**.

Note:

- Cloning of saveset is not supported from a non-Data Domain device to a DD Cloud Tier device. Cloning fails with the following error:

```
Failed to get mmd reservation with err: Clone saveset(s)
operation from a non-Data Domain device to a DD Cloud Tier device
is not supported.
```

- When you perform a fresh installation of NetWorker Server 19.1, RPS cloning is disabled by default.
- If you have NetWorker 19.1 server installed, Dell EMC recommends that you maintain parity with the NetWorker 19.1 storage node in both the RPS Enabled and Disabled mode. However, if you have compatibility and migration challenges and want to maintain earlier versions of the storage node, that is, the N-2 version, then Dell EMC recommends that you use the RPS Disabled mode for cloning. In cases where VMware vProxy save sets are used for cloning, RPS Enabled mode is supported by default. Therefore, Dell EMC recommends that you maintain NetWorker server and storage node compatibility and configure clone actions with the appropriate storage node based on the workloads.

Cloning save sets from the command prompt

Use the `nsrclone` command to clone save sets and volumes from a command prompt, or to script clone operations.

Script clone operations for any of the following scenarios:

- To control the conditions before cloning occurs. For example, following a specific event or test, or as part of a workflow.
- To control the actions after cloning has been successful. For example, deleting files, or moving data as part of a workflow.
- To control the cloning as part of an enterprise management scheduler that is independent of NetWorker scheduling or NMC.

- To create multiple clones. For example, clone 1 on disk, clone 2 to tape, each with specific dependencies, timing, and logic.

Note: When using the scripted cloning feature, use the latest versions of NetWorker software. This minimizes the complexity of the logic in the cloning script.

The `nsrclone` command requires specific privileges which are assigned based on session authentication. NetWorker supports two types of session authentication. Token-based authentication, which requires you to run the `nsrlogin` before you run the command and authenticates the user that runs the command against entries that are defined in the External Roles attribute of a User Group resource. Classic authentication, which is based on user and host information and uses the user attribute of a User Group resource to authenticate a user. Classic authentication does not require an authentication token to run the command. For example, if you run the command without first running `nsrlogin`, NetWorker assigns the privileges to the user based on the entries that are specified in the Users attribute of the User Group resource. When you use `nsrlogin` to log in as a NetWorker Authentication Service user, NetWorker assigns the privileges to the user based on the entries that are specified in the External Roles attributes of the user Group resource. The *NetWorker Security Configuration Guide* provides more information about privileges [Using nsrlogin for authentication and authorization](#) provides more information.

Mounting clone source volumes on remote storage nodes

When the volume that contains the original data resides on a storage node that is not the NetWorker server, mount the source volume in a device on the storage node before you try the clone operation.

NetWorker displays the following error message in the `daemon.raw` and the **Logs** window in the **NetWorker Administration** window when the source volume is not mounted before the clone operation:

```
Server server_name busy, wait 30 second and retry
```

Cloning volumes from the command prompt

Volume cloning is the process of reproducing complete save sets from a storage volume to a clone volume. Use the `nsrclone` command to clone save set data from backup or archive volumes.

Procedure

1. Optionally, use the `nsrlogin` command to authenticate a user and generate a token for the [Using nsrlogin for authentication and authorization](#) provides more information.
2. Use the `mminfo` command or the **NetWorker Administration** window to determine the name of the volume that contains the save sets that you want to clone.
 - To use the NetWorker **Administration** window, perform the following steps:
 - a. Click **Media**.
 - b. In the expanded left pane, select either **Disk Volumes** or **Tape Volumes**.
 - c. In the right pane, record the name that appears in the **Volume Name** column.
 - To use the `mminfo` command to display volumes. For example, to display a list of all the available volumes, type the following command:

```
mminfo -mv
state volume written (%) expires read mounts capacity
```

```

valid next type
bu_iddnwserver2.iddlab.local.001 46 MB 100% 04/04/2015 0
KB 0 0 KB 16193908 0 adv_file
bu_iddnwserver2.iddlab.local_c.002 0 KB 0% undef 0 KB 0
0 KB 4294384030 0 adv_file

```

3. From a command prompt on the NetWorker server, use the `nsrclone` command to clone the save sets on a volume. For example to clone save sets to volume in the default clone pool, type:

```
nsrclone -v -b Default backup.001
```

where:

- *backup.001* is the name of the volume that contains the source save sets.
- The clone pool that the clone operation uses to write the clone save sets is the Default clone pool.

Cloning save sets from the command prompt with filters

You can use the `nsrclone` command on the NetWorker server to manually clone save sets, based on user defined criteria or identifiers. Use the `mminfo` command to determine which identifiers you want to use to define a list of save sets to clone. Identifiers include the client name, the save set name, the backup level, and the number of valid copies or clones not yet created in the target pool.

UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrclone` and the `mminfo` commands.

The following examples illustrate the `nsrclone` command:

 **Note:** Optionally, use the `nsrlogin` command to authenticate a user and generate a token for the `nsrclone` and `mminfo` commands. [Using nsrlogin for authentication and authorization](#) provides more information.

- To clone all save sets created in the last 24 hours for clients mars and jupiter with save set names `/data1` and `/data2` for only backup level full, type:

```
nsrclone -S -e now -c mars -c jupiter -N /data1 -N /data2 -l full
```

- To clone all save sets that were not copied to the default clone pool in a previous partially aborted `nsrclone` session, type:

```
nsrclone -S -e now -C 1
```

- To clone all save sets that were not copied to the default clone pool in a previous partially aborted `nsrclone` session and then assign the save sets a retention policy value that differs from the original save set, type:

```
nsrclone -S -e now -C 1 -y 12/12/2016
```

Using query filter (-q) to specify savesets

You can clone specific savesets that are a part of policy, workflow or action by using filters to the `nsrclone` command. You can provide any combination of the 3 new query options.

- Standalone query options:

```
nsrclone -q "policy=Gold" -S -t "1 month ago"
```

```
nsrclone -q "workflow=NMC server backup" -S -t "1 month ago"
```

```
nsrclone -q "action=backup" -S -t "1 month ago"
```

- Combination of 2 options:

```
nsrclone -q "action=backup" -q "policy=Gold" -S -t "1 month ago"
```

```
nsrclone -q "workflow=NMC server backup" -q "policy=Gold" -S -t "1 month ago"
```

```
nsrclone -q "workflow=NMC server backup" -q "action=backup" -S -t "1 month ago"
```

- Combination of 3 options:

```
nsrclone -q "action=backup" -q "policy=Gold" -q "workflow=NMC server backup" -S -t "1 month ago"
```

User can add pre-existing group query options along with the newly added query options. For example, -q group=<>

```
nsrclone -q "group=DPG" -q "action=backup" -q "policy=Gold" -q "workflow=NMC server backup" -S -t "1 month ago"
```

The following table provides the descriptions of the options that are used in the `nsrclone` command example.

Table 81 List of `nsrclone` options and their descriptions

Options	Description
-S	Specifies that the subsequent <code>nsrclone</code> options are save set identifiers and not volumes names.
-C less_than_clone_copies_value	Specifies the upper non-inclusive integer limit such that only save sets with a lesser number of clone copies in the target clone pool are considered when <code>nsrclone</code> searches for save sets to clone. Use this option when you retry an aborted clone operation. Note: NetWorker does not consider the following save sets when calculating the copies value for a save set: - Original save set. - AFTD read-only mirror clone. NetWorker counts the read or write master clone only because there is only one physical clone copy between the related clone pair. - Recyclable, aborted, incomplete, and unusable clone save sets. Requires the -t or -e option.
-l level_or_range_value	Specifies the backup level to search for when <code>nsrclone</code> determines which save sets to clone: - Manual—For ad-hoc or client-initiated save sets. - full—For level full save sets. - incr—For level incremental save sets.

Table 81 List of nsrclone options and their descriptions (continued)

Options	Description
	You can specify more than one level by using multiple <code>-l</code> options. Requires the <code>-t</code> or <code>-e</code> option.
<code>-N save_set_name</code>	Specifies the save set name to search for when <code>nsrclone</code> determines which save sets to clone. Use multiple <code>-N</code> options to specify more than one save set name. Requires the <code>-t</code> or <code>-e</code> option.
<code>-c client_name</code>	Specifies the name of the client to search for when <code>nsrclone</code> determines which save sets to clone. Use multiple <code>-c</code> options to specify more than one client name. Requires the <code>-t</code> or <code>-e</code> option.
<code>-y date</code>	Specifies the retention policy date to assign to the clone save set. Use a time and date format that is accepted by the <code>nsr_getdate</code> program. UNIX man page and the <i>NetWorker Command Reference Guide</i> provides detailed information about <code>nsr_getdate</code>.

Updates to group, policy, workflow, and action fields in mediadb for a cloned save set

When a cloned save set is queried for group, policy, workflow or action fields in mediadb, output of the query can be categorized based on how cloning was carried out (whether it was run through action or command line).

Action or policy clones- For concurrent, sequential, and scheduled clone types, the cloned copy displays group, policy, workflow, and action for policy tree under which clone action is created. For concurrent and sequential clones these fields are same as backup copy.

Command line clones- Policy, workflow, and action fields are empty and only the group which is inherited from backup is populated.

Staging save sets

Staging is the process of transferring backup data from one storage device, usually an AFTD to another device, and then removing the data from the original device. Staging

save sets from the primary data device ensures that there is always sufficient disk space available on the primary device to store data.

To manage the staging process, manually stage individual save sets from a command prompt or you can configure a Staging resource that automatically stages the data. The Staging resource defines the criteria that the stage process uses to determine when the data device requires data staging and which save sets are eligible to stage and in what order.

Based on the configuration of the Staging resource, the staging process performs the following high level activities:

1. Performs file system checks at an interval that is defined in the **File system check interval** attribute to determine:
 - If the percentage of used disk space on the source device exceeds the value that is defined in the **High water mark** attribute of the Staging resource.
 - If the length of time that the save sets have resided on the disk exceeds the value that is defined in the **Max storage period** attribute of the Staging resource.
2. Creates a list of save sets on the source device that are eligible to move to a destination device.
3. Clones the eligible save sets from the source device to the destination device, and then updates the media database with information about the save sets on the destination device. The save set on the destination device retain the same attributes values, for example retention policy, as the original save set.
4. Removes the original save sets from the source device, recovers disk space on the source volume for staged save sets, and then removes information about the original save sets from the media database.

Note: When the staging process encounters an error after successfully cloning some save sets, the staging process only removes successfully staged save sets from the source volume before the process ends. Only a single set of save sets will exist on either the source or destination volumes after staging.

Staging data allows you to accommodate multiple service levels. You can configure a staging policy that keep the most recent backups on one storage device for fast recovery and move other backups with less demanding recovery requirements to more cost-effective slower storage. For example, you can store the initial backup data on a high performance file type or advanced file type device to reduce backup time. At a later time, outside of the normal backup period, you use the staging process to move the data to a less expensive but more permanent storage medium, such as magnetic tape. After the backup data moves to the other storage medium, NetWorker deletes the backup data from the file or advanced file type device so that sufficient disk space is available for the next backup. Staging does not affect the retention policy of backup data and the staged data is still available for recovery on the destination device.

You can stage a save set from one disk to another as many times as required. For example, you could stage a save set from disk 1 to disk 2 to disk 3, and finally to a remote tape device. When the save set is staged to a tape, it cannot be staged again. However, you could still clone the tape.

Staging bootstrap backups

You can direct bootstrap backups to a disk device such as an AFTD or FTD device.

However, if you stage a bootstrap backup to a volume on another device, NetWorker reports the staging operation as complete although the “recover space” operation has not started, and the bootstrap remains on the original device. Therefore, if the staged bootstrap is accidentally deleted, you can recover the bootstrap from the original disk.

The *NetWorker Server Disaster Recovery and Availability Best Practices Guide* describes how to recover a bootstrap from the original disk.

Also, if the bootstrap data is not staged from the original disk, the data on the original disk is subject to the same retention policies as any other save set backup and is, therefore, deleted after the retention policy has expired.

Creating a staging resource

To prevent an AFTD from becoming full, configure a Staging resource to automatically move save sets to another medium and make disk space available. The Staging resource defines when NetWorker starts the stage or reclaim disk space operation on the source device, and the criteria that NetWorker uses to determine which data to stage.

Procedure

1. In the **Administration** window, click **Devices**.
2. In the left navigation pane, select **Staging**.
3. From the **File** menu, select **New**.

The **Create Staging** dialog box appears, starting with the **General** tab.

4. In the **Name** box, type a name for the staging policy.
5. In the **Comment** attribute, type a description for the staging policy.
6. In the **Enabled** attribute, select **Yes** to enable the staging policy or **No** to disable the staging policy.

When you select **Yes**, NetWorker automatically starts the staging policy, based on the configuration settings that you define.

7. In the **Devices** attribute, select the check boxes next to each source device from which you want to stage data.

You can assign multiple devices to a single staging policy. However, you cannot assign a single device to multiple staging policies.

8. From the **Destination Pool** list, select the destination clone pool that contains the volumes to which NetWorker stages the data.

If you select the clone pool that only uses remote storage node devices, you must also modify **Clone storage nodes** setting on the **Configuration** tab of the storage node resource for the NetWorker server to include the storage node name. [Determining the storage node for writing cloned data](#) on page 503 Provides details.

9. In the **Configuration** group box, specify the criteria that starts the staging policy.

The following table summarizes the available criteria that you can define for the staging policy.

Table 82 Staging criteria options

Option	Configuration steps
High water mark (%) Low water mark (%)	Use these options to start the stage policy based on the amount of used disk space on the file system partition on the source device. You must define a value higher than the value defined in the Low water mark (%) attribute.

Table 82 Staging criteria options (continued)

Option	Configuration steps
	High water mark (%)—Defines the upper used disk space limit. When the percentage of used disk space reaches the value that is defined in the High water mark (%) attribute, NetWorker starts the stage operation to move save sets from the source disk. Low water mark (%)—Defines the lower used disk space limit. When the percentage of used disk space reaches the value that is defined in the Lower water mark (%) attribute, NetWorker stops moving save sets from the source disk. Note: When staging and backup operations occur concurrently on the source disk device, NetWorker does not accurately display the disk volume usage total in the Written column in output of the <code>mminfo -mv</code> command or in the Used column on the Media window of the NetWorker Administration application.
Save set selection	Use this option to rank the order in which NetWorker stages the save sets, based on save set size or age. Available values include: • largest save set—Stage the save sets in order of largest save set size to smallest save set size. • oldest save set—Stage the save sets in order of oldest save set to most recent save set. • smallest save set—Stage the save sets in order of smallest save set size to largest save set size. • youngest save set—Stage the save sets in order of most recent save set to least recent save set.
Max storage period Max storage period unit	Use this option to start the stage operation based on the amount of time that a save set has resided on the volume. • Max storage period—Defines the number of hours or days that a save set can reside on a volume before the stage process considers the save eligible to move to a different volume. • Max storage period unit—Defines the unit of measurement for the value in the max storage period attribute. Available values are Hours and Days. The maximum storage period setting is used along with the file system check interval. Once the maximum storage period is reached, staging does not begin until the next file system check.
Recover space interval Recover space unit	Use this option to determine when the stage operation removes the successfully staged save set from the source volume. • Recover space interval—Defines the frequency in which NetWorker starts of the recover space operation, which removes successfully stage data from the source volume.

Table 82 Staging criteria options (continued)

Option	Configuration steps
	• Recover space unit—Defines the unit of measurement for the value in the recover space interval attribute. Available values are Hours and Days.
File system check interval	Use this option to define when NetWorker automatically starts the staging process. • File System Check Interval—Defines the frequency in which NetWorker starts the staging process. At every file system check interval, if either the high water mark or the maximum storage period has been reached, then staging begins. • File system check unit—Defines the unit of measurement for the value in the file system check interval attribute. Available values are Hours and Days.

10. Optionally, to start the staging process immediately:
 - a. Select the **Operations** tab.
 - b. From the **Start Now** list, select the component of the staging process to perform immediately, for all source devices that are assigned to the staging policy:
 - **Recover space**—To recover space for save sets with no entries in the media database and to delete all recycled save sets.
 - Select **Check file system**—To check the file system and stage eligible save set data to a destination volume.
 - Select **Stage all save sets**—To stage all save sets to a volume in the destination pool.

After the staging operation is complete, this option returns to the default setting (blank).

11. Click **OK**.

Editing staging configurations

You can edit all settings for a Staging resource except for the name of the resource. To edit the name of a resource, first delete the resource, and then re-create the resource with the new name.

Procedure

1. In the **Administration** window, click **Devices**.
2. In the right pane, perform one of the following tasks:
 - To modify multiple attributes in a single configuration resource by using the **Staging** window, right-click the staging configuration and select **Properties**.
 - To modify a specific attribute that appears in the resource window, place the mouse in the cell that contains the attribute that you want to change, then right-click. The menu displays an option to edit the attribute. For example, to modify the **Comment** attribute, right-click the resource in the **Comment** cell and select **Edit Comment**.

 **Note:** To modify a specific attribute for multiple resources, press and hold the **Ctrl** key, select each resource, and then right-click in the cell that contains the attribute that you want to change. The menu displays an option to edit the attribute.

3. In the left pane, select **Staging**.
4. Click **OK**.

Copying a Staging resource

Procedure

1. In the **Administration** window, click **Devices**.
2. In the left pane, select **Staging**.
3. In the right pane, right-click the staging policy, and select **Copy**.
The **Create Staging** dialog box appears with the same settings as the original staging policy except for the name.
4. Type the name for the new staging policy in the **Name** box.
5. Select the checkboxes next to the source devices for the staging policy in the **Devices** list.

You can assign multiple devices to a single staging policy. However, you cannot assign a single device to multiple staging policies.

6. Edit other settings for the staging policy as necessary.
7. Click **OK**.

Deleting a staging policy

You can delete any staging policy except for the default staging policy. Disable the default staging policy if you do not want to perform staging.

Procedure

1. In the **Administration** window, click **Devices**.
2. In the left pane, select **Staging**.
3. Remove all devices from the staging policy:
 - a. In the right pane, right-click the staging policy, and select **Properties**.
 - b. Clear the checkboxes next to all the devices in the **Devices** list.
 - c. Click **OK**.
4. In the right pane, right-click the staging policy, and select **Delete**.
A confirmation message appears.
5. Click **Yes**.

Manual staging from the command prompt

Use the `nsrstage` command to stage individual save sets from a command prompt.

The `nsrstage` command requires specific privileges which are assigned based on session authentication. NetWorker supports two types of session authentication. Token-based authentication, which requires you to run the `nsrlogin` before you run the command and authenticates the user that runs the command against entries that are defined in the External Roles attribute of a User Group resource. Classic

authentication, which is based on user and host information and uses the user attribute of a User Group resource to authenticate a user. Classic authentication does not require an authentication token to run the command. For example, if you run the command without first running `nsrlogin`, NetWorker assigns the privileges to the user based on the entries that are specified in the Users attribute of the User Group resource. When you use `nsrlogin` to log in as a NetWorker Authentication Service user, NetWorker assigns the privileges to the user based on the entries that are specified in the External Roles attributes of the user Group resource. The *NetWorker Security Configuration Guide* provides more information about privileges, and [Using nsrlogin for authentication and authorization](#) provides more information.

Staging save sets from the command prompt

You can use the `nsrstage` command to stage save sets to another volume, based on the ssid.

About this task

If the save set has been cloned and you stage the save set from the command prompt, the cloned versions of the save set are removed when the original save set is removed. To keep the cloned save sets after you remove the original save set, specify a clone ID with the save set ID to indicate the source volume of the staging.

Procedure

1. Optionally, use the `nsrlogin` command to authenticate a user and generate a token for the `nsrstage` and `mminfo` commands. [Using nsrlogin for authentication and authorization](#) provides more information.
2. Use the `mminfo` command to determine the ssid and cloneid of a save set.

For example:

```
mminfo -avot -r "volume,ssid,cloneid,name"
```

3. Use the `nsrstage` command to migrate the save sets to another volume.

For example:

```
nsrstage -m -S ssid/cloneid
```

Note: When you do not use the `-b` option to specify a destination clone pool, the `nsrstage` command migrates the save sets to a volume in the Default Clone pool.

UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrstage` and `mminfo` commands.

Common NetWorker staging commands and issues

This section describes how to run common staging ptasks from the command prompt and how to resolve common staging issues.

How to migrate all save sets created by specific date

```
nsrstage -m -S 'mminfo -r ssid -q 'savetime>last saturday'
```

How to use the *-f* and *-d* option in the `nsrstage` command

```
mminfo -r ssid -q 'savetime>last saturday' >inputfile.txt
nsrstage -m -d -f inputfile.txt
```

How to recover space from volume by using `nsrstage` command

For example, to recover space from volume `volume.012`:

```
nsrstage -C -V volume.012
```

How to remove incomplete or aborted save sets that the staging process does not migrate

The stage operation does not move aborted or incomplete save sets to a tape device. To remove the save sets from the source device, perform the following steps:

1. Manually delete the save set from the media database by typing: `nsrmm -d -S ssid`
2. Remove the save sets from the source device by typing: `nsrstage -C -V volume`

How to resolve the '`nsrstage: device `(staging_volume)` is not enabled`' error

Staging fails with this error when either the source or destination device is not ready. The following error message might also appear:

```
Error: 'nsrd: media warning: (staging_volume) reading: Badfile
number'
```

When you see these errors, ensure the following:


- The source device is not in service mode.
- The destination tape device or jukebox is properly synchronized.

Archiving data

The archive process captures files or directories as they exist at a specific time, and writes the data to archive storage volumes, which are not automatically recycled. After the archive process completes, you can delete (*groom*) the original files from the disk to conserve space.

The client archive program (`nsrarchive`) creates an archive. The client `nsrexecd` service starts this archive.

The following figure illustrates how the NetWorker software archives data.

Figure 49 Overview of archive operation

where:

1. Client file systems
2. Backup data tracking structures
3. Data
4. Media database information
5. File index information

Archive save sets

Archive save sets are similar to backup save sets. The main difference is that there is no retention period for archive save sets, so the archive save sets never expire.

By default, the archive backup level is always set to full.

Licensing

You must purchase and license the archive feature separately from other NetWorker software components. The *NetWorker Licensing Guide* provides more information on licensing procedures.

Encryption of archive data

If the NetWorker client is set up for encryption with the aes ASM, then archive data is also encrypted.

Limitations

The following limitations apply to the archive feature:

- You cannot archive the `WINDOWS ROLES AND FEATURES` save set.
- The Client Direct feature does not support archiving.

Storage of archived data

To archive data, you must configure a device, either stand-alone or in an autochanger or silo, that is connected to a NetWorker server or storage node. If you are cloning archives, at least two devices must be available. Also, archive data must be written to archive pools instead of backup or clone pools.

The archive volume must be loaded and mounted in the server device to complete an archive operation.

Information about archive data is tracked in the media database for the NetWorker server.

Configuring pools to index archive data

The settings for the archive pool that is used to store archive data determine whether you index archive data.

When you index archive data, information about individual files in the archive save set is tracked in the client file index. The client file index entries that are generated during an archive are backed up to volumes from the default pool during the next scheduled backup. You can browse and recover individual files from indexed archive save sets. However, indexed archive data can result in a large client file index that never expires.

When you perform nonindexed archiving, entries are not added to the client file index. You must recover the entire save set instead of browsing to and recovering individual files.

Default archive pools

The following default pools are available for archived data:

- Indexed Archive pool
- PC Archive pool
- Archive pool

The Indexed Archive pool and the PC Archive pool support indexed archiving. The Archive pool supports nonindexed archiving.

You cannot change the settings for these default pools, although you can create custom archive pools.

If you do not specify a pool to store archived data, the NetWorker software uses the Indexed Archive pool by default.

Custom archive pools

You can create custom archive pools in the **Media** window of the Administration interface. The **Store index entries** checkbox on the **Configuration** tab for the media pool determines whether the archive data written to the volumes in the pool are indexed. Select the checkbox to perform indexed archiving, or clear the checkbox to perform nonindexed archiving.

Enabling archiving

After you license the archive service and type the enabler code in the NetWorker server, all clients for that server are enabled for the NetWorker archive feature by default. You can specify which clients and users have permission to archive data.

Before you begin

Ensure that the NetWorker Server is in diagnostic mode. To enable diagnostic mode, from the **View** menu, select **Diagnostic mode**.

Procedure

1. To control whether a client can archive data, in the **Client Properties** window, on the **Globals (2 of 2)** tab, perform one of the following actions on the **Archive services** box.
 - Clear the checkbox to disable archiving for the client.
 - Select the checkbox to enable archiving for the client.

You must select or clear the **Archive services** checkbox for all Client resources that are associated with the client. You might have multiple Client resources for

a single client. For example, if both the NetWorker module software and the NetWorker client software are installed on the same computer, there are multiple Client resources.

2. Add users that should have permission to perform archiving to the Archive Users user group in the **Server** window of the Administration interface.

The *NetWorker Security Configuration Guide* provides details.

Archiving data from Windows

You can manually archive data from a NetWorker client on Windows by using the NetWorker User program.

About this task

Note: Manual archives from a Windows client do not enforce global or local file (`nsr.dir`) directives. However, local directives (`networkr.cfg`) that are created with the NetWorker User program are enforced.

Procedure

1. In the NetWorker User program (`winworkr.exe`), click **Archive**.
The **Archive Options** dialog box appears.
2. Type a comment in the **Annotation** attribute.
The annotation uniquely identifies the archive save set during retrieval. Consider adopting a consistent naming convention so that you can easily identify archives, based on the annotation name.
3. From the **Archive Pool** list, select the archive pool for the data.
4. To clone each archive save set, select the **Clone** checkbox, and then select the destination archive clone pool from the **Clone Pool** list.
5. To check the integrity of the archive data on the storage volume, select the **Verify** checkbox.
6. To remove the archived files from the disk after archiving completes, select the **Groom** checkbox.
7. Click **OK**.
The **Archive** browse dialog box appears.
8. Select the checkbox next to the directories and files to archive, and clear the checkbox next to the directories and files that you do not want to archive.
9. From the **File** menu, select **Start Archive**.
The **Archive Status** dialog box displays the status of the archive process. When the archive process completes, a confirmation message appears if you selected the **Groom** checkbox.
10. Click **Yes** to continue with deletion of archived files from the local disk.

Archiving data from UNIX

To perform a manual archive from a UNIX client, use the `nsrchive` command. The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrchive` command.

Recovering archived data

The steps to recover archived data depend on the client operating system and whether the data is indexed.

Required permissions for archive recovery

To recover archive data, a user must be a member of the Archive Users user group and must have read permissions for the archive data.

The **Public archives** checkbox on the **Setup** tab of the **NetWorker Server Properties** dialog box controls whether all users with read permissions for the data can recover the data, or if only the user who owns the data can perform recovery. Select the checkbox to allow all users with read permissions to recover the data, or clear the checkbox to require users to own data that they want to recover.

 **Note:** The user that recovers archived data becomes the owner of the data. Some operating systems allow you to change the ownership of archived data to the original owner during the recovery.

Recovering indexed archive data from a Windows client

You can recover indexed archive data from a Windows client the same way that you recover backup or clone data.

About this task

Indexed archive data must be stored on a volume in one of the following pools:

- Indexed Archive pool
- PC Archive pool
- Custom archive pool with the **Store index entries** checkbox selected in the pool properties

Procedure

1. In the NetWorker User program (`winworkr.exe`), click **Recover**.
The **Source Client** dialog box appears.
2. Select the source client with the data to recover, and click **OK**.
The **Destination Client** dialog box appears.
3. Select the destination client for the recovered data, and click **OK**.
The **Recover** browse dialog box appears.
4. Select the checkbox next to the files and directories to recover.
5. Click **Start**.

Recovering nonindexed archive data from a Windows client

When you recover nonindexed archive data, you must recover the entire save set instead of individual directories and files.

About this task

Nonindexed archive data must be stored in the default Archive pool or in a custom archive pool with the **Store index entries** checkbox cleared in the pool properties.

You can recover nonindexed archive data either by using the Archive Retrieve feature or the Save Set Recover feature in the NetWorker User program (`winworkr.exe`).

[Performing a save set recover with NetWorker User](#) on page 572 provides details on save set recovery.

Procedure

1. Mount the archive volume in the storage device.
2. In the NetWorker User program, select **Operation > Archive Retrieve**.
The **Source Client** dialog box appears.
3. Select the source client with the data to recover, and click **OK**.
The **Archive Retrieve** dialog box appears.
4. In the **Annotation string** box, type all or part of the annotation string that you specified for the save set when it was archived.
Leave the box empty to view a list of all archived save sets for the client.
5. Click **OK**.
The **Save Sets** dialog box appears.
6. To view a list of volumes that are required to retrieve the data from this archived save set, click **Required Volumes**.
7. To type a new path for the location of the recovered data and to indicate what the NetWorker server should do when it encounters duplicate files, click **Recover Options**.
8. Select the archived save set to recover and click **OK**.
The **Retrieve Status** dialog box displays the status of the recovery.

Recovering archive data from a UNIX client

Use the `nsrretrieve` program to retrieve archive data for a UNIX client. You must specify the files or directories to recover, or recover the entire save set on a UNIX client. You cannot browse archive data on UNIX.

Procedure

1. Mount the archive volume in the storage device.
2. Open a command prompt, and type the `nsrretrieve` command using the following syntax:

```
nsrretrieve -s NetWorker_server -A annotation -S ssid/cloneid
-i{N|Y|R} path
```

where:

- *NetWorker_server* is the hostname of the NetWorker server.
- *-A annotation* specifies the annotation string for the archive save set. You must specify at least one annotation or save set ID.

Consider an example where archive A is annotated with `Accounting_Fed` and archive B is annotated with `Accounting_Local`. If you type `nsrretrieve -A Accounting`, then no match is found and the archive data is not recovered. If you type `nsrretrieve -A ting_L`, then the recovery process recovers the data from Archive B.

- *-S ssid/cloneid* specifies the archive save set to recover. To recover a cloned archive save set, specify both the save set ID and the clone ID. You must specify at least one annotation or save set ID.

- `-i{N|Y|R}` specifies how the NetWorker server should handle a naming conflict between a recovered file and an existing file:
 - `iN` does not recover the file when a conflict occurs.
 - `iY` overwrites the existing file when a conflict occurs.
 - `iR` renames the file when a conflict occurs. The recover process appends a `.R` to each recovered file name.
- `path` specifies the file or directory to recover. When you do not specify a path, NetWorker recovers all data in the archive save set.

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about additional options for the `nsrretrieve` command.

Troubleshooting NetWorker archiving and retrieval

This section explains how to troubleshoot issues with the Archive Module.

Remote archive request from server fails

If a remote archive request from the NetWorker server fails, ensure that the username for the archive client (for example, `root`) appears in the Archive Users attribute of the Client resource for the archive client.

You can also grant NetWorker administrator privileges for `root@client_system` in the Administrator attribute in the Server resource. However, be aware that NetWorker administrators can recover and retrieve data owned by other users on other clients.

Multiple save sets appear as a single archive save set

When you combine multiple save sets in an archive, such as `/home` and `/usr`, NetWorker stores the archived data in a single archive save set. To retrieve archives separately, archive the save sets separately.

Wrong archive pool is selected

If multiple archive pools exist in the NetWorker configuration, the archive operation will write the archive data to a volume in the last archive pool that was created on the NetWorker server .

Second archive request does not execute

If you create two archive requests with the same name, NetWorker will only perform the first request.

To ensure that NetWorker performs all of the archive requests, do not create two archive requests with the same name.

The nsrarchive program does not start immediately

If you run the `nsrarchive` command from a command prompt, the archive operation does not start immediately. Wait a short time until the archive starts. Do not press `[Ctrl]+[D]` multiple times to stop the archive operation.

Archive request succeeds but generates error when nsrexecd is not running

If the `nsrexecd` process is not running on a remote client during an archive request operation, NetWorker reports that the archive operations completed successfully, but the following error message appears in the `daemon.raw` file and the archive fails:

```
Failed to get port range from local nsrexecd: Service not available.
```

To resolve this issue, ensure that you start the `nsrexecd` daemon on a UNIX client or the NetWorker Remote Exec service on a Windows client before you perform an archive operation.

CHAPTER 8

Backup Data Management

This chapter contains the following topics:

• Overview of backup data management	526
• Viewing volume and save set details	526
• Managing volumes	537
• Changing save set status	540
• Changing the save set retention time	540
• Removing expired save sets	541

Overview of backup data management

After a backup occurs, there are several options to manage the save sets and volumes on backup storage.

The following backup data management features are available:

- View detailed status information about the save sets and volumes.
- Change the mode of a volume, for example, from Appendable to Read-only.
- Change the recycle policy for a volume to achieve greater control over the recycling of the volume.
- Relabel a library volume after all the save sets for the volume expire.
- Mark a volume as full for offsite storage.
- Remove a volume from the media database and online indexes, for example, if the volume is physically damaged.
- Change the status of a save set to Normal or Suspect.
- Clone save sets or volumes to create a copy of the backup data.
- Stage save sets to move data from one type of media to another.
- Archive data from a client, which copies the data to NetWorker storage and then removes the data from the client.
- Remove expired save sets so that you can recycle volumes and reclaim backup storage.

Viewing volume and save set details

The **Media** window of the NetWorker Administration interface provides details on volumes and save sets, including both backup and archive volumes and save sets. You can view save set details for a specific volume, or you can search for the save sets to view.

Viewing disk volume details

Procedure

1. In the **Administration** window, click **Media**.
2. In the left pane, select **Disk Volumes**.

A list of disk volumes for the server appears in the right pane. The following table lists the information that appears for each volume.

Table 83 Disk volumes window

Category	Description
Volume Name	Name of the volume, which is the same as the name that appears on the volume label in the NetWorker Administration interface. At the end of the name, one of the following designations might appear:

Table 83 Disk volumes window (continued)

Category	Description
	(A) indicates an archive volume. (R) indicates a read-only volume.
Media Type	The type of media for the volume.
Used	The amount of space currently in use on the volume, which is shown in KB, MB, or GB, as appropriate. The value of <code>full</code> indicates that there is no more space on the volume or an error has occurred.
Mode	Whether the volume is appendable, read-only, or recyclable: - Appendable volumes contain empty space. Data that meets the acceptance criteria for the pool to which this volume belongs can be appended. - Read-only volumes contain read-only save sets. No new data can be written to the volume. However, the save sets are still subject to retention settings, and the volume is recycled when the retention periods for all the save sets on the volume expire. When the mode is read-only, the Mode field appears blank. An (R) appears next to the volume name. - Recyclable volumes contain save sets that have all exceeded their retention periods.
Expiration	The expiration date for the volume. If the recycle policy is set to manual instead of automatic, then <code>manual</code> appears in this column. To change the expiration date for the volume, use the <code>nsrmm</code> command from the command prompt, or right-click the volume, select Recycle, and then select Manual on the Recycle dialog box.
Pool	Name of the pool to which the volume belongs.

Table 83 Disk volumes window (continued)

Category	Description
Location	An administrator-defined description of the physical location of a volume.

Viewing tape volume details

Procedure

1. In the **Administration** window, click **Media**.
2. In the left pane, select **Tape Volumes**.

A list of tape volumes for the server appears in the right pane. The following table lists the information that appears for each volume.

Table 84 Volume details

Category	Description
Volume Name	Name of the volume, which is the same as the name that appears on the volume label in the NetWorker Administration interface. At the end of the name, one of the following designations might appear: • (A) indicates an archive volume. • (R) indicates a read-only volume. • (W) indicates that the volume is a write once, read many (WORM) device.
Barcode	Barcode label for the volume, if one exists.
Used	The amount of space currently in use on the volume, which is shown in KB, MB, or GB, as appropriate. The value of <code>full</code> indicates that there is no more space on the volume and the end-of-tape marker has been reached, or that an error has occurred.
% Used	An estimate of the percentage that is used, based on the total capacity of the volume, and on the Media type setting of the device resource. A value of 100% indicates that the value is equal to or exceeds the estimate for this volume. A value of <code>full</code> indicates that the volume is full and you cannot write any more data to the volume, regardless of the estimate of the volume capacity.
Mode	Whether the volume is appendable, read-only, or recyclable: • Appendable volumes contain empty space. Data that meets the acceptance criteria for the pool to which this volume belongs can be appended. • Read-only volumes contain read-only save sets. No new data can be written to the volume. However, the save sets are still subject to retention settings, and the volume is recycled when the retention periods for all the save sets on the volume expire.

Table 84 Volume details (continued)

Category	Description
	When the mode is read-only, the Mode field appears blank. An (R) appears next to the volume name. Recyclable volumes contain save sets that have all exceeded their retention periods. You can also manually set the volume mode to full from the command prompt by using the <code>nsrjb</code> command with the <code>-o</code> option for libraries, and the <code>nsrmm</code> command with the <code>-o</code> option for stand-alone drives. When you set the volume mode to full, there is no more space for data in the volume, and the save sets have not yet exceeded the retention periods. The UNIX man pages of those commands and the <i>NetWorker Command Reference Guide</i> provide more information on the commands.
Expiration	The expiration date for the volume. If the recycle policy is set to manual instead of automatic, then <code>manual</code> appears in this column. To change the expiration date for the volume, use the <code>nsrmm</code> command from the command prompt, or right-click the volume, select Recycle, and then select Manual on the Recycle dialog box.
Pool	Name of the pool to which the volume belongs.
Location	An administrator-defined description of a physical location of the volume.

Viewing save set details for a volume

You display information about save sets on a volume.

About this task

Perform the following steps to view and print information about save sets on a volume, export the data to an HTML, CSV, or Post Script file, and filter the save set output for a particular time period.

Procedure

1. In the **Administration** window, click **Media**.
2. In the left navigation pane, select either **Disk Volumes** or **Tape Volumes**.
A list of volumes appears in the right pane.
3. To modify or save the information that appears in the window, perform one of following tasks:
 - To print the information that appears in the window, right-click the column header, and select **Print**.
 - To limit the output that appears in the window to a date range, right-click the column header and select **Show Filters**. Use the **From** and **To** drop downs to select the dates in the range. To remove the filters, click **Clear All**.
 - To export the data to a file, right-click the column header, and then select **Export**. From the menu, select the export format.

- To remove a column and the column details from the details window, right-click the column header that you want to remove, and then select **Remove This Column**.
 - To customize the columns that appear in the details window, right click in the column header, and select **Choose Table Columns**. Perform the following tasks:
 - Check the columns that you want to appear, and clear the columns that you want to hide.
 - Select a column from the box to choose a column on which to sort the save set details.
 - Select a column, and then use the up and down buttons to change the order in which the columns appear.
 - Click **Restore Defaults** to reset the save set details table to the default settings.
4. To view information about the save sets on a volume, right-click a volume, and then select **Show Save Sets**.

The **Volume Save Sets** window appears.

The following table lists the information that appears for each save set.

Table 85 Save Set details

Column	Description
Client	Name of the NetWorker client computer that created the save set.
Save Set	Pathname of the file system that contains the save set. This column also includes clone information. If the save set has a clone, the pathname is marked has clones and the cloned save set is marked clone save set.
SSID	Save set ID number.
Checkpoint ID	Checkpoint ID number.
Save Time	Date and time when the save set was created.
Clone Retention Time	Date and time when the clone expires.
Level	Level of backup that generated the save set. This refers only to scheduled backups. For manual backups, the level is blank.
Status	Status of the save set, such as whether the save set is browsable or recoverable.
Size	Size of the save set.
Flags	Flags that provide additional details about the save set. The first flag indicates which part of the save set is on the volume:

Table 85 Save Set details (continued)


Column	Description
	• c indicates that the save set is completely contained on the volume. • h indicates that the save set spans volumes and the head is on this volume. • m indicates that the save set spans volumes and a middle section is on this volume. • t indicates that the save set spans volumes and the tail section is on this volume. The second flag provides the save set status: • b indicates that the save set is in the online index and is browsable. • r indicates that the save set is not in the online index and is recoverable. • E indicates that the save set is eligible for recycling and may be overwritten at any time. • a indicates that the save set aborted before completion. Aborted save sets with targets of AFTD or DD Boost devices never appear in the Volume Save Sets dialog box or in <code>mminfo</code> reports because such save set entries are immediately removed from the media database. • i indicates that the save set is still in progress. The third flag is optional and provides the following information for the save set: • N indicates that the save set is an NDMP save set. • R indicates that the save set is a raw partition backup (such as for a supported module). • P indicates that the save set is a snapshot backup. The fourth flag is optional. If the fourth flag appears, the value is s to indicate that the save set is an NDMP

Table 85 Save Set details (continued)

Column	Description
	save set backed up by the <code>nsrdsa_save</code> command to a NetWorker storage node.
VBA	Provides details about a VBA save set.

5. To modify the information that appears in the window, perform one of following tasks:
 - To print the information that appears in the window, right-click the column header, and select **Print**.
 - To limit the output that appears in the window to a date range, right-click the column header and select **Show Filters**. Use the **From** and **To** drop downs to select the dates in the range. To remove the filters, click **Clear All**.
 - To export the data to a file, right-click the column header, and then select **Export**. From the menu, select the export format.
 - To add a new column of information, right click the column header, select **Add Column**, and then select a column option.
 - To remove a column and the column details from the details window, right-click the column header that you want to remove, and then select **Remove This Column**.
 - To customize the columns that appear in the details window, right-click in the column header, and select **Choose Table Columns**. Perform the following tasks:
 - Check the columns that you want to appear, and clear the columns that you want to hide.
 - Select a column from the box to choose a column on which to sort the save set details.
 - Select a column, and then use the up and down buttons to change the order in which the columns appear.
 - Click **Restore Defaults** to reset the save set details table to the default settings.

The following figure provides an example of the **Volume Save Sets** window, after you right-click on the column header.

Figure 50 Volume Save Sets window

6. Click **OK** on the **Volume Save Sets** dialog box.

Viewing save set details from a search

You can search for save sets associated with a policy or workflow in the **Media** window of the Administration interface. The search steps depend on whether you are searching for a normal save set or a VMware Backup Appliance save set.

About this task

You can Print the save sets, Set the Filter to show details of particular time period. Export to data to PDF, HTML, CSV and Post Script Add and remove column and Choose Table Columns

Based on the requirement Column can be sorted on Ascending or Descending Order

 **Note:** You cannot search for save sets that were created in releases prior to NetWorker 9.0.x.

Searching for save sets

Procedure

1. On the **Administration** window, click **Media**.
2. In the left pane, select **Save Sets**.
3. In the right pane, select **All Save Sets**.
4. On the **Query Save Set** tab, specify one or more of the search criteria in the following table.

Table 86 Query criteria

Criterion	Description
Client Name	Type the name of the client that is associated with the save set.
Save Set	Type the name of the save set.
Save Set ID	Type the identifier of the save set.
Volume	Select the volume on which the save set is stored from the list.
Pool	Select the media pool for the volume on which the save set is stored from the list.
Checkpoint ID	Type the identifier of the checkpoint for partial save sets.
Copies	To limit the save set results to the number of copies of the save set: a. From the Copies list, select whether the number of copies is less than (>), equal to (=), or greater than (<) a number that you specify. b. Specify the number in the second box.

Table 86 Query criteria (continued)

Criterion	Description
Save Time	Select the start and end dates and times for the save time of the save set.
Clone Retention Time	Select the start and end dates and times for the retention time of a cloned save set.
Status	Select All to view save sets of any status. Select Select from to view save sets of a specific status, and then select the checkbox next to one or more of the following statuses: • Browsable • Recoverable • Recyclable • Scanned-in • Suspect • Aborted • In-Progress • Checkpoint Enabled
Type	Select All to view save sets of any type. Select Select from to view save sets of a specific type, and then select the checkbox next to one or more of the following statuses: • Normal • Raw • Data Domain • Synthetic Full • Rehydrated • NDMP • Snapshot • ProtectPoint
Maximum Level	Select the maximum level of the backup. Save sets that meet the selected level and backups of levels below the selected level appear in the results.

5. Click the **Save Set List** tab.

A list of save sets that meet the search criteria appears with details for each save set. The following table provides more information.

Table 87 Save set search results view

Column	Description
Client	Name of the client.
Save Set	Name of the save set.
SSID	Save set identifier.
Clone ID	Clone identifier if the save set is a cloned save set.
Level	Backup level.
Status	Status of the save set, such as Recyclable or Recoverable.
Type	Type of backup, such as Normal or Synthetic Full.
Media	The media that contains the save set.
Volume Name	Name of the volume on which the save set is stored.
Pool	Name of the media pool for the volume on which the save set is stored.
Size	Size of the save set.
Files	Number of files in the save set
Save Time	Date and time at which the save set was saved to backup storage.
Clone Retention Time	Retention period for a cloned save set.
Checkpoint ID	Identifier of the checkpoint for a partial save set.

Searching for VMware Backup Appliance save sets

Procedure

1. In the **Administration** window, click **Media**.
2. In the left pane, select **Save Sets**.
3. In the right pane, select **VMware Backup Appliance Only**.
4. On the **Query Save Set** tab, specify one or more of the search criteria in the following table.

Table 88 Query criteria

Criterion	Description
VBA Name	Select the checkbox next to VBA Name above the list, and then select the VBAs from the list.
VM Name	Type the name of the virtual machine.
vCenter Name	Type the name of the vCenter for the VBA.

Table 88 Query criteria (continued)

Criterion	Description
Policy	Select the policy that generated the VBA save set.
Save Set ID	Type the save set identifier.
Volume	Select the volume on which the save set is stored from the list.
Pool	Select the media pool for the volume on which the save set is stored from the list.
Copies	To limit the save set results to the number of copies of the save set: - From the Copies list, select whether the number of copies is less than (>), equal to (=), or greater than (<) a number that you specify. - Specify the number in the second box.
Save Time	Select the start and end dates and times for the save time of the save set.
Status	Select All to view VBA save sets with any status. Select Select from to view VBA save sets of a specific status, and then select the checkbox next to one or more of the following statuses: • Recyclable • Recoverable • Suspect • Scanned-in • In-Progress

5. Click the **Save Set List** tab.

A list of VBA save sets that meet the search criteria appears with details for each save set. The following table provides more information.

Table 89 VBA save set search results window

Column	Description
VBA Name	Name of the VBA.
VM Name	Name of the virtual machine.
vCenter	Name of the vCenter for the VBA.
Policy	Name of the policy that generated the save set.
SSID	Save set identifier.

Table 89 VBA save set search results window (continued)

Column	Description
Clone ID	Clone identifier if the save set is a cloned save set.
Status	Status of the save set, such as Recyclable or Recoverable.
Media	Type of Media.
Volume Name	Name of the volume on which the save set is stored.
Pool	Name of the media pool for the volume on which the save set is stored.
Size	Size of the save set.
Save Time	Date and time at which the save set was saved to backup storage.
Clone Retention Time	Retention period for a cloned save set.

Managing volumes

A volume is a physical piece of media such as a tape cartridge or disk. On file type devices, a volume is a directory on a file system. Volume management tasks include changing the mode or recycle policy for the volume, relabeling the volume, removing volumes from the media database and online indexes, and marking a volume as full for offsite storage.

If a volume is not mounted when a backup is started, then one of three messages appears, suggesting that one of these tasks be performed:

- Mount a volume.
- Relabel a volume (only when Auto Media Management is enabled).
- Label a new volume (only when Auto Media Management is enabled).

During file recovery, the NetWorker server requests the volume name. If multiple volumes are needed to recover the files, the server lists all the volumes in the order of which they are needed. During the recovery process, the server requests each volume, one at a time. If a library is used, the server automatically mounts volumes that are stored in the library.

To manage volumes, you must have the correct permissions that are associated with the NetWorker server and its storage nodes.

Changing the volume mode

You can manually change the mode of a volume to a different mode such as read-only, recyclable, or appendable.

About this task

When the volume mode is read-only, no new data can be written to the volume, but the save sets are still subject to retention settings. However, a read-only volume is not a write-protected volume. When the retention period for all the save sets on the volume expire, the volume is recycled. Recyclable volumes contain save sets that have

all exceeded their retention periods. Appendable volumes can receive additional backup data.

Procedure

1. Unmount the volume by right-clicking the device in the **Devices** window and selecting **Unmount**.
2. In the **Administration** window, click **Media**.
3. In the left pane, select either **Disk Volumes** or **Tape Volumes**.
A list of volumes appears in the right pane.
4. Right-click the volume and select **Change Mode**.
The **Change Mode** dialog box appears.
5. Select a mode and click **OK**.
The new volume mode appears in the **Mode** column.
6. (Optional) Mount the volume by right-clicking the device in the **Devices** window, and selecting **Mount**.

Changing the volume recycle policy

You can override the retention policy for a volume by changing the recycle policy from automatic to manual. You may want to set the recycle policy to manual to keep save sets on a volume longer than the specified retention period. If you reset a volume to the automatic recycle policy, then the original retention policy applies to the volume.

Before you begin

Unmount the volume by right-clicking the device in the **Devices** window, and selecting **Unmount**.

About this task

 **NOTICE** A volume that has been set to manual recycle retains that setting, even after the volume is relabeled. You must explicitly reset the volume to automatic recycle.

Procedure

1. In the **Administration** window, click **Media**.
2. In the left pane, select either **Disk Volumes** or **Tape Volumes**.
A list of volumes appears in the right pane.
3. Right-click the volume, and select **Recycle**.
The **Recycle** dialog box appears.
4. Select either the **Auto** or **Manual** recycle policy.
5. Click **OK**.

After you finish

Mount the volume by right-clicking the device in the **Devices** window, and selecting **Mount**.

Marking a tape volume as full for offsite storage

When you remove a tape volume from a library to store the volume offsite, mark the volume as full so that the NetWorker software does not request the volume. Marking the volume as full also marks the volume as read-only. You can also specify the

physical location of the volume for reference purposes in the NetWorker Administration interface.

Procedure

1. Unmount the tape volume by right-clicking the volume in the **Devices** window, and selecting **Unmount**.
2. Use the `nsrjb` command for libraries or the `nsrmm` command for stand-alone drives from the command prompt to mark the volume as full:
 - For libraries, type `nsrjb -o full valid`, where *valid* is the volume identifier.
 - For stand-alone drives, type `nsrmm -o full valid`, where *valid* is the volume identifier.
3. Specify the physical location of the volume for reference purposes:
 - a. In the **Administration** window, click **Media**.
 - b. Select **Tape Volumes**.
A list of volumes appears in the right pane.
 - c. Right-click the volume in the right pane and select **Set Location**.
The **Set Location** dialog box appears.
 - d. Type a description for the physical location of the volume.
 - e. Click **OK**.

Removing volumes from the media database and online indexes

You may need to remove a volume from the media database and online indexes to eliminate physically damaged or unusable volumes from the NetWorker server.

About this task

When you remove the volume from the media database and online indexes, you can recover data from the volume by using the scanner program if the volume is undamaged.

If there is a clone of the volume, you cannot delete the volume entry from the media database. This is because the NetWorker server accesses the cloned volume rather than the original volume. As a result, removing volume entries from the media database is not an effective way to reduce index size, although it does reduce the size of the online indexes by deleting index entries that are associated with specific volumes.

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrmm` and `mminfo` commands.

Procedure

1. In the **Administration** window, click **Devices**.
2. In the left pane, select **Libraries**.
A list of libraries appears in the right pane.
3. Select the library in the left pane or double-click the library in the right pane.
The library drives and mounted volumes appear in the right pane, as well as the library slots and volumes.
4. Right-click the volume, and select **Unmount**.
You can only delete unmounted volumes.

5. Right-click the volume, and select **Delete**.

The **Delete** dialog box appears.

6. Specify the locations from which to remove the volume:

- Select **File and Media Index Entries** to remove the volume from both the media database and the online indexes.
- Select **File Index Entries Only** to remove the volume only from the online indexes.

Do not remove the indexes of save sets on bad volumes. In addition, do not remove both the client file index and media database entries simultaneously unless the volume is damaged or destroyed.

7. Click **OK**.

After you finish

After you remove a bad volume, perform an index consistency check by using the `nsrck` command in the command prompt. The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrck` command.

Changing save set status

You can manually change the status of a save set to either **suspect** or **normal**. Change the status to suspect if there may be a problem with the save set, for example, if a recovery from the save set failed.

About this task

The status of a save set may change to suspect automatically if the volume label of the volume for the save set cannot be read when the volume is ejected and the option to verify that the label is selected for the device.

Procedure

1. In the **Administration** window, click **Media**.
2. In the left pane, select either **Disk Volumes** or **Tape Volumes**.
3. Right-click the volume for the save set and select **Show Save Sets**.

The **Volume Save Sets** dialog box appears.

4. Select the save set.
5. Click **Change Status**.

The **Change Save Set Status** dialog box appears.

6. Select either the **normal** or **suspect** status for the save set.
7. Click **OK** on the **Change Save Set Status** dialog box.
8. Click **OK** on the **Volume Save Sets** dialog box.

Changing the save set retention time

You can change the expiration of a save set, including a cloned save set in three ways.

About this task

- Extend the retention time to a new expiration date.
- Keep the selected save set indefinitely, which sets the retention time to forever.

- Expire the save set immediately.


Perform the following steps to change the retention time on save sets:

Procedure

1. In the **Administration** window, click **Media**.
2. In the left navigation pane, select **Disk Volumes** or **Tape Volumes**.
3. Right-click the volume for the save set and select **Show Save Sets**.
The **Volume Save Sets** box appears.
4. Select the save set, and then click **Change Expiration**.

The **Change Expiration** window appears. The following figure provides an example of the **Change Expiration** window.

Figure 51 Change Expiration window



5. Perform one of the following tasks:
 - To define a new retention date, select **New Retention Time**, and then click on the calendar to select the date.
 - To keep the save sets indefinitely, leave the default selection **Keep the selected save sets indefinitely**.
 - To expire the save sets immediately, select **Expire the selected save sets now**.
6. Click **OK**.

The browse and retention attributes for the save set change.

Removing expired save sets

After the retention period for a save set expires (and the retention period for all the save sets that depend on the save set expire), the expire action, which is a part of the server maintenance workflow, marks the save set as recyclable in the media database. The NetWorker server tracks save set dependencies regardless of whether the dependent save sets are stored on the same or different volumes.

The activities that the expire action performs when a save set and all depend save sets expire, differs for advanced file type devices (AFTD) or Data Domain devices and tape volumes:

- **Tape volume**—Entries for save sets that are marked browsable are removed from the client file indexes. The status of the save set changes to recyclable in the media database. When all the save sets on the volume are recyclable, the mode of

the volume changes to recyclable. You can relabel and overwrite a recyclable volume to reclaim backup storage.

- AFTD or Data Domain devices—Entries for save sets that are marked browsable are removed from the client file index and media database. Entries that are recoverable are removed from the media database. The expire action removes the data that are associated with the save sets from the disk volume and reclaims the disk space.

The NetWorker server maintains one file index for each client computer (regardless of the number of client resources for the client), and one media database that tracks data from all clients and all save sets.

Note: For AFTD devices, there might be some instances where the save sets are not removed. This might happen, if you have any incremental or level backup dependency with previous backups. The save sets are removed only after all the dependent incremental or level backup gets into the recyclable mode.

Save set management on tape devices

Review the following information about save set status management for tape volumes.

A volume can contain save sets from multiple backup sessions, all with different retention policies. The mode of a tape volume might not change to recyclable in the media database for a long time. All data on the volume remains available for recovery by using either save set recovery or the `scanner` program. All entries for recyclable save sets remain in the media database.

You can also manually delete save set entries from the media database. However, the data on that volume is still available for recovery by using the `scanner` program. The `scanner` program retrieves the information that is needed to re-create entries in either the client file index, in the media database, or in both places:

- If you re-create the entries in the client file index, a user with the proper permissions can recover data by using the NetWorker client computer.
- If you re-create the save set entries in the media database, a UNIX root user or a member of the Windows Administrators group can recover data by using save set recovery.

Entries for a save set are automatically removed from the media database when NetWorker relabels the volume. You cannot recover data after NetWorker relabels a volume.

NOTICE When NetWorker relabels a volume for reuse within the same pool, the volume identification (the volume name as it appears on the volume label) remains unchanged. Although the volume has the same label, information that is required by the NetWorker server to locate and restore data on the volume is destroyed. All existing data is inaccessible and is overwritten.

If a volume contains one or more deduplication save sets, the resource for the deduplication node that was used to create the backup must exist when the save sets pass their retention time. If the resource for the deduplication node has been deleted, NetWorker cannot mark the volume as recyclable in the media database or relabel the volume. Furthermore, when deduplication save sets pass their retention time, the NetWorker server begins the process of deleting the deduplicated data from the deduplication node. Therefore, deduplication data may not be recoverable by using the `scanner` program when the deduplication save set has passed its retention time.

CHAPTER 9

Recovery

This chapter contains the following topics:

• Recovering data	544
• Recovery roadmap	544
• Prerequisites for recovering a NetWorker client or storage node	545
• Downloading the NetWorker software and documentation	545
• NetWorker recovery overview	546
• Recovery types	546
• Recover programs	550
• Recovering the data	557
• Recovering deduplication data	578
• vProxy recovery in NMC	578
• NMC function to collect vProxy log bundle information	600
• Recovering file system data on Windows	601
• Recovering data on OS-X clients	603
• Recovering client files on a different NetWorker server	610
• Recovering as NMC User with Non-Admin Privileges	612
• Preparing the NMC Sever for Recovery	612
• Recover the NMC Server database	614


Recovering data

You can recover NetWorker data by using the `recover` command, the NetWorker User program on Windows, or the **NMC Recovery** wizard on the NMC server.

Recovery roadmap

The following figure provides a high-level roadmap of the recovery paths for a NetWorker Client and Storage Node host.

Figure 52 Recovery roadmap



1. The "Special recoveries on Windows hosts" chapter describes how to perform a Bare Metal Recovery (BMR) of a NetWorker Client or Storage Node. The *NetWorker Module for Microsoft Administration Guide* describes how to perform a BMR of a Microsoft application data.
2. The "Special recoveries on Windows hosts" chapter describes how to recover the Active Directory Domain Service (AD DS) on a Windows domain controller.
3. The "Recovery" chapter describes how to recover file system data on Window, Linux, UNIX, and OS X hosts.

4. The NetWorker application documentation describes how to recover application data. For example, the *NetWorker Module for Microsoft Administration Guide* describes how to recover Microsoft application data that was backed up by using the NetWorker Module for Microsoft.

 **Note:** The *NetWorker Server Disaster Recovery and Availability Best Practices Guide* describes how to perform a disaster recovery of a host.

Prerequisites for recovering a NetWorker client or storage node

Before recovering a NetWorker client or storage node, perform the following steps.

Procedure

1. Verify that the same operating system as the source host is installed on the target host.
2. Verify that the NetWorker server is functioning and available on the network.
3. Obtain the following information:
 - NetWorker server hostname.
 - NetWorker client or storage node software version and patch level on the computer before the disaster occurred.
 - Link names to the NetWorker directories you must recover. An example of a typical link from a NetWorker directory to a user directory is `/nsr` to `/var/nsr`.

Downloading the NetWorker software and documentation

To obtain the latest NetWorker software and documentation, perform the following steps.

Procedure

1. Review the online NetWorker documentation, such as the *NetWorker Administration Guide*, *NetWorker Installation Guide*, and *NetWorker Release Notes*, for the latest information.
2. Obtain the required NetWorker cumulative hotfix media kits that provide customers with the opportunity to install the latest version of NetWorker including important hotfixes. Cumulative builds are released approximately once a month and each build contain a rollup of the fixes in each previous build.

If additional hotfixes are required in an environment where a cumulative build is installed, hotfixes can be generated for use with the latest cumulative version. The cumulative releases for specific NetWorker versions are available at the Online Support website.
3. Open the NetWorker Cumulative Hotfix document for details regarding fixes that are in each build, knowledge base articles that are related to the fixes in each build, and download instructions.

NetWorker recovery overview

Use the `recover` command, the NetWorker User program on Windows, or the **NMC Recovery** wizard on the NMC server to recover backup and clone data.

 **Note:** NetWorker 9.0.x and later does not support the recovery of archive data. Use an older version of the NetWorker client software to recover archive data.

Hosts in a NetWorker recovery operation

All recovery operations use three types NetWorker hosts to perform a recovery:

- **Administering host**—The NetWorker host that starts the recovery operation. The administering host can be the source host, the destination host, or another NetWorker host in the datazone.
- **Source host**—The NetWorker host from which the backup was run.
- **Destination host**—The NetWorker host that receives the recover data. The destination host can be the source host or another NetWorker host in the datazone.

Recovery types

NetWorker provides you with two types of recoveries.

- **Local recover**—A single NetWorker host is the administering, source, and destination host.
- **Directed recover**—The administering host is the source host or any other NetWorker host in the datazone. The destination host is not the source host. Use a directed recovery:
 - To centralize the administration of data recoveries from a single host.
 - To recover the data to a shared server, when the user cannot recover the data themselves.
 - To recover data to another host because the source host is inoperable or the network does not recognize the source host.
 - To transfer files between two NetWorker hosts.

Directed recoveries

A directed recovery enables a user to recover data to a NetWorker host that differs from the source of the backup, while retaining the original file ownership and permissions.

A directed recovery is a restricted NetWorker function available only to user accounts that have the necessary privileges that are required to perform the directed recovery operation.

A user with directed recovery privileges can:


- Browse the backup data of all NetWorker clients.
- Recover the data to any NetWorker client.

The following figure provides an example of a directed recovery.

In this figure:

- Saturn is the administering host.
- Neptune is the NetWorker server.
- Mars is the destination host.
- Pluto is the source host (not shown).
- The OS of destination host is the same as a source host.

Figure 53 A directed recovery from a remote client



In this figure, the numbers represent the following:

1. A user on Saturn sends a request to the NetWorker server to browse backup data from Pluto. If the user has the privileges that are required to perform a directed recovery of data on Pluto, the user can select the data to recover, and then starts the recovery operation.
2. The NetWorker server mounts the volume that contains the data in a local tape device.
3. The NetWorker server recovers the requested backup data to Mars.

Directed recover requirements

The following table summarizes the requirements for each host in a directed recover session.

Table 90 General recover requirements

Host	Requirements
Destination	Ensure that the destination host: • Is the same platform as the source host, for example, Linux to Linux, AIX to AIX, or Windows to Windows.

Table 90 General recover requirements (continued)

Host	Requirements
	• Uses the same file system as the source host, for example, XFS to XFS, UFS to UFS, or NTFS to NTFS. • Contains an entry for the administering host in the <code>servers</code> file. The <i>NetWorker Security Configuration Guide</i> provides more information about client-tasking rights and how to modify the <code>servers</code> file. • Is configured to accept directed recoveries from a remote host. Ensure that the Disable Directed Recover attribute is set to the default value No, in the NSRLA database. Editing a client NSRLA database on page 829 describes how to edit the NSRLA database. • Has the required access rights to receive data. ■ If you run the <code>nsrlogin</code> command on the administering host to create an authenticated recover session, ensure that the External Roles attribute of a user group with Remote Access All Clients privileges contains one of the following entries: – User DN for the authenticated user – Group DN for a group that contains the authenticated user ■ If you do not run the <code>nsrlogin</code> command on the administering host to create an authenticated session, the root user or the Administrator user on the destination host must appear in one of the following configurations: – A member of a NetWorker User Group with Remote Access All Clients privileges. Add an entry to the User attributes for the Root or Admin account in this format. – Added to the Remote Access attribute of the source host. For example: The source client is mars. The destination client, venus, is a Windows host. The Remote Access attribute for the client mars contains: Administrator@venus
Source	Ensure that the source host: • Is the same platform as the destination host, for example, Linux to Linux, AIX to AIX, or Windows to Windows. • Uses the same file system as the destination host, for example, XFS to XFS, UFS to UFS, or NTFS to NTFS. • Has the required access rights to enable the administering host to browse the data. ■ If you run <code>nsrlogin</code> on the administering host to create an authenticated recover session, ensure that the Remote access attribute on the source host contains one of the following entries: – User DN for the authenticated user – Group DN for a group that contains the authenticated user ■ If you do not perform a <code>nsrlogin</code> on the administering host to create an authenticated session, ensure that Remote access attribute on the source host

Table 90 General recover requirements (continued)

Host	Requirements
	contains the root user or the Administrator user of the administering host. For example: The source client is mars and the administering client is venus. The Administrator account on venus starts the recover program. The value in the Remote Access attribute for the client mars is: Administrator@venus
Administering	Ensure that the administering host: • Is a client of the NetWorker server that contains the backup information. The administering client can be a different platform from the source and destination clients. • Has the required access rights to perform the recover operation. ■ If you run the <code>nsrlogin</code> command on the administering host to create an authenticated recover session, ensure that the External Roles attribute of the Operators, the Application Administrators, the Database Administrators, or the Database Operators user group contains one of the following entries: – User DN for the authenticated user – Group DN for a group that contains the authenticated user ■ If you do not perform a <code>nsrlogin</code> on the administering host to create an authenticated session, ensure that Users attribute of the Operators, the Application Administrators, the Database Administrators, or the Database Operators user group contains the root user or the Administrator user of the administering host in the Users attribute. Note: If you do not use the Operators, the Application Administrators, the Database Administrators, or the Database Operators user group, ensure that you add the required user information to a user group that has the following privileges: ■ Remote Access All Clients ■ Operate NetWorker ■ Monitor NetWorker ■ Operate Devices and Jukeboxes ■ Backup Local Data ■ Recover Local Data ■ Recover Remote Data You must have operator privileges in the Operators user group to perform a selective file restore from a Microsoft Windows deduplication backup. Microsoft provides complete documentation for working with the Windows deduplication functionality.

Windows requirements

NetWorker enables you to perform directed recoveries of data to a local drive on Windows destination host, when you enable Windows File and Print Sharing option on the destination host. You cannot perform a directed recovery to a CIFS share.

When you use the `recover` command on a Windows destination host and the NetWorker server is also a Windows host, change the account that starts the NetWorker Backup and Recovery service on the NetWorker server:

- When the NetWorker server and the destination host are in the same domain, start service with a domain user that is a member of the local Administrators group.
- When the NetWorker server and destination host are not in a domain, or are not in the same domain, start the service with a local user that meets the following requirements:
 - The same username exists as a local user on the destination host.
 - The local user must have the same password on both hosts.
 - The local user on the NetWorker server is a member of the local Administrators group.

UNIX specific requirements

Review this information before you recover non-ASCII directories to a different directory on UNIX hosts.

- If the remote directory is an existing non-ASCII directory, the locale of the administering client must match the locale of the destination client.
- If the remote directory does not exist, NetWorker creates the relocation directory on the destination file system, which is based on the locale of the administering client.

Local recoveries

When you perform a local recovery, the administering host is also the source and destination host. Local recoveries are the simplest way to recover NetWorker data.

Ensure that user account that performs the recovery operation meets the following requirements:

- Belong to a NetWorker User Group that has the Recover Local Data privilege. If you use `nsrlogin`, add the DN or the user or group to the External Roles. If you do not use `nsrlogin`, add the account in `user@host` to the Users attribute. The *NetWorker Security Configuration Guide* provides more information.
- Have operating system ownership of the recovered files. The root user on UNIX, and a Windows Administrator have this privilege.
- Have write privileges to the local destination directories. The root user on UNIX, and a Windows Administrator have this privilege.

Recover programs

NetWorker provides you with the following tools to recover data.

- NetWorker Recover program—Recover GUI for OS-X hosts.

- **NMC Recovery** wizard—Recover wizard that you start from the NMC server. The **NMC Recovery** wizard provides a NetWorker datazone with a centralized recovery method.
- The `recover` command—CLI tool available on Windows, UNIX, and OS-X. Use the `recover` command to recover data from a command prompt. To perform multiple recovery operations in parallel, use multiple `recover` commands.
- NetWorker User program—Recover GUI for Windows hosts. Use the NetWorker User program to recover file system data when the administering client is Windows.
- The `scanner` command—CLI tool available on Windows, UNIX, and OS-X. Use the `scanner` command to recover data from a volume by save set ID (SSID) to the host that starts the program. To perform multiple recovery operations in parallel, use multiple `scanner` commands.

 **Note:** The NetWorker User, NetWorker Recover, and **NMC Recovery** wizard programs only recover data sequentially.

Using the NetWorker User program

Use the NetWorker User program to recover file system data when the administering client is Windows. To recover application data for Microsoft applications that are protected with NMM (NetWorker Module for Microsoft Applications) use the NetWorker Module for Microsoft Applications Client User program. The *NetWorker Module for Microsoft Administration Guide* provides more information.

 **Note:** The NetWorker log file in `\install_path\logs\networkr.raw` contains a record of every file that was part of an attempted recovery from the NetWorker User program. This file is overwritten with the next recovery. To save the information in the file, rename the file or export the information by using the `nsr_render_log` program.

Using the NetWorker Recovery program

Use the NetWorker Recovery program to recover file system data when the administering client is Mac OS-X.

Using the Recovery Wizard

NetWorker includes a new Recovery Wizard that allows you to recover data to NetWorker 8.1 and later clients from a centralized location, the NMC GUI. The Recovery Wizard supports browsable, save set, and directed recoveries. The Recovery Wizard does not support cross-platform recoveries.

Use the Recovery Wizard to configure scheduled and immediate recoveries of:

- File system backups.
- NDMP backups, when you use a NetWorker server 8.1.1 or later and NMC server 8.1.1 or later.

 **Note:** When you use NetWorker server 8.1 and earlier, the Recovery Wizard does not display NDMP clients in the Select Recovery Hosts window.

- Block Based Backups (BBB), when BBB is enabled for a client and BBB are available for recovery.
- BBB that you cloned to tape.

You can also use the Recovery wizard to configure an immediate recover of a Snapshot Management backup.

When you create a recover configuration by using the Recovery Wizard, NetWorker saves the configuration information in an NSR recover resource in the resource database of the NetWorker server. NetWorker uses the information in the NSR recover resource to perform the recover job operation.

When a recover job operation starts, NetWorker stores:

- Details about the job in the nsrjobsd database. [Using nsrrecomp](#) on page 706 describes how to query and report on recovery status.
- Output sent to stderr and stdout in a recover log file. NetWorker creates one log file for each recover job. [Troubleshooting the Recovery Wizard](#) on page 553 provides more information.

NOTICE NetWorker removes the recover log file and the job information from the job database based on value of the *Jobsdb retention in hours* attribute in the properties of the NetWorker server resource. In NetWorker 9.0.1, the default jobsdb retention is 72 hours.

Recovery Wizard requirements

Review this section before you use the Recovery Wizard.

Ensure that:

- The destination host is a client of the NetWorker Server.
- For a directed recover, the **Remote Access** attribute of the source client must contain the hostname of the destination client.
- The source and destination clients are running the NetWorker 8.1 or later software.
- **Note:** You can recover data from a pre-8.1 backup after you update the source host to NetWorker 8.1 or later.
- The account that you use to connect to the NMC Server has Configure NetWorker privileges. The *NetWorker Security Configuration Guide* provides more information.
- The required configuration is in place to perform a directed recover. [Directed recoveries](#) on page 546 provides more information.
- **Note:** On Windows, scheduled recovery to CIFS shared folder fails as recovery is spawned as system user and it does not have the required privileges to access the shared folder.

Create a new recover configuration

The **Recovery** wizard allows you to create and save a configuration that you can reuse or modify later.

Procedure

1. Use NMC to connect to the NetWorker server.
2. Click **Protection** from the left navigation pane, then select **Clients**.
3. Right-click the client from which you want to recover the data, then select **Recover**.

The **Recovery** wizard appears.

4. Browse through the **Recovery** wizard screens and define the configuration for the recover job.

Online help describes how to use the **Recovery** wizard.

Note: To avoid the over consumption of memory, NetWorker limits the number of files that you can view when you browse a directory that contain a large number of files, for example, 200,000 files. When NetWorker determines that displaying the number of files will exhaust memory resources, NetWorker will display a partial list of the files and a message similar to the following appears: Expanding this directory has stopped because the result has too many entries

Modifying a saved recover configuration

The Recovery Wizard allows you to save partial recover configurations and complete the configuration at a later time.

Procedure

1. Use NMC to the NetWorker server.
2. Click **Recover** on the Administration window toolbar. The **Recover** window appears. [Recover window](#) on page 62 provides more information about the **Recover** window.
3. In the **Configured recovers** window, right-click the saved recover configuration, select **Open Recover**.

Reusing recover configurations

When you define a recover configuration, the Recovery Wizard provides you with the option to save the recover configuration or delete the configuration after the recover completes. When you save the configuration, you can reuse the configuration information to perform a new recover job.

Before you begin

Connect to the NMC server from an NMC client. Ensure that the account you use to connect to the NMC server has Configure NetWorker privileges. The *NetWorker Security Configuration Guide* provides more information.

Procedure

1. Connect to the NetWorker server.
2. Click **Recover** on the **Administration** window toolbar. The **Recover** window appears. [Recover window](#) on page 62 provides more information about the **Recover** window.
3. In the **Configured recovers** window, right-click the saved recover configuration, select **Recover Again**.
4. Change the configuration as required and save the configuration with a new name.

Troubleshooting the Recovery Wizard

At the start time for a Recovery resource, `nsrd` uses an `nsrtask` process on the NetWorker server to start the recover job. The `nsrtask` process requests that the `nsrjobd` process on the NetWorker server run the recovery job on the destination client, then `nsrtask` monitors the job.

Once the recover job starts:

- The log files on the NetWorker server contain stdout and stderr information for the recover job. NetWorker stores the logs files in the following location, by default:

- **Windows:** C:\Program Files\EMC NetWorker\nsr\logs\recover
- **UNIX:** /nsr/logs/recover



Note: NetWorker names the log file according to the name of the recover resource and the time of the recovery job:
recover_resource_name_YYYYMMDDHHMMSS

- The jobsdb contains job status information for the recover job.

Debugging recover job failures from NMC

To troubleshoot a recovery issue by using NMC, configure the Recovery resource to display greater detail in the log file, then retry the recover configuration in debug mode:

Procedure

1. In the **Recover** window, right-click the recover configuration and select **Recover Again**.
2. Click the **Back** button until you reach the **Select the Recover Options** window.
3. Select **Advanced Options**.
4. Increase the value in the **Debug level** attribute to enable debugging. The higher the value, the more the debug output that appears in the recover log file.
5. Click **Next** until you reach the **Perform the Recover** window.
6. In the **Recover name** field, provide a new name for the recover configuration.
7. Click **Run Recover**.
8. Monitor the status of the recover job in the option in the **Recover** window.
9. When the recover completes, review the recover log file.

Debugging recovery failures from command line

To troubleshoot recovery issue from the command line, use the `nsradmin` and `nsrtask` programs.

Procedure

1. From a command prompt on the NetWorker Server, type `nsradmin`.
2. From the `nsradmin` prompt:
 - a. Set the resource attribute to the **Recover** resource. For example:

```
. type: nsr recover
```

- b. Display the attributes for the **Recover** resource that you want to troubleshoot. For example:

```
print name: recover_resource_name
```

Where *recover_resource_name* is the name of the Recover resource.

- c. Make note of the values in the **recover**, **recovery options**, and **recover stdin** attributes. For example:

```
recover command: recover;
recover options: -a -s nw_server.corp.com -c mnd.corp.com -
```

```
I - -i R;
recover stdin:
"<xml>
<browsetime>
May 30, 2013 4:49:57 PM GMT -0400
</browsetime>
<recoverpath>
C:
</recoverpath>
</xml>"
```

where:

- `nw_server.corp.com` is the name of the NetWorker server.
- `mnd.corp.com` is the name of the source NetWorker client.

3. To confirm that the `nsrd` process can schedule the recover job:

- a. To start the recover job, update the **Recover** resource:

```
update: name: recover_resource_name;start time: now
```

where *recover_resource_name* is the name of the Recover resource.

- b. Exit the `nsradmin` application.
- c. Confirm that the `nsrtask` process starts.
- d. If the `nsrtask` process does not start, review the `daemon.raw` file on the NetWorker server for errors.

4. To confirm that the NetWorker server can run the `recover` command on the remote host, type the following command on the NetWorker server:

```
nsrtask -D3 -t 'NSR Recover' recover_resource_name
```

Where *recover_resource_name* is the name of the Recover resource.

5. When the `nsrtask` command completes, review the `nsrtask` output for errors.

6. To confirm that the Recovery UI sends the correct recovery arguments to the `recover` process:

- a. Open a command prompt on the destination client.
- b. Run the `recover` command with the **recover options** that the Recover resource uses. For example:

```
recover -a -s nw_server.corp.com -c mnd.corp.com -I - -i R
```

- c. At the Recover prompt, specify the value in the **recover stdin** attribute.

 **Note:** Do not include the “,” or the ; that appears with the **recover stdin** attribute.

- d. If the `recover` command appears to stop responding, review the `daemon.raw` file for errors.
- e. When the `recover` command completes, review the recover output for errors. If the `recover` command fails, then review the values that are specified in the Recover resource for errors.

7. Use the `jobquery` command to review the details of the Recover job. From a command prompt on the NetWorker server, type: `jobquery`.
8. From the `jobquery` prompt, perform one of the following steps:

- a. To set the query to the Recovery resource and display the results of all recovery jobs for a Recovery resource, type:

```
print name: recover_resource_name
```

Where *recover_resource_name* is the name of the Recover resource.

- b. To set the query to a particular jobid and display the results of the job, type:

```
print job id: jobid
```

Where *jobid* is the jobid of the Recover job that you want to review.

 **Note:** Review the `daemon.raw` file on the NetWorker server to obtain the jobid for the recovery operation.

Common recovery error messages

This section contains a summary of common recovery error messages and resolutions.

Unable to connect to the server. Remote system error - unknown error

This error appears in the **Select the Recovery Hosts** window when the Wizard cannot contact the host that you selected as the source or destination host.

To resolve this issue, ensure that:

- The host is powered on.
- The NetWorker Remote Exec service (nsrexecd) is started.
- Name resolution for the host is working correctly.

Host *destination_hostname* is missing from the remote access list of *source_hostname*. Press [Yes] to update the remote access list of *source_hostname* with *destination_hostname*

This message appears in the **Select the Recovery Hosts** window when you select a destination host that does not have the correct permissions to receive directed recovery data.

To resolve this issue, click **Yes**. The Recovery Wizard will update the Remote access attribute in the properties of the source host with the hostname of the destination host.

If you click **No**, then you cannot proceed in the recovery wizard until you select a destination host that is in the Remote access attribute of the source host.

This host is either improperly configured or does not support this operation

This message appears in the **Select the Recovery Hosts** window after you select a source or destination host when the source or destination host is running NetWorker 8.0 or earlier.

Destination_host_name does not support *recovery_type*

This message appears in the **Select the Recovery Hosts** window after you select a destination host and the destination host does not support the recovery type that you selected. To resolve this issue, select a destination host that supports the recovery type.

Using the recover command

Use the `recover` command to perform the data recovery from a command prompt.

There are two recovery methods:

- Interactive mode—enables the user on the administering host to browse, and select files and directories from the source backup.
- Non-interactive mode—enables the user on the administering host to recover a directory or file immediately, without browsing the client file index for file information. Use non-interactive mode when you know the path to recover and do not need to browse through the backup data find it.

Scanner recovery

The `scanner` program enables you to recover data directly from a NetWorker volume.

Use the `scanner` program in the following scenarios:

- To perform a by-file-selection recovery, when the save set information is not in the client file index.
- To recover data directly from a tape.
- To recover data from an incomplete save set.

Recovering the data

Use one of the recovery applications to recover data.

NetWorker provides you with a number of recovery methods:

- Browsable recovery—By selecting individual files and folders.
- Save set recovery—By recovering all data in a save set.
- Scanner recovery—By recovering the data directly from the media
- VSS File Level Recovery—By recovering Windows System State data with VSS File Level Recovery (FLR).

Determining the volume for recovering cloned data

You can specify whether to use the original volume or a cloned volume to recover data in some recovery scenarios. In other scenarios, NetWorker decides which volume to use.

The following table provides details on when you can select the volume from which to recover data and when NetWorker selects the volume.

Table 91 Volume selection by recovery method

Recovery method	Volume selection
NMC Recovery wizard	Choose whether to specify the volumes or to allow NetWorker to select the volumes on the Obtain the Volume Information page of the wizard.

Table 91 Volume selection by recovery method (continued)

Recovery method	Volume selection
NetWorker User program	You can select the volume when you perform a save set recovery. NetWorker selects the volume when you perform a browsable recovery.
<code>recover</code> command	You can specify the clone pool for a browsable recovery or the clone ID for a save set recovery. If you do not specify the clone pool or the clone ID, then NetWorker selects the volume.

When NetWorker selects the volume from which to recover data, the recovery operation uses the following logic:

1. The highest priority is assigned to the volume (clone or original volume) that has a complete, non-suspect save set status. A complete save set that is suspect has a higher priority than an incomplete non-suspect save set.
2. If the volumes still have equal priority, then priority is assigned to the mounted volume.
3. If the volumes are mounted, then priority is based on the media type. The media types from highest to lowest priority are:
 - Advanced file type device
 - File type device
 - Other (such as tape or optical)
4. If the volumes are not mounted, then priority is based on the media location. The media locations from highest to lowest priority are:
 - Volumes in a library.
 - Volumes that are not in a library but are onsite (or, the `offsite` flag is not set).
 - Volumes that are offsite (or, the `offsite` flag is set).

To specify that a volume is offsite, use the `nsrmm` command. For example:

```
nsrmm -o offsite -V volume_id
```

where `volume_id` is the ID of the volume to mark offsite.

The volumes that are required for recovery appear in the **Required Volumes** window of the **NMC Recovery** wizard and the NetWorker User (Windows) programs.

Recovering access control list files

NetWorker allows a user to browse and recover files with associated access control lists (ACLs) in directories for which the user is not the primary owner. To recover files

with associated ACLs, enable the ACL passthrough attribute on the NetWorker server. The feature is enabled by default.

About this task

When the ACL passthrough attribute is disabled, the following message appears when a non-owner tries to browse ACL files in a directory: `Permission denied (has acl)`

To enable ACL passthrough, perform the following steps:

Procedure

1. On the **Administration** window, click **Server**.
2. In the left pane of the **Server** window, right-click the NetWorker server.
3. From the **File** menu, select **Properties**.
4. Select the **Configuration** tab.
5. In the **Recover** section, select **ACL passthrough**.

Browsable recovery

A file selection recovery method, or browsable recovery inspects the client file index that NetWorker creates for the source host, to gather information about backups. When the recovery process reviews entries in the client file index, you can browse the backup data and select the files and directories to recover. The retention policy that NetWorker applies to a backup determines the earliest versions of files and file systems that are available for recovery. [Backup retention](#) on page 386 provides more information about browse and retention policies.

Use a browsable recovery in the following scenarios:

- To recover a file or directory when you are not certain of its exact name or location.
- To recover a small number of files or directories. When you select many files and directories, the process of marking the files for recovery and the recovery process can take some time to complete, particularly from the NetWorker User program.
- To perform a directed recovery.
- To recover only the files that you select in one or more directories, not all files in a directory.

Adding information about recyclable save sets to the client file index

Each NetWorker client, including the NetWorker server, has a client file index (CFI). The CFI is a database that contains information about the files that are in a save set.

When NetWorker adds save set information into the media database and CFI, NetWorker assigns the save set a retention date, which is based on the retention policy that is assigned to the backup, clone, or archive. Browsable information about the save set remains in the CFI until the current date is equal to the retention date.

When the current date is equal to the retention date, NetWorker expires the save set and identifies the save set as no longer required for recovery, or as eligible for recycling. When the status of the save set is eligible for recycling, NetWorker removes the information about the save set from the CFI, and you cannot perform a browsable recovery of the save set data. Some applications, such as the NetWorker Module for Databases and Applications, require that a save set is browsable to perform a recovery.

You can make expired save set files browsable for recovery by adding the save set information back into the client file index.

Determining the status of a save set

Use the save set query feature in NetWorker Administration to determine the status of a save set.

About this task

Perform the following steps to determine the status of a save set and record the information that you require to add the save set information back into the client file index (CFI) for an expired save set.

Procedure

1. Connect to the NetWorker server that contains the data in NMC.
2. On the **Administration** window, click **Media**.
3. In the left pane, select **Save Sets**.
4. In the right pane, select **All Save Sets**.
5. On the **Query Save Set** tab, specify one or more of the search criteria in the following table.

Table 92 Query criteria

Criterion	Description
Client Name	Type the name of the client that is associated with the save set.
Save Set	Type the name of the save set.
Save Set ID	Type the identifier of the save set.
Volume	Select the volume on which the save set is stored from the list.
Pool	Select the media pool for the volume on which the save set is stored from the list.
Checkpoint ID	Type the identifier of the checkpoint for partial save sets.
Copies	To limit the save set results to the number of copies of the save set: a. From the Copies list, select whether the number of copies is less than (>), equal to (=), or greater than (<) a number that you specify. b. Specify the number in the second box.
Save Time	Select the start and end dates and times for the save time of the save set.
Clone Retention Time	Select the start and end dates and times for the retention time of a cloned save set.
Status	Select All to view save sets of any status.

Table 92 Query criteria (continued)

Criterion	Description
	Select Select from to view save sets of a specific status, and then select the checkbox next to one or more of the following statuses: • Browsable • Recoverable • Recyclable • Scanned-in • Suspect • Aborted • In-Progress • Checkpoint Enabled
Type	Select All to view save sets of any type. Select Select from to view save sets of a specific type, and then select the checkbox next to one or more of the following statuses: • Normal • Raw • Data Domain • Synthetic Full • Rehydrated • NDMP • Snapshot • ProtectPoint
Maximum Level	Select the maximum level of the backup. Save sets that meet the selected level and backups of levels below the selected level appear in the results.

6. Click the **Save Set List** tab.

Review the results of the query in the **Save Set List** window for the save set that you want to recover. If the value in the status column is not browsable, then record the values in the SSID, Clone ID, and level columns.

 **Note:** When the level value is anything other than full, ensure that you record the SSID and Clone ID for the previous full backup and all level backups in between.

The following table summarizes some of the status attributes assigned to the save set that are relevant to the process of adding save set information back into a CFI.

Table 93 Save set status

Status	Definition
Browsable	The save set is browsable. The save set has not exceeded the defined retention policy.
Recoverable	Information about the save set Information appears only in the media database. NetWorker does not allow information about some save sets, for example the bootstrap save set to appear in the CFI for browsing.
Recyclable	The save set has expired is eligible for recycling. The save set has exceeded the defined retention policy.
Incomplete	The save set did not complete. NetWorker does not store save set information about an incomplete save set in a CFI.

Using nsrmm to modify the save set properties

Modify the save set properties with the `nsrmm` command.

Procedure

1. When the save set is recyclable:

- a. Modify the save set entry to make it **recoverable** with the `nsrmm` command:

```
nsrmm -e MM/DD/YYYY> -S ssid/cloneid
```

where:

- *MM/DD/YYYY* is the date that is chosen to make the save set browsable from.
- *ssid/cloneid* is the save set ID/cloneid.

For example:

```
nsrmm -e "11/21/2009" -S 4294078835/1257402739
```

When more than one SSID was recorded, repeat this step for all SSIDs.

- b. Modify the save set to the **not recyclable** status:

```
nsrmm -o notrecyclable -S ssid/cloneid -y
```

where *ssid/cloneid* is the save set ID/cloneid.

For example:

```
nsrmm -o notrecyclable -S 4294078835/1257402739 -y
```

When more than one SSID was recorded, repeat this step for all SSIDs.

- c. Verify that the save set status is recoverable:

```
mminfo -q ssid=ssid -r sumflags
```

Recoverable save sets have an **r**, in addition to other values in the sumflags output.

For example:

```
mminfo -q ssid=4294078835 -r sumflags cr
```

When more than one SSID was recorded, repeat this step for all SSIDs.

2. Query the media database to confirm that the index save set for a client is recoverable:

```
mminfo -avot -N index:client_name
```

where *client_name* is the name of the client to which this save set is located.

3. Confirm that the value in the **fl** column is **cr** for an index backup with the time frame of the client save set to be restored.



NOTICE If the index save set is not recoverable, the save set expires when the NetWorker software cross checks the indexes. For example, when the NetWorker server runs the `nsrim -X` command.

4. Record the values in the date and time columns.

Repopulating the client file index

Use the `nsrck` or `scanner` command to repopulate the client file index with information about files in a save set.

Repopulate the client file index by using the `scanner` program

Use the `scanner` program to repopulate the client file index with information about files and directories for a specific save set.

About this task

The entries assume the browse policy of the original save set. For example, suppose a save set originally had a browse time of one month and a retention time of three months. However, the browse and retention times have expired. When you restore the save set entry by using the scanner program, the save set then remains browsable for one month and recoverable for three months.

To Repopulate the client file index by using the `scanner` program, perform the following steps:

Procedure

1. Ensure the **idle device timeout** value of the device containing the volume is **0**. Refer to [Unmounting volumes automatically \(idle device timeout\)](#) on page 155 for details.
2. Query the media database using the `mminfo` program for save set information:

For example:

```
mminfo -avq ssid=ssid -r  
volume,client,name,ssid,mediafile,mediarec
```

where *ssid* is the associated save set id for the data you want to recover.

3. Use the information from the `mminfo` command for the save set to run the **scanner** program. When the save set spans more than one volume, scan the volumes in the order in which in which they were written:

```
scanner -v -i -S ssid -f mediafile -r mediarec device
```

where:

- *mediafile* is the starting file number for the save set, obtained from the `mminfo` output.

- *mediarec* is the starting record number for the save set, obtained from the *mminfo* output.
 - *device* is the name of the device the volume is loaded in, for example */dev/rmt0.1* or *\\.\Tape0*.
4. When the save set spans multiple volumes, the **scanner** program prompts for a new volume as needed.

Repopulating the client file index by using the **nsrck** program

Use the **nsrck** program to repopulate the client file index with information about all save sets for the client up to the date and time specified.

Procedure

1. Ensure that the volume containing the index backup is available.
2. Use the **nsrck** command to repopulate the client file index:

```
nsrck -L 7 -t MM/DD/YYYY client_name
where:
```

- where *client_name* is the name of the client with the data to be recovered.
- *MM/DD/YYYY* is the backup date of the save set.

For example:

```
<NetWorker_install_path>\nsr\bin>nsrck -L 7 -t
"11/21/2009" swift nsrck: checking index for 'swift'
9343:nsrck: The file index for client 'swift' will be
recovered.Requesting 1 rec over session(s) from server
Recover completion time: 11/20/2009 1:45:55 PM nsrck:
<NetWorker_install_path>\nsr\index\swift contains 12
records occupying 2 KB nsrck: Completed checking 1
client(s)
```

When you recover a client file index from a time and date in the past, **nsrck** adds the full contents of the index from that time and date to a temporary subdirectory of the client file index directory. When a time value is not specified, everything for the specified date (up to 23:59) is included. After the index has been read from the backup media, the required index data is integrated fully into the client file indexes and the temporary subdirectory is removed. The “required index data” includes the indexes from the date specified to the first full backup that occurred prior to the date specified.

Be aware that if a save set from the specified date runs into the next day, which would be Nov 22, 2009 in this example, then the index required to browse the save set will not be recovered. To recover this index, you would have to specify Nov 22, 2009 as the recovery date as shown in the following command:

```
nsrck -t "11/22/2009" -L7 swift
```

A check on the required index date may be necessary if index backups are set to be taken once daily. When the back up of the index does not take place until the following day, the date of the following day must be specified.

3. Confirm that the client save sets are now browsable:

```
mminfo -q ssid=ssid -r sumflags
```

Browsable save sets contain a *b*, in addition to other values in the *sumflags* output.

For example:


```
NetWorker_install_path\nsr\bin>mminfo -q ssid=4294078835 -r
sumflags
cb
```

4. Perform a file-by-file recovery by using the NetWorker User program (Windows), the **recover** command or the NMC Recovery Wizard.

Adding information about a save set in the client file index and media database

When a volume contains a save set that does not appear in the media database or client file index, use the `scanner` command to restore save set information into the media database and client file indexes.

About this task

Procedure

1. Log in as root or a Windows Administrator.
2. Load the first volume that contains the save set information into an available device. Ensure the **Idle Device Timeout** value for the device is 0. Refer to [Unmounting volumes automatically \(idle device timeout\)](#) on page 155 for details.
3. At the command prompt, run the **scanner** and specify the name of the device that contains the volume:

```
scanner device_name
```

4. Use the output from the `scanner` program to determine:
 - If the volume contains the save set that you want to scan.
 - If you want to scan the contents of the volume in the online indexes.
 - If the save set spans multiple volumes.
5. Use the `scanner` command to add the save information into the media database and CFIs:
 - To repopulate media database and CFIs with the save set information for all save sets on the volume, type `scanner -i device_name`
 - To repopulate the media database and client file index with the save set information for a specific save set , type `scanner -i -S ssid`

```
device_name
```

NOTICE When the volume contains data from an earlier version of NetWorker, there may be no pool information on the volume. In this case, the volume is considered to belong to the Default pool. To assign the volume to another pool, use the `-b pool_name` option in this step. When the volume already belongs to a pool, the `-b` option will have no effect.

Performing a browsable recover with NetWorker User

Perform these steps on the administering host.

Procedure

1. Open the NetWorker User program.
To recover data that was encrypted with the current AES pass phrase, no special action is required. However, to recover data that was encrypted with an AES pass phrase that is different than the current pass phrase, start the

`recover` command specify the `-p pass_phrase`. To enter multiple pass phrases with the `-p` option, type: `recover -p pass_phrase1 -p pass_phrase2 -p pass_phrase3`.

NOTICE When an incorrect pass phrase or no pass phrase is entered, encrypted data is not recovered. Instead, the file names are created without data. However, if unencrypted data is also selected for recovery, it is recovered.

2. Select the NetWorker server when you are prompted.
3. From the **Operations** menu, select **Recover/Directed**. To perform a save set recover, select **Save Set Recover**.
4. Select the source host that has the data you want to recover, then click **OK**.
5. Select the destination host for the recovered data, then click **OK**.
6. Mark the files and directories to recover, in the **Recover** window.

Note: When a drive letter is not present on the destination client, the drive appears with a red question mark.

7. Select optional recover options. The following table summarizes the available recovery options.

Table 94 Optional browsable recovery options

Recover option	Details
Change the browse time	The Recovery window appears with the latest version of the backup files. To change the browse date and time for all files in the Recovery window: Select View > Change Browse Time. On the Change Browse Time window, select a new day within the calendar. Select Previous Month or Next Month to change from the current month. In the Time field, change the time of day by typing an hour, minute, and the letter a (for a.m.) or p (for p.m.). Use the 12-hour format. Click OK.
View all versions of a selected file or directory	The Recovery window appears with the latest version of the backup files. When you mark a file system object for example, a file or directory, you recover the last backup version. To view earlier versions of file system objects:

Table 94 Optional browsable recovery options (continued)

Recover option	Details
	Highlight the file or directory that you want to review. Select View > Versions. Select a previous version. Select Change Browse Time. When prompted to change the browse time, click OK. Mark the new version of the file system object.
Search for file system objects	To search for file system objects in the defined browser time: From the File menu, select Find. Type the name of the file or directory. Use wildcards to expand the search. Without wildcards, partial file names result in no match being found.
Relocate the recovered file system objects	By default, NetWorker recovers file system objects to their original location. To relocate the files to a different location: Select Options > Recover Options In the Relocate Recovered Data To field, type the path on the destination host to recover the data, then click OK. For NDMP data restores, the target path is a string and must match the path as seen by the NAS filer in its native OS. Otherwise, NetWorker recovers the files to the original location and overwrites the existing file host with the same name. <i>NetWorker Network Data Management Protocol (NDMP) User Guide</i> provides details about NDMP recoveries.
View volumes required for recovery	Before you start the recovery operation, monitor which volumes NetWorker requires to recover the selected file system objects. To view the required volumes, select View > Required Volumes. Ensure that the listed volumes are available or NetWorker to mount into an available device.

Table 94 Optional browsable recovery options (continued)

Recover option	Details
Resolve name conflicts	By default, the Naming Conflict window appears each time there is a file name conflict during a recovery. To specify the method to automatically resolve all name conflicts: Select Options > Recover Options. Select a conflict resolution option: - Rename the recovered files. By default, the recover operation appends a tilde (~) to the beginning of the name of the recovered file <i>~file name</i>. When a file named <i>~file name</i> already exists, the recovered file is renamed <i>~00_file name</i>, and so forth, to <i>~99_file name</i>. When this fails, the recover process does not automatically rename the file and prompts the user is to specify a name for the file. - Discard recovered file: Discards the recovered file and keeps the existing file. - Overwrite existing file: Replaces the file on the file system with the recovered version. - Overwrite and replace a reboot: Replaces the file on the file system with the recovered version after you restart the destination host. NDMP recoveries do not support resolving name conflicts NDMP recoveries always overwrite existing files. Relocate the NDMP data to a different location to avoid data loss. <i>NetWorker Network Data Management Protocol (NDMP) User Guide</i> describes how to perform NDMP recoveries

- Click **Start** to begin the recovery. It takes the NetWorker server a few moments to recover the files, depending on file size, network traffic, server load, and tape positioning. During this time, messages appear so that you can monitor the progress of the recovery.

When the recovery is successful, a message similar to this appears:

```
Received 1 file(S) from NSR server server
Recover completion time: Tue Jan 21 08:33:04 2009
```

NOTICE When an error occurs while recovering Microsoft Exchange Server or Microsoft SQL Server data by using VSS, you must restart the recovery process. When the recovery fails due to a problem with VSS or a writer, an

error message appears. Use the Windows Event Viewer to examine the event logs for more information. VSS recovery error messages are also written to the NetWorker log file.

Performing a browsable recover by using the recover command

Use the `recover` command in interactive mode to access the client file index of the source client and recover individual files and folder from a command prompt. Interactive mode enables you to browse and select files and directories from a save set. NetWorker supports a local or directed browsable recovery from a command prompt. You cannot recover the Windows `DISASTER_RECOVERY:\` save set in interactive mode.

Before you begin

The `recover` command requires specific privileges which are assigned based on session authentication. NetWorker supports two types of session authentication. Token-based authentication, which requires you to run the `nsrlogin` before you run the command and authenticates the user that runs the command against entries that are defined in the External Roles attribute of a User Group resource. Classic authentication, which is based on user and host information and uses the user attribute of a User Group resource to authenticate a user. Classic authentication does not require an authentication token to run the command. For example, if you run the command without first running `nsrlogin`, NetWorker assigns the privileges to the user based on the entries that are specified in the Users attribute of the User Group resource. When you use `nsrlogin` to log in as a NetWorker Authentication Service user, NetWorker assigns the privileges to the user based on the entries that are specified in the External Roles attributes of the user Group resource. The *NetWorker Security Configuration Guide* provides more information about privileges provides more information.

For Windows hosts only, to ensure that you use the NetWorker `recover.exe` command and not the Windows OS `recover` command, perform one of the following tasks:

- Ensure that `NetWorker_install_path\bin` appears before `%SystemRoot%\System32` in the `$PATH` environment variable.
- When you start the `recover` command include the path to the binary. For example: `NetWorker_install_path\bin\recover.exe`.

About this task

Perform the following steps on the destination host in the data zone.

Procedure

1. Use the `mminfo` command to display information about the save set of the data that you want to recover. For example, type:

```
mminfo -r volume,savetime,client,ssid,cloneid,name
```

Output similar to the following appears:

Table 95 Save set information

volume	date	client	ssid	pool	name
backup.001	05/03/2015	bu_iddnwserve r	3644194209	Default	C:\ddlib

Table 95 Save set information (continued)

clone.001	05/03/2015	bu_iddnwserve	3644194209	Default Clone	C:\ddlib r

The `mminfo` command provides you with information that you require to recover the save set. For example, the name of the volume that contains the save set, the date that the save set was created and the name of the pool that contains the volume. NetWorker assigns each backup and clone save set the same save set ID (SSID) and unique clone ID (cloneid). To recover from a clone volume, the name of the clone pool is required.

2. Ensure that the volume which contains the save set is available for a device in the datazone.
3. Use the `recover` command to select and then to recover the data from the backup save set or the clone save set.

For example, type:

```
recover -t date -c source_host -R destination_host -b pool_name
-i i_recover_option
```

where:

- *date* is the date that NetWorker created the save set.
i **Note:** When you do not specify a date, the `recover` command displays the latest version of each file in the save set.
- *source_host* is the original data host.
i **Note:** When you do not specify source host, NetWorker assumes that the source client is the host where you run the recover program.
- *destination_host* is the host on which to recover the data.
- *pool_name* is the name of the pool that contains the volume. Use this option when you want to recover data from a clone volume.
- `-i i_recover_option` specifies how NetWorker handles a naming conflict between a recovered file and an existing file.
 - `iN` does not recover the file when a conflict occurs.
 - `iY` overwrites the existing file when a conflict occurs.
 - `iR` renames the file when a conflict occurs. The recover process appends a `.R` to each recovered file name.
- i **Note:** The `recover` command requires the `-i` option when you use the `-R` option to perform a directed recovery.

For example, to recover the data from a clone volume from a clone operation that was performed on July 20, 2015, type:

```
recover -t 07/20/2015 -b Default Clone
```

The **Recover** prompt appears.

4. Select the files or directories and perform the recover:
 - a. Specify the directory to browse:

```
recover> cd path
```

For example: `cd /var/adm`

b. Select the file or directory for recovery:

```
recover> add file_name
```

For example: `add system.log`

Note: On Windows, to recover files or directories that begin with a dash (-) such as `-Accounting`, try one of the following options:

- Type `add ./-Accounting` to recover the `-Accounting` file or directory and its contents.
- Use the `cd` command to change directories to `-Accounting`. Type `add .` to add the current directory and the directory contents for recovery.
- When the current directory is `/temp` and `-Accounting` resides in the `/temp` directory, type `add /temp/Accounting`. This input adds `-Accounting` and the contents of the directory to the recovery list.

c. To view the files or directory that you marked for recovery, type:

```
recover> list
```

d. To view the list of the volumes that NetWorker requires to recover the data, type:

```
recover> volumes
```

e. To recover the files to a location that differs from the original location, type:

```
recover> relocate path
```

5. To start the recovery operation, type:

```
recover> recover
```

When the recovery process completes, messages similar to the following appear:

```
Received 1 file(s) from NSR server `bu-idd-nwserver2'
Recover completion time: Tue Aug 21 08:33:04 2015
recover>
```

6. To close the `recover` program, type `quit`.

Save set recovery

The save set selection recovery method, or save set recover enables you to recover data without browsing and selecting the files for recovery. Unlike a browsable recovery, a save set recover does not inspect the client file index for information about each selected file.

When you perform a save set recovery, NetWorker recovers the last full backup first, then recovers incremental backups in the chronological backup order. [Backup levels](#) on page 359 provides information about the relationship between each backup level.

Use a save set recovery in the following scenarios:

- To recover many files or all the data in a save set, for example, if there is a total disk failure. When you perform a save set recovery, you do not select individual files or directories for recovery.
- To recover data from a recyclable save set. [Backup retention](#) on page 386 provides more information about browse and retention policies. [Adding information about recyclable save sets to the client file index](#) on page 559 describes how to repopulate the client file index entries for recyclable (expired) save sets.
- To recover data on a host with limited memory resources. A save set recovery requires less memory than a browsable recovery.

Performing a save set recover with NetWorker User

Perform the following steps on the administering host.

About this task

 **NOTICE** Only members of the Windows Administrators group have permission to perform a save set recovery.

Procedure

1. Open the NetWorker User program.

To recover data that was encrypted with the current AES pass phrase, no special action is required. However, to recover data that was encrypted with an AES pass phrase that is different than the current pass phrase, start the `recover` command specify the `-p pass_phrase`. To enter multiple pass phrases with the `-p` option, type: `recover -p pass_phrase1 -p pass_phrase2 -p pass_phrase3`.

 **NOTICE** When an incorrect pass phrase or no pass phrase is entered, encrypted data is not recovered. Instead, the file names are created without data. However, if unencrypted data is also selected for recovery, it is recovered.

2. Select the NetWorker server when you are prompted.
3. Select **Operation > Save Set Recover**.
4. Select the source host that has the data that you want to recover, and then click **OK**.
5. In the **Save Sets** window, select the name of the save set from the **Save Set Name** list.
6. Select the version of the save set . Optionally, select the cloned version of a save set.
7. Select optional recover options. The following table summarizes the recover options that are available with a save set recovery.

Table 96 Optional save set recovery options

Recover option	Description
Specify file system objects	By default, NetWorker recovers all selected files and directories. To recover only certain file system objects in a save set:

Table 96 Optional save set recovery options (continued)

Recover option	Description
	Click Files... Specify the files and directories to recover, one full path per line. Click OK.
View required volumes	Before you start the recovery operation, monitor which volumes NetWorker requires to recover the selected file system objects. To view the required volumes, select Required Volumes. Ensure the listed volumes are available for NetWorker to mount into an available device.
Relocate the recovered file system objects	By default, NetWorker recovers file system objects to their original location. To relocate the files to a different location: Select Recover Options. In the Relocate Recovered Data To field, type the full path of the directory where the data should be relocated and then click OK. For NDMP data restores, the target path is a string and must match the path as seen by the NAS filer in its native OS. Otherwise, the recover process uses the original location and overwrites existing files with the same name. <i>NetWorker Network Data Management Protocol (NDMP) User Guide</i> provides details about NDMP recoveries.
Resolve name conflicts	By default, the Naming Conflict window appears each time there is a file name conflict during a recovery. To specify the method to automatically resolve all name conflicts: Select Options > Recover Options. Select a conflict resolution option: Rename the recovered files. By default, a tilde (~) is appended to the beginning of the name of the recovered file <i>~file name</i>. When a file named <i>~file name</i> already exists, the recovered file is renamed <i>~00_file name</i>, and so forth, to <i>~99_file name</i>. When this fails, the recover process does not automatically

Table 96 Optional save set recovery options (continued)

Recover option	Description
	rename the file and prompts the user to specify a name for the file. • Discard recovered file: Discards the recovered file and keeps the existing file. • Overwrite existing file: Replaces the file on the file system with the recovered version. • Overwrite and replace a reboot: Replaces the file on the file system with the recovered version after you restart the destination host. NDMP recoveries do not support resolving name conflicts. NDMP recoveries always overwrite existing files. Relocate the NDMP data to a different location to avoid data loss. <i>NetWorker Network Data Management Protocol (NDMP) User Guide</i> describes how to perform NDMP recoveries

8. Click **OK** to begin the recovery. The NetWorker server takes a few moments to start the file recovery, depending on file size, network traffic, server load, and tape positioning. When NetWorker starts to recover the files, messages appear that enable you to monitor the progress of the recovery.

When the recovery is successful, a message similar to the following appears:

```
Received 1 file(S) from NSR server server Recover
completion time: Tue Jan 21 08:33:04 2009
```

Performing a save set recover from the command prompt

Use the `recover` command in non-interactive mode to perform a save set recover data from a command prompt. Non-interactive mode enables you to recover a directory or file immediately, without browsing the client file index for file information. Use non-interactive mode to recover data when you know the path to recover and you do not need to browse through the directory contents of the save set. NetWorker only supports a local save set recover. You cannot perform directed recover by using a save set recover.

Before you begin

The `recover` command requires specific privileges which are assigned based on session authentication. NetWorker supports two types of session authentication. Token-based authentication, which requires you to run the `nsrlogin` before you run the command and authenticates the user that runs the command against entries that are defined in the External Roles attribute of a User Group resource. Classic authentication, which is based on user and host information and uses the user attribute of a User Group resource to authenticate a user. Classic authentication does not require an authentication token to run the command. For example, if you run the command without first running `nsrlogin`, NetWorker assigns the privileges to the user based on the entries that are specified in the Users attribute of the User Group resource. When you use `nsrlogin` to log in as a NetWorker Authentication Service

user, NetWorker assigns the privileges to the user based on the entries that are specified in the External Roles attributes of the user Group resource. The *NetWorker Security Configuration Guide* provides more information about privileges

Procedure

1. Connect to the target host with the root account on UNIX or the Administrator on Windows.
2. Use the `mminfo` command to display information about the save set of the data that you want to recover.

For example, type: `mminfo -av -r volume,savetime,client,ssid,cloneid,name`
Output similar to the following appears:

Table 97 Save set information

volume	date	client	ssid	clone id	name
backup.001	05/03/2015	bu_iddnwserve r	3644194209	1362492833	C:\ddlib
clone.001	05/03/2015	bu_iddnwserve r	3644194209	1362493448	C:\ddlib

The `mminfo` command provides you with information that you require to recover the save set. For example, the name of the volume that contains the save set and the date that the save set was created. NetWorker assigns each backup and clone save set the same save set ID (SSID) and unique clone ID (cloneid).

3. Ensure that the volume which contains the save set is available for a device in the datazone.
4. Use the `recover` command to recover the data from the backup save set or the clone save set.

Note: To perform concurrent recoveries from an advanced file type by either using multiple `-S` options to identify multiple save sets, or starting multiple `recover` commands.

- To recover the all the data from a backup save set, type the following command:

```
recover -S ssid - i_recover_option
```

where:

- `ssid` is the SSID of the backup save set.
- `- i_recover_option` specifies how NetWorker handles a naming conflict between a recovered file and an existing file.
 - `iN` does not recover the file when a conflict occurs.
 - `iY` overwrites the existing file when a conflict occurs.
 - `iR` renames the file when a conflict occurs. The recover process appends a `.R` to each recovered file name.

For example:

```
recover -S 3644194209 -iR
```

- To recover the all the data from a clone save set, type the following command:

```
recover -S ssid/cloneid
```

where:

- *ssid* is the SSID of the backup save set.
- *cloneid* is the cloneid of the clone save set.

For example:

```
recover -S 3644194209/1362493448
```

Note: When you do not specify the cloneid of the save set, the `recover` command recovers the data from the backup save set.

- To recover a single directory from the clone save set and relocate the data to a new directory location, type the following command:

```
recover -S ssid/cloneid -d destination_dir original_dir
```

where:

- *ssid* is the SSID of the backup save set.
- *cloneid* is the cloneid of the clone save set.
- *destination_dir* is the location to which you want to recover the data.
- *original_dir* is the directory that is contained in the save set that you want to recover.

For example, to recover the directory `/var/adm` on the backup save set to the `/usr/mnd` directory, type the following command:

```
recover -S 3644194209/1362493448 -d /usr/mnd /var/adm
```

- To recover data that was encrypted with the current AES pass phrase, no special action is required. However, to recover data that was encrypted with an AES pass phrase that is different than the current pass phrase, start the `recover` command specify the `-p pass_phrase`. To enter multiple pass phrases with the `-p` option, type: `recover -p pass_phrase1 -p pass_phrase2 -p pass_phrase3`.

NOTICE When an incorrect pass phrase or no pass phrase is entered, encrypted data is not recovered. Instead, the file names are created without data. However, if unencrypted data is also selected for recovery, it is recovered.

Using the scanner program to recover data

You can use the `scanner` command to recover data from a volume by save set ID (SSID) to the host that starts the program. Ensure that the operating system of the NetWorker host that runs the `scanner` command is the same operating system as the source client.

Before you begin

The `scanner` command requires specific privileges which are assigned based on session authentication. NetWorker supports two types of session authentication. Token-based authentication, which requires you to run the `nsrlogin` before you run the command and authenticates the user that runs the command against entries that are defined in the External Roles attribute of a User Group resource. Classic authentication, which is based on user and host information and uses the user attribute of a User Group resource to authenticate a user. Classic authentication does

not require an authentication token to run the command. For example, if you run the command without first running `nsrlogin`, NetWorker assigns the privileges to the user based on the entries that are specified in the Users attribute of the User Group resource. When you use `nsrlogin` to log in as a NetWorker Authentication Service user, NetWorker assigns the privileges to the user based on the entries that are specified in the External Roles attributes of the user Group resource. The *NetWorker Security Configuration Guide* provides more information about privileges

About this task

 **NOTICE** You cannot use the `scanner` command recover data from a NetWorker Module, NDMP or DSA save set.

Procedure

1. Optionally, use the `nsrlogin` command to authenticate a user and generate a token for the
[Using nsrlogin for authentication and authorization](#) provides more information.
2. Ensure the value in the **Idle device timeout** attribute of the device that contains the volume is 0. [Unmounting volumes automatically \(idle device timeout\)](#) on page 155 provides more information.
3. Use the `mminfo` program to query the media database for save set information.

For example:

```
mminfo -avq ssid=ssid -r
volume,client,name,ssid,mediafile,mediarec
```

where `ssid` is the save set ID associated with the data.

4. Use the save set information from the `mminfo` command to run the `scanner` program:

- To recover all files in a save set on Windows, type:

```
scanner -v -Sssid -f mediafile -r mediarec device | path\uasm
-rv
```

where:

- `ssid` specifies the save set ID value that you obtained from the `mminfo` output.
- `mediafile` specifies the starting file number of the save set that you obtained from the `mminfo` output.
- `mediarec` specifies the starting file record number of the save set that you obtained from the `mminfo` output.
- `device` is the name of the device that contains the volume. is the name of the device the volume is loaded in, for example `/dev/rmt0.1` or `\.\Tape0`
- `path` is the path on the NetWorker host that contains the `uasm` binary. For example, on Windows:

```
C:\Program Files\EMC NetWorker\nsr\bin
```

Scanner command examples

Recovering a single file to a different location on Windows

To recover a single file in the save set on Windows to a different location, type:

```
scanner -v -S ssid -f mediafile -r mediarec device | path\uasm -rv -m
source_dir=dest_dir filename
```

where:

- *source_dir* is the directory where the data resided during the backup.
- *dest_dir* is the directory where the data is relocated during the recovery.
- *filename* is the name of the file or directory to recover.

Recover a complete save set on UNIX

To recover all files in a save set on UNIX, type:

```
scanner -v -S ssid -f mediafile -r mediarec device -x path/uasm -rv
```

Recovering a single file to a different location on UNIX

To recover a single file in the save set on UNIX and to a different location, type:

```
scanner -v -S ssid -f mediafile device -x path/uasm -rv -m
source_dir=dest_dir filename
```

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `scanner` program.

VSS File Level Recovery

VSS File Level Recovery (FLR) provides the ability to browse, select and restore any System State file from the backup of the volume where it resides. There are changes to how Windows VSS-based backups and restores behave. The major changes include:

- System state files are now backed up as part of the volumes where they reside.
- All file system backups require that all system writers affected by the backed up volumes be included to ensure the backups are VSS consistent. You can use the command line flag `VSS:*=off`, to remove this VSS requirement.
- The Exclude file list specified by system state writers, and directives specified by unsupported application writers continue to work and are excluded from file system backups.

Recovering deduplication data

The *NetWorker Data Domain Boost Integration Guide* provides more information on how to recover deduplication data.

vProxy recovery in NMC

You can use the **Recovery** wizard in NMC to perform image level recovery, which allows you to recover full virtual machines and VMDKs. You can also use the **Recovery** wizard to perform file-level restore from a primary or cloned backup on a Data Domain device, but only as an administrator.

In NMC's **NetWorker Administration** window, click **Recover**. From the **Recover** window, launch the **Recovery** wizard by selecting **Recover > New**.

Entering management credentials for the Data Domain resource (instant recovery and User mode file-level restore only)


Before you perform an instant recovery of a virtual machine or file-level restore (User mode), ensure that you provide the management credentials for the Data Domain resource. For instant recovery, these credentials are required when performing the recovery using the NMC **Recover** wizard or the **VM Backup and Recovery** interface in the **vSphere Web Client**.

Procedure

1. In the NMC Administration window, click **Devices**.
The **Devices** window displays.
2. In the expanded left navigation pane, select **Data Domain Systems**.
3. In the right details pane, right-click the Data Domain system, and then select **Properties**.

The **NSR Data Domain Properties** window displays.

Figure 54 NSR Data Domain Properties



4. In the **Access** pane, type the management credentials.
 - a. In the **Management host** field, specify the hostname of the Data Domain system that is used for management commands.
 - b. In the **Management user** field, specify the username for a Data Domain user that has admin access. For example, sysadmin. The Management user should have Data Domain administrator privileges.
 - c. In the **Management password** field, specify the password of the management user.

- d. In the **Management port** field, specify the management port. By default, the port is 3009.

Note: The *NetWorker Data Domain Boost Integration Guide* provides information about the Cloud unit field and use of the Cloud tier device.

5. If required, in the **Configuration** pane, update the export path. It is recommended that you leave this field blank, which sets the export path to the default path. The short name of the NetWorker server is the default path.

If you do type a path in this field, ensure that the path has NFS permissions. When you log in to the Data Domain resource, browse to the NFS section and add the Mtree device path (the path to the NetWorker backup device) as a valid NFS path.

6. To save the changes, click **OK**.

File-level restore as an Active Directory user

In order to perform a file-level restore as an Active Directory (AD) user in the NMC **NetWorker Administration** window's **Recovery** wizard or the **Dell EMC Data Protection Restore Client**, you need to register a tenant user and provide the FLR AD user required permissions.

About this task

Note: The following steps include high level information for running `authc_config`. More detailed steps might be required if configuring AD authentication in the NetWorker environment. The *NetWorker Security Configuration Guide* provides more information.

Procedure

1. Log in to your AD server, and make note of objects related to this user (such as the Organizational Units (OU) and Groups in the Active Directory server) that will be added as a tenant user of NetWorker. You can use any third party AD viewer/browser to obtain the required objects and their properties.

For example,

- a. Create an OU **proxy** inside the domain.
- b. Inside the OU **proxy**, create a group named **vmware** and another OU **user**.
- c. Inside the OU **user**, create a user named **ADuser** and make this user a member of the group **vmware**.
- d. If you plan to use the command line to register the AD user to NetWorker, as described in step two below, make note of the values for the following variables as you will require these values for the registration commands:
 - config-name
 - config-server-address
 - config-domain
 - config-user-dn
 - config-user-dn-password
 - config-user-object-class
 - config-user-search-path

- config-user-id-attr
- config-group-search-path
- config-group-name-attr
- config-group-object-class
- config-group-member-attr
- config-active-directory

e. If you plan to use the **NetWorker Management Console** to register the AD user to NetWorker, also described in step two below, make note of the following values as you will require these values for the registration:

- Domain
- Port number
- Provider Server name
- User DN
- Group Object Class
- Group Search Path
- Group Name Attribute
- Group Member Attribute
- User Object Class
- User Search Path
- User ID Attribute

2. Register the AD domain user to NetWorker by using the command line.

a. Create a tenant user on NetWorker by running the `authc_config` command, as in the following:

```
authc_config -u administrator -e add-tenant -D tenant-  
name=your tenant name -D tenant-alias your selected aliases -p  
password
```

For example, to create a tenant user `ADuser` with the alias `FLR`, run

```
authc_config -u administrator -e add-tenant -D tenant-  
name=ADuser -D tenant-alias FLR -p password
```

b. Obtain the tenant ID by running the `authc_config` command using the `find-tenant` parameter. For example:

```
authc_config -u administrator -e find-tenant -D  
tenantname=ADuser -p password
```

c. Register the AD domain user to NetWorker by running the `authc_config` command using the `add-config` parameter and using the values obtained in Step1d, as in the following:

```
authc_config -u administrator -e add-config -D config-tenant-  
id=tenant ID number -D config-name=your tenant name -D config-  
server-address=ldap IPv4 addressOU=proxy,DC=domain name,DC=com  
-D config-domain=domain name -D config-user-  
dn=CN=ADuser,OU=user,OU=proxy,DC=domain name,DC=com -D config-  
user-dn-password=password -D config-user-object-  
class=inetOrgPerson -D config-user-search-path=OU=user -D
```

```
config-user-id-attr=cn -D config-group-search-path=OU=user -D
config-group-name-attr=cn -D config-group-object-class=group -
D config-group-member-attr=member -D config-active-directory=y
-p password
```

3. In the **NetWorker Management Console**, click the **Server** tab.
4. On the **Server** window, select **User Groups**.
5. Add a user group (for example, **ADuser group**) with the following permissions:
 - View Security Settings
 - View Application Settings
 - Remote Access All Clients
 - Operate NetWorker
 - Monitor NetWorker
 - Operate Devices and Jukeboxes
 - Recover Local Data
 - Recover Remote Data
 - Backup Local Data
6. Edit the new user to add the required **AD user** and **AD group** in the **External Roles** field. For example, for a user named **ADuser** with the domain **rideblr**, add the following in the **External Roles** field:


```
CN=Aduser,OU=user,OU=proxy,DC=rideblr,DC=com
CN=vmware,OU=proxy,DC=rideblr,DC=com
```
7. Log in to the **Dell EMC Data Protection Restore Client** as the AD user, in the format `<tenant>\<domain>\<userid>`. For example, `default\rideblr.com\ADuser`.


Results

You can now perform file-level restore as an Active Directory user.

Recovering a virtual machine using the NMC Recovery wizard

About this task


When you click **Recover** in NMC's **NetWorker Administration** window and select **Recover > New** from the menu, the **Recovery** wizard launches. **Virtual Machine Recovery** is the second recovery type displayed.

Figure 55 Virtual machine recovery in the NMC Recovery wizard

After selecting the **Virtual Machine Recovery** type, you can perform recovery of individual virtual machines, or (for revert and virtual machine recovery options) recovery from multiple virtual machines.

Procedure

1. In the **Select the Recovery Type** page, select **Virtual Machine Recovery**, and then select a vCenter server to recover from using the **Source vCenter server** drop-down. Click **Next**.
2. In the **Select the Virtual Machine to Recover** page, enter the name of the source virtual machine(s) to recover from, or perform a search for the virtual machine. Additionally, you can use the tabs on this page to choose a single virtual machine or multiple virtual machines from a selected backup, or browse the source vCenter to determine the required virtual machine source. When you locate and choose the desired virtual machine(s), click **Next**.

Figure 56 Select the Virtual Machine to Recover

3. In the **Select the Target Backups** page, select the virtual machine backup(s) you want to restore from the **Available Backups** pane. This pane lists both primary backups and, if available, clone copies. If you selected recovery from multiple virtual machines, you can switch between virtual machines to browse each machine's available backups by using the **Virtual Machine Name** drop-down. Click **Next**.

Figure 57 Select the Target Backup (individual virtual machine)

Available Backups

Name: obr-16

Backup Date	Location	DataStore	Size	Data Domain Copies	Other Media Copies
3/15/16 8:26:30 AM	/HPRMSPH_PVIM-back...	vmx5500-2-jun38	69 GB	1	0
3/15/16 1:49:24 AM	/HPRMSPH_PVIM-back...	vmx5500-2-jun38	69 GB	1	0

Figure 58 Select the Target Backup (multiple virtual machines)

Available Backups

Name: W2K8R2-HM11-2

Backup Date	Location	DataStore	Size	Data Domain Copies	Other Media Copies
6/1/16 12:37:25 AM	New Datacenter/SA...	vsanDatastore, vsanD...	26 GB	1	0
6/10/16 9:00:08 PM	New Datacenter/SA...	vsanDatastore	20 GB	1	0
6/9/16 9:00:10 PM	New Datacenter/SA...	vsanDatastore	20 GB	1	0
6/8/16 11:23:07 PM	New Datacenter/SA...	vsanDatastore	20 GB	1	0
6/8/16 11:02:18 PM	New Datacenter/SA...	vsanDatastore	20 GB	1	0
6/8/16 10:24:37 PM	New Datacenter/SA...	vsanDatastore	20 GB	1	0
6/8/16 9:55:23 PM	New Datacenter/SA...	vsanDatastore	20 GB	1	0

4. In the **Select the Virtual Machine Recovery method** page, select from one of the available recovery options:

- Revert (or rollback) a virtual machine
- Instant Recovery of a virtual machine (direct restore from a Data Domain device)
- Virtual Machine recovery (recovery to a new virtual machine)
- Virtual Disk recovery (recover VMDKs to an existing virtual machine)
- Emergency recovery (recovery to an ESX host)
- File Level recovery (recover files from VMDKs to a file system, or as a download).

Figure 59 Select the Virtual Machine Recovery method

Recover Configuration

Select the Virtual Machine Recovery Method

Select the recovery method to restore data from the select virtual machine.

Recovery Method

- ☐ Revert a Virtual Machine
Revert the selected virtual machine back to a point-in-time. If CBT is enabled, then it will be used to move only the data that has changed.
- ☐ Instant Recovery
Use instant access to recover the selected virtual machine, as a new virtual machine.
- ☐ Virtual Machine Recovery
Recover the selected virtual machine, as a new virtual machine.
- ☐ Virtual Disk Recovery
Recover one or more disks to an existing virtual machine.
- ☐ Emergency Recovery
Recover the selected virtual machine to an ESX host.
- ☒ File Level Recovery
Recover individual files and folders back to the same or a different virtual machine.

< Back | Next > | Close

Results

Subsequent wizard options change based on the recovery option selected, as described in the following sections.

Revert (or rollback) a virtual machine backup

The first virtual machine recovery option available in the NMC Recovery wizard is to revert, or rollback, a virtual machine backup. With a Revert a virtual machine backup recovery, you use an existing virtual machine to rollback the VMDKs as a virtual machine.

About this task

Note: When you revert a virtual machine, the current virtual machine is removed in the process. You cannot use the **Revert a Virtual Machine** recovery option when the ESXi has been removed from the vCenter and then added back to the vCenter. In this case, use the **Virtual Machine recovery** option instead.

To complete the Recovery wizard with the reverting a virtual machine method, perform the following.

Procedure

1. In the **Select the Virtual Machine Recovery Method** page:
 - a. Select **Revert a Virtual Machine**.
 - b. Click **Next**.

The **Select Options to Revert a Virtual Machine** page displays


2. In the **Revert Type** pane of the **Select Options to Revert a Virtual Machine** page:
 - a. Select **Revert both VM configuration and data** to revert both the configuration information (such as operating system, virtual machine size) and data for a virtual machine. When you select this revert type, the **Delete existing disk on disk configuration mismatch** option appears in the **Revert Options** pane to allow you to overwrite an existing disk if a configuration mismatch occurs.
 - b. Select **Revert VM Data Only** to revert only the virtual machine data without changing the virtual machine configuration.
3. In the **Revert Options** pane of the **Select Options to Revert a Virtual Machine** page, choose from the following options
 - a. Select **Revert all disks on this virtual machine** to rollback all VMDKs, or select **Revert one or more disks only** and then select a specific disk drive to rollback only that disk.
 - b. Select the **Power on virtual machine** checkbox to power on the virtual machine after the restore.
 - c. Select **Delete existing disk on disk configuration mismatch** if you want to be presented with the option of deleting the existing disk if a disk configuration mismatch is detected. Note that this option only appears when you select the **Revert both VM configuration and data** revert type in step two.
 - d. Click **Next**.

Note: If the virtual machine is currently powered on, a dialog displays requesting confirmation to power off the virtual machine. Additionally, if a change has occurred in the virtual machine configuration since the backup, a warning message displays.

Figure 60 Choose Disks to Revert

Note: The entire VMDK will be rolled back unless you have CBT enabled, in which case only the changed blocks will be moved.

4. In the **Select Alternate Recovery Sources** page:
 - a. Select the original backup or a clone copy if one is available.
 - b. If recovering from a clone that is not on a Data Domain device, or recovering from a Data Domain Cloud Tier device, specify the DD Boost clone pool.
 - c. Click **Next**.

Figure 61 Select Alternate Recovery Sources

5. In the **Perform the Recovery** page:

- a. Specify a name for the recovery and check the summary at the bottom of the page to ensure all the details are correct.
- b. Click **Run Recovery**.

Results

The **Check the Recovery Results** page will display the duration of the recovery, and a log file entry when the reversion is complete.

Instant Recovery of a virtual machine

The next virtual machine recovery option available in the NMC Recovery wizard is instant recovery of a virtual machine backup. With instant recovery, the virtual machine backup is read directly from the Data Domain device and the VMDKs will be restored directly on a Data Domain device. You can perform one instant recovery session at a time.

Before you begin

Before you begin, make note of the following:

- For the Data Domain resource, ensure that you provide the management credentials and, if required, enter the export path appropriately.
- Ensure that the free space on the Data Domain system is equal to or greater than the total disk size of the virtual machine being restored, as the restore does not take into account the actual space required after deduplication occurs. If there is insufficient disk space, an error appears indicating "Insufficient disk space on datastore," and creation of the target virtual machine fails.
- Ensure that you have at least one proxy that is not restricted to a specific datastore. For the vProxy, select **Properties** and then select **Configuration**, and verify that datastores is left blank.
- Do not perform an instant recovery of virtual machines in resource pools and other similar containers that are part of a currently running protection group.

About this task


To complete the Recovery wizard with the instant recovery method, perform the following steps:

Procedure

1. In the **Select the Virtual Machine Recovery Method** page:
 - a. Select **Instant Recovery**.
 - b. Click **Next**.
2. In the **Configure the Instant Recovery Options** page:
 - a. Select the location where you want to restore the virtual machine in the vCenter environment.

This does not have to be the original location, and can also be on a different vCenter server.
 - b. Ensure that you select the **Power on virtual machine** and **Reconnect to network** options.
 - c. Click **Next**.

Figure 62 Configure the Instant Recovery



3. In the **Select Alternate Recovery Sources** page:
 - a. Select the original backup, or a clone copy if one is available.
 - b. If recovering from a clone that is not on a Data Domain device, or recovering from a Data Domain Cloud Tier device, specify the DD Boost clone pool.
 - c. Click **Next**.
4. In the **Perform the Recovery** page:
 - a. Specify a name for the recovery.
 - b. Check the summary at the bottom of the page to ensure all the details are correct.
 - c. Click **Run Recovery**.

Results

The **Check the Recovery Results** page will display the duration of the recovery, and a log file entry when the instant recovery is complete. When the instant recovery is complete and ready for use, you can then storage vMotion the virtual machine to a datastore, or perform a file level recovery to the target file system, and then stop the completed instant recovery to free up those resources.

To stop an instant recovery in NMC:

1. Navigate to the **Recover** window.
2. Right-click the entry for the recovery within the Recover sessions pane.
3. Select **Stop** from the drop-down.

Note: To optimize use of NetWorker and Data Domain resources, it is strongly recommended that you stop the instant recovery session once you satisfy your recovery objectives.

Virtual machine recovery

The next virtual machine recovery option available in the NMC Recovery wizard is to perform a recovery of a virtual machine backed up with the vProxy Appliance to a new virtual machine.

About this task

Note: Recoveries of virtual machines backed up with the VMware Backup Appliance should still be performed with the **EMC Backup and Recovery** user interface in the **vSphere Web Client**.

To complete the Recovery wizard with the virtual machine recovery method, perform the following.

Procedure

1. In the **Select the Virtual Machine Recovery Method** page:
 - a. Select **Virtual Machine Recovery**.
 - b. Click **Next**.
2. In the **Configure the Virtual Machine Recovery** page, select the location where you want to restore the virtual machine in the vCenter environment
 - a. In the **Destination** pane, select the option to recover the new virtual machine to the original location, or browse to select a new location on the same vCenter server or a different vCenter server.
 - b. In the **Recovery Options** pane, choose a vProxy for the virtual machine recovery from the **Select vProxy** drop-down, specify the name of the new virtual machine, and then optionally select the virtual machine file datastore and folder where you want to recover the files. You can recover the virtual machine to a Blue folder by using the **VM Folder** drop-down, as shown in the following figure. The folder can be the default folder, or a new folder.

Figure 63 Configure the virtual machine recovery

Recover Configuration

Configure the Virtual Machine Recovery

Select the destination for the backup data, the recovery options, and the data to recover. In the Recovery Data table, you can edit the information in the Data Store column.

Destination

Recover the virtual machine to the original location on the vCenter server
 /R176161.lss.emc.com/MahatCluster01/Vml10_1.vms

☐ Browse the vcenter server to select a recovery location

vCenter: /R176161.lss.emc.com

Recovery Options

Name: Vml10_1 Power on virtual machine: ☐

Virtual Machine Files: datastore_7676 Reconnect to network: ☒

VM Folder: vm (Path: /vm) Select vProxy:

Recovery Data

Recover	Disk Name	Datastore
☒	Hard disk 1	datastore_7676

< Back Next > Close

If you have a single disks, or multiple disks with multiple datastores, you can perform the following:

- Choose to recover a collection of all the available hard drives.
- Select a different datastore than the original datastore.
- Select a different datatore for each disk you want to recover.
- Specify the datastore where the virtual machine configuration files reside.

Optionally, select the **Power on virtual machine** and **Reconnect to network** options to power on and reconnect after the recovery, and then click **Next**.

3. In the **Select Alternate Recovery Sources** page:
 - a. Select the original virtual machine backup, or a clone copy if one is available.
 - b. If recovering from a clone that is not on a Data Domain device, or recovering from a Data Domain Cloud Tier device, specify the staging pool.
 - c. Click **Next**.

Note: If selecting a clone from **Select Alternate Recovery Sources**, additionally review the "Selecting alternate recovery sources" section.

4. In the **Perform the Recovery** page:
 - a. Specify a name for the recovery and check the summary at the bottom of the page to ensure all the details are correct.
 - b. Click **Run Recovery**.

Results

The **Check the Recovery Results** page will display the duration of the recovery, and a log file entry when the virtual machine recovery is complete.

Virtual Disk Recovery

The next virtual machine recovery option available in the NMC Recovery wizard is to perform a virtual disk, or VMDK, recovery. With VMDK recovery, the disks from the virtual machine backup are recovered to an existing virtual machine.

About this task

To complete the Recovery wizard with the virtual disk recovery method, perform the following.

Procedure

1. In the **Select the Virtual Machine Recovery Method** page:
 - a. Select **Virtual Disk Recovery**.
 - b. Click **Next**.
2. In the **Configure the Virtual Disk Recovery** page:
 - a. Select the virtual machine where you want to restore the VMDKs. This can be the original virtual machine, or another existing virtual machine.
 - b. Select the desired disks from the **Recovery Data** pane, and select a datastore.
 - c. Click **Next**.

Figure 64 Configure the Virtual Disk Recovery

Recover	Disk Name	Datastore
☒	Hard disk 1	datastore1 (2)
☒	Hard disk 2	datastore1 (2)

3. In the **Select Alternate Recovery Sources** page:
 - a. Select the original virtual disk backup, or a clone copy if one is available.
 - b. If recovering from a clone that is not on a Data Domain device, or recovering from a Data Domain Cloud Tier device, specify the staging pool.

- c. Click **Next**.
4. In the **Perform the Recovery** page:
 - a. Specify a name for the recovery.
 - b. Check the summary at the bottom of the page to ensure all the details are correct.
 - c. Click **Run Recovery**.

Results

The **Check the Recovery Results** page will display the duration of the recovery, and a log file entry when the disk recovery is complete.

 **Note:** When you start a VMDK recovery, the virtual machine will be powered off automatically without issuing a warning message.

Emergency Recovery

The next virtual machine recovery option available in the NMC Recovery wizard is an Emergency Recovery. An Emergency Recovery is required when you need to restore the virtual machine to an ESXi host.

Before you begin

Emergency Recovery requires a vProxy set up on the ESXi host prior to running the recovery.

Additionally, ensure that you disconnect the ESXi host from the vCenter server.

About this task

 **Note:** During an Emergency Recovery, the vProxy gets associated with the ESXi host and is unavailable for other operations on the vCenter server. Wait until the recovery completes before initiating any other operations on the vProxy.

To complete the Recovery wizard with the Emergency Recovery method, perform the following:

Procedure

1. In the **Select the Virtual Machine Recovery Method** page:
 - a. Select **Emergency Recovery**.
 - b. Click **Next**.
2. In the **Configure the Emergency Recovery** page:
 - a. Specify the target ESXi server in the vCenter environment.
 - b. Click **Connect**.

Figure 65 Configure the Emergency Recovery

Target ESX Host

ESX Server: 10.31.183.205
 Username: root
 Password: *****
 Connect

Recovery Options

Power on virtual machine: ☐
 Reconnect to network: ☒
 Name: RHEL
 Virtual Machine Files: datastore1

Proxy Selection

☐ Take control of an existing proxy already deployed within the vCenter
 Existing Proxy:
☒ Use a proxy already deployed in the ESX host but not registered with NetWorker
 Proxy Hostname:
 Username: admin
 Password:
 ⚠ This proxy will be un-registered after recovery.
☐ Invoke the proxy deployment manager to deploy and register a proxy for this ESX host
 Launch Proxy Deployment Manager
 ⚠ This proxy will be un-registered after recovery.

Recovery Data

Disk Name	Datastore
Hard disk 1	datastore1
Hard disk 2	datastore1

< Back Next > Close

The **Proxy Selection** and **Recovery Data** panes get populated with the ESXi server details.

3. In the **Proxy Selection** pane, if a proxy is not discovered, add a new proxy which is deployed in vCenter but not added to NetWorker.
4. For the disks in the **Recovery Data** pane:
 - a. Select a datastore.
 - b. Optionally, select the **Power on virtual machine** and **Reconnect to network** options.
 - c. Click **Next**.
5. In the **Select Alternate Recovery Sources** page:
 - a. Select the original disk backup, or a clone copy if one is available.
 - b. If recovering from a clone that is not on a Data Domain device, or recovering from a Data Domain Cloud Tier device, specify the staging pool.
6. In the **Perform the Recovery** page:
 - a. Specify a name for the recovery and check the summary at the bottom of the page to ensure all the details are correct.
 - b. Click **Run Recovery**.

Results

The **Check the Recovery Results** page will display a progress bar with the duration of the recovery, and a log file entry when the Emergency Recovery is complete.

Note: The progress bar may not update correctly when you perform an Emergency Recovery directly to the ESXi host.

File Level recovery (Admin mode only)

The final virtual machine recovery option available in the NMC Recovery wizard is File Level recovery. With file level recovery, you can recover individual files from backups

of virtual machines or VMDKs to a primary or secondary vCenter server, and for application-consistent backups, you can also restore the transaction log from Data Domain to the SQL database.

Before you begin

NetWorker only supports file level recovery operations from a primary or cloned backup if the save set is on a Data Domain device. If a cloned backup does not exist on the Data Domain device, you must manually clone a save set from the tape device to Data Domain before launching the **Recovery** wizard.

For the Data Domain resource, ensure that you provide the management credentials and, if required, type the export path appropriately. The section [Entering management credentials for the Data Domain resource \(instant recovery and User mode file-level restore only\)](#) provides detailed steps.

Additionally, if recovering to a virtual machine on a secondary vCenter, ensure that a vProxy appliance has been deployed on the secondary vCenter server and configured with the NetWorker server.

About this task

File level recovery in the NMC **Recovery** wizard can only be performed by an administrator.

 **Note:** For file-level recovery of high-density file systems (more than few hundred files/folders), it is recommended to use either the **NetWorker Management Web UI** or the **Dell EMC Data Protection Restore Client** (User or Admin mode, as applicable) instead of the NMC **Recovery** wizard.

To complete the Recovery wizard with the file level recovery method, perform the following:

Procedure

1. In the **Select the Virtual Machine Recovery Method** page:
 - a. Select **File Level recovery**.
 - b. Click **Next**.
2. In the **Select Alternate Recovery Sources** page:
 - a. Select the primary backup to recover from, or select the **Recover the Virtual machine from a clone on a Data Domain device** option.
 - b. Select the clone copy that you want to recover files from.
 - c. Click **Next**.

 **Note:** If selecting a clone from **Select Alternate Recovery Sources**, additionally review the section "Selecting alternate recovery sources".

Figure 66 Select Alternate Recovery Sources for file level recovery

Recover Configuration

Select Alternate Recovery Sources

Select the source for the recover data. You can recover data from a backup, or from a cloned copy of the backup on a Data Domain device or a non-Data Domain device.

☒ **Recover the virtual machine from this backup**

DO System	Networker Device Name	Volume	Pool	Enabled
dd890-2.lss.emc.com	MyBackup	MyBackup.001	MyBackup	☒

☐ **Recover the virtual machine from a clone on a Data Domain device**

DO System	Networker Device Name	Volume	Pool	Enabled
dd890-2.lss.emc.com	MyClone	MyClone.001	MyClone	☒

< Back Next > Close

3. In the **Select the target Virtual Machine** page:

- a. Select the virtual machine that you want to recover the files to.

By default, the virtual machine that you selected for recovery in the **Select the Virtual Machine to Recover** page is displayed.

- b. To recover to another virtual machine in the vCenter, or recover to a virtual machine on a secondary vCenter, select **Browse the vCenter server to select a Virtual Machine to recover to**, and choose a vCenter from the drop-down to browse that vCenter's tree and select a different virtual machine.

- c. Click **Next**.

Note: Cross-platform recovery, for example from a Windows to a Linux virtual machine, is not supported.

4. In the **Mount The Saveset** page:

- a. Provide the username and password of the virtual machine where the files will be restored to.
- b. Click **Start Mount**.
- c. If performing file level recovery as a domain user, provide the AD user details —no operating system or local account is required if you have configured the AD/domain user.

Figure 67 Mount the save set for file level recovery

Mount the Saveset

Wait for the selected saveset to be mounted on the target virtual machine.

Mount Credentials

Provide credentials for the target virtual machine. If the FLR agent is not already installed it will be automatically installed, and these credential must have admin privileges. Select the 'Keep FLR agent after installing' checkbox to keep the FLR agent installed on the target virtual machine for future FLR recoveries.

Username:

Password:

☒ Keep FLR Agent after installing

Mount Results

Status: End of mount request processing; mountpoint is 'lopt/emc/vproxyra/fr/mountpoints/FLR859936944' (no error)

Duration: 00:00:48

Mount Log:


155854.nsvproxy_fr: Starting Saveset Mount...

155858.nsvproxy_fr: Mounting Saveset id '7db15042-00000006-fc11e330-5811e330-00075000-1ac8d56' on vCenter 'vm85-vct1.sas.emc.com' at VM 'vm85-nw1 - dev 8b424 (Linux/vm-2410)' using NetWorker server 'vm85-nw1.sas.emc.com' for user 'root'.

When the **Mount Results** pane shows that the mount has succeeded, click **Next**.

Note: This user should have privileges to install the **FLR Agent**, which is required to perform file level recovery. For Linux virtual machines, this requires the root user account or an equivalent sudo local user account, as described in the section "FLR Agent installation on Linux platforms" of [FLR Agent requirements](#).

5. In the **Select the Files and Folders to Recover** page:
 - a. Browse through the folder structure to select the files you want to recover.
 - b. Click **Next**.

Figure 68 Select the files and folders to recover

6. In the **Select the Restore Location** page:
 - a. Select the folder that you want to recover the files to, or create a folder.
 - b. Click **Next**.

Note: NetWorker does not currently support creating folders with spaces in the folder name.

7. In the **Perform the Recovery** page:
 - a. Specify a name for the recovery.
 - b. To ensure all the details are correct, check the summary at the bottom of the page
 - c. Click **Run Recovery**.

Results

The **Check the Recovery Results** page displays the duration of the recovery, and a log file entry when the file level recovery is complete.

Selecting alternate recovery sources in the NMC Recovery wizard

The NMC Recovery wizard contains a step for each virtual machine recovery method where you can select an alternate source to recover from, for example, a clone copy on a Data Domain or non-Data Domain device. If the primary source is present, it is recommended that you recover from the primary source. However, if both the primary source and clone copies are present and enabled and you want to recover from a clone copy, perform the following.

Procedure

1. In the **Select Alternative Recovery Sources** page, select the clone you want to recover from, either a clone on a Data Domain device or non-Data Domain device.

Additionally, make note of the name indicated in the **Volume** column for all of the volumes you do not want to recover from, as you will require this information in steps 5 and 6.

2. Click **Close** to display the **Save Progress** dialog, and then specify a name for the recover and click **Save** to save your progress.
3. In the NMC **Administration** window, click **Devices** to display the **Devices** window.
4. In the left navigation pane, select **Devices**. The list of devices displays in the right pane.
5. For each volume you do not want to recover from that you made note of in step 1, locate the corresponding device, and make note of that device name.
6. For each device you identify as corresponding with those volumes, right-click the device and select **Unmount** from the drop-down, and then also select **Disable** from the drop-down.

 **Note:** Ensure that no backups are currently running to these devices prior to unmounting.

7. In the NMC **Administration** window, click **Recover** to display the **Recover** window, and locate the saved recovery
8. Right-click the saved recovery and select **Open Recover**.

The Recovery wizard re-opens on the **Select Alternative Recovery Sources** page.

9. In the **Recovery Source** pane of the **Select Alternative Recovery Sources** page, select either **Recover the virtual machine from a clone on a Data Domain device**, or **Recover the virtual machine from a clone on a non-Data Domain device**. Click **Next**.

 **Note:** If you want to recover from a clone on a non-Data Domain device, manually change the staging pool to a different pool, and ensure that your selected pool does not already contain copies for this backup. If the primary source is present and you select a clone to recover from using the same staging pool that contains the existing copy, the recovery may become unresponsive.

10. In the **Perform the Recovery** page, specify a name for the recovery and check the summary at the bottom of the page to ensure all the details are correct. Click **Run Recovery**.

The **Check the Recovery Results** page will display the duration of the recovery, and a log file entry when the recovery is complete.

11. In the NMC **Administration** window, click **Devices** to return to the **Devices** window, and in the left navigation pane, select **Devices** to display the list of devices in the right pane.
12. For each device that you unmounted and disabled in step 6, right-click the device and select **Enable** from the drop-down, and then select **Mount** from the drop-down.

Monitoring and verifying Virtual Machine recoveries

After selecting Run Recovery to complete the Recovery wizard, there are multiple ways you can monitor the progress of the virtual machine recovery, and then verify when the recovery is complete.

NMC Recover and Monitoring windows

To monitor the progress of the virtual machine recovery, use the **Recover sessions** pane in the **Monitoring** window, or the **Currently Running** pane of the **Recover** window.

To verify that the virtual machine recovery is complete, use the **Configured Recovers** pane in the **Recover** window.

Check the Recovery results in the NMC Recovery wizard

The final step of the **Recovery** wizard also allows you to check the recovery results. Upon completion of the virtual machine recovery, an entry for the log file appears in the **Recovery log** pane. Click **Export log** to save and view the log file.

Recovery configuration information storage

When you create a recover configuration by using the Recovery wizard, NetWorker saves the configuration information in an NSR recover resource in the resource database of the NetWorker server. NetWorker uses the information in the NSR recover resource to perform the recover job operation.

When a recover job operation starts, NetWorker stores:

- Details about the job in the nsrjobsd database.
- Output sent to stderr and stdout in a recover log file. NetWorker creates one log file for each recover job.

NOTICE NetWorker removes the recover log file and the job information from the job database based on value of the *Jobsdb retention in hours* attribute in the properties of the NetWorker server resource. The default jobsdb retention is 72 hours.

NMC function to collect vProxy log bundle information

NetWorker 18.x features an NMC function to collect vProxy log bundle information from a virtual machine. To collect log bundle information, perform the following steps in NMC:

1. From NMC's **NetWorker Administration**, open the **Devices** window.
2. From the left pane, select **VMware Proxies** to display the virtual machine proxy devices.
3. Right-click the virtual machine that you want to collect log bundle information from, and then from the menu, click **Log Bundle**.
4. (Optional) Collect the recycled logs from the pop-up window selection.

NOTE: Since the temporary log bundle download occurs on the NetWorker server, ensure that there is sufficient space on the drive where the NetWorker server is installed.

Also, note that NMC cannot collect the log bundle when accessed from a remote machine that cannot communicate with vProxy

Recovering file system data on Windows

This section provides detailed information about how to recover Windows data without using BMR.

Recovering Windows volume mount points

A volume mount point (or mount point) is a disk volume that is grafted into the namespace of a host disk volume. This allows multiple disk volumes to be linked into a single directory tree, similar to the way DFS links network shares into a unified structure.

Assigning a drive letter to a mount point is optional. Many disk volumes can be linked into a single directory tree, with a single drive letter assigned to the root of the host volume.

Recovering mount points

Perform separate recovery operations to recover the mount point and the mounted volume's data.

About this task

 **NOTICE** The NetWorker Save Set Recovery feature does not support recovery of mount points. To recover mount points and their data, use these special procedures.

Procedure

1. Manually create the mountpoint, if it does not exist already.
2. Start the NetWorker User program and recover the data under the mount point.

[Using the NetWorker User program](#) on page 551 provides more information about performing data recoveries.

Recovering nested mount points

Procedure

1. When the mount points do not already exist, manually create the top-level mount point, then work down the hierarchy and create each successive mount point.
2. Start the NetWorker User program and recover the data under the mount points.

Recovering Windows DHCP and WINS databases

Use the following procedures to perform an offline recovery of the DHCP and WINS databases.

 **NOTICE** When you recover from a save set `ALL` backup, the recovery operation automatically recovers the DHCP and WINS, and these procedures are not required.

Recover a DHCP database

Procedure

1. Use the NetWorker User program to recover the %SystemRoot%\System32\dhcp directory.
2. Use the Microsoft DHCP administrative tools to restore the DHCP database. The Microsoft documentation provides detailed instructions about Microsoft DHCP administrative tools.

Recovering a WINS database

About this task

 **NOTICE** Microsoft documentation describes how to use the Microsoft WINS administrative tools to recover the databases.

Procedure

1. Use the NetWorker User program to recover the backup configured in the WINS backup procedure. [DHCP and WINS databases](#) on page 414 provides more information.
2. Use Microsoft WINS administrative tools to restore the WINS database.

Recovering DFS

Review this section for information about how to recover DFS.

DFS topology information

Domain-based DFS topology information is backed up as part of AD, which is a component of the `WINDOWS ROLES AND FEATURES` save set on domain controllers. Registry-based DFS topology information is backed up as part of the Windows registry, which is a component of the DFS host server's `WINDOWS ROLES AND FEATURES` save set.

Restoring a DFS

About this task

Restore DFSR through the `WINDOWS ROLES AND FEATURES` save set.

Procedure

1. Restore the DFS topology information:
 - To restore a domain-based system, restore the `WINDOWS ROLES AND FEATURES` save sets on the domain controller.
2. On the DFS host server:
 - a. Restore the DFS root.

 **Note:** You cannot restore individual DFS links. If the DFS root has lost a link, restore the entire DFS root in which that link resided.

- b. If required, restore any local DFS destination directories.

3. If required, restore the remote DFS destination directories.

Authoritative restores of DFS Replication writers

You must perform authoritative restores of the DFS Replication writers from the command line. Restores from the NetWorker User program GUI are not authoritative.

To perform an authoritative restore of the DFS Replication writer, use the `-U` option with the `recover` command.

The following examples assume that you have two DFSR shares, `E:\Share1` and `E:\Share2`.

- To restore all the DFSR shares (two shares in this example), type the following command:

```
recover -s server -U -N "WINDOWS ROLES AND FEATURES:\DFS
Replication service writer"
```

- To restore just one DFSR share (Share1 in this example), type the following command:

```
recover -s server -U -N "WINDOWS ROLES AND FEATURES:\DFS
Replication service writer:Share1"
```

Non-authoritative DFS Replication writer granular recovery

Windows Distributed File System Replication (DFSR) granular recovery is supported on Windows 2008 R2 SP1, Windows Server 2012, and Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019.

DFSR Shared Directories supports granular DFSR folder and file recoveries on computers that run Windows 2008 R2 SP1, Windows Server 2012 and later operating systems. You do not have to recover the entire `WINDOWS ROLES AND FEATURES` save sets to restore DFSR shared directories. If you perform a file level non-VSS granular recovery, then the recovered file is treated as new version of the file by DFS.

You must use volume backup to correctly back up a DFSR namespace. Also, namespaces are skipped when specifying the `ALL` save set. You must back up namespaces directly by specifying the path of the namespaces as separate save sets in the Save Set attribute.

For recovery of namespace data, use the NetWorker User program and select individual files or folders of the NetWorker Client resource.

Recovering data on OS-X clients

Use the `recover` command or the NetWorker Recover application to recover files on a OS-X host.

Recovering files and directories from the command prompt

Use the `recover` command to recover individual files and directories from the command prompt on an OS-X client.

About this task

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `recover` command.

Procedure

1. From the Mac OS-X Terminal application, type:

```
$ recover -s NetWorker_server
```

Note: If you do not specify the `-s NetWorker_server` option, the `save` command contacts the NetWorker server that is defined in the `/nsr/res/servers` file.

2. At the `recover` prompt:

- a. To browse the files and directories, use common UNIX shell commands such as `cd` and `ls`.
- b. To specify the files and directories that you want to recover, use the `add` command.

For example:

```
recover> add directory_name
```

- c. Optionally, to automatically overwrite existing files, use the `force` option at the `recover` prompt.
- d. To start the recovery operation, type `recover`:

```
recover> recover
```

NOTICE Do not recover any OS-X operating system start files. For example, do not recover the OS-X operating system kernel, `/mach_kernel`.

Recovering files and directories by using the NetWorker Recover GUI

Use the NetWorker Recover application to recover data from a NetWorker server.

Connecting to the NetWorker server


Perform the following steps on the OS-X client.

Procedure

1. Start the NetWorker Recover application.
2. Use NMC to connect to the NetWorker server.
 - When you start the NetWorker Recover GUI for the first time, the **Connect to Server** dialog appears. Specify the NetWorker server that contains the backup data for the client:
 - In the **Available Servers** field, select the NetWorker server, and click **Connect**.
The **Available Servers** field displays a list of host names that appear in the `/nsr/res/servers` file on the Mac client. To query the network for other NetWorker servers, click **Update**.
 - In the **Server Address** field, specify the hostname or IP address of the NetWorker server, and click **Connect**.
The following figure shows the **Connect to Server** dialog box.


Results

Figure 69 Connect to Server



- When you close the NetWorker Recover GUI, subsequent recover operations will connect to the last NetWorker server selected, by default. To change the NetWorker server, perform one of the following steps:
 - In the **SERVERS** section on the side bar, select the NetWorker server, then click **Connect**.
 - On the **Go** menu, select **Connect to Server**. The **Connect to Server** dialog box appears.

After you successfully connect to a NetWorker server, the **NetWorker Recover** window appears.

Figure 70 NetWorker Recover window

Changing the source NetWorker Client


After you connect to the NetWorker server, the browse view displays a list of files and folders that you can recover from the last local host backup.

About this task

NetWorker Recover provides you with the ability to recover of files from a host that is not the local host. [Directed recoveries](#) on page 546 provides detailed information about directed recovery requirements.

To change the source host, perform one of the following actions:

- From the **Go** menu, select **Browse Client**. A list of clients for the current NetWorker server appear in a drop down. To establish a browse session with a new host, select the source host from the drop down.
- On the side bar, in the **SERVERS** section, select the NetWorker server. The browse view displays a list of clients. To establish a browse session with a new host, double-click the source host. The following figure provides an example of browse session window after you select a NetWorker server from the **SERVERS** section.

Figure 71 List of clients available for a NetWorker server

Note: The **Clients** filter bar, located above the list of client names, enables you to filter the client list by operating system. For example, select **All** to show all clients of the NetWorker server, or select **OS-X Clients** to display OS-X hosts.

Changing the browse time

By default, the browser view displays files and directories from the last backup. To browse or recover files from an earlier backup, use one of the following methods to change the browse time:

About this task

- On the tool bar, select **Browse Time**. The **Browse Time** view appears, which displays the current browse time. Use the controls to specify a new date and time.
- From the **Go** menu, select **Browse Time**. Select one of the preconfigured options from the drop down. To use a calendar and clock to choose the date and time, select **Other**.

Selecting objects to recover and recovering the data

The NetWorker Recover feature support the ability to perform a browsable recovery or a save set recovery.


Procedure

- Display a list of file system objects in the browser view.
 - To perform a browsable recovery, on the side bar in the **Devices** section, select a file system. **NetWorker Recover** queries the client file index and displays the objects that you can recover.
- Note:** To show hidden files, from the **View** menu, select **Show Hidden Files**.
- To perform a save set recover, on the side bar in the **SAVE SETS** section, select a save set. **NetWorker Recover** queries the media database and displays each instance of the save set, including cloned save sets.

Note: The **Save Sets** filter bar, located above the list of save sets enables you to filter the save set list by save set type. For example, to show all the original save set instance, select **Save Sets** or to display cloned save set instances, select **Cloned Save Sets**.

2. To search browser view for the files you want to recover:
 - a. Type the text string in the **Search** field in the upper right of the **NetWorker Recover** window.
 - b. Use the **Search Scope** bar to narrow the scope of the search result. The following figure displays some of the search criteria you can use.

Figure 72 Search browse view




When you select an object in the **Search Result** view, NetWorker Recovery displays the path to the object in the **Status** bar at the bottom of the browser view.

3. To display information about an object, right-click the object, and select **Get Info**.
4. To mark objects in the browser view for recovery, select the checkbox next to each object that you want to recover. You can only mark one save set or clone instance at a time.

NetWorker Recover adds each item that you mark to the **RECOVERY SETS** section on the side bar. A number appears next to each recovery set in the sidebar, which represents the total number of items that are selected for recovery.

5. To view or select different versions of a marked file, perform the following steps:
 - a. Right-click the file and select **File Versions**. The **Versions** side bar appears. The following figure provides an example of the **Versions** side bar.

Figure 73 Versions side bar



- b. To recover a specific version of a file, perform one of the following actions:
 - Drag and drop the file from the **Versions** side bar to the browser view.
 - Drag the file to a folder for recovery.
 - Right-click the file to select **Mark for recovery**.
6. To review a summary list of the marked files, in the **RECOVERY SETS** section on the side bar, perform one of the following actions:
 - Select **Files** to display a list of objects that you marked for a browsable recovery.
 - Select **Save Sets** to display a list of objects that you marked for a save set recovery.

The **Recover Files** browse view displays a list of marked files and the list of volumes that the recovery operation requires.
7. To view the status of the required volumes, click **Volume Status**. Ensure that the status of the required volume indicates online, then close the dialog box.
8. To start the recover operation, click the **Recover** button in the toolbar. The **Recover** window appears.
9. In the **Recover** window, select the recovery options.
 - To recover the objects to a directory that differs from the original location, perform one of the following actions:
 - In the **Relocate files to** field, type the path on the destination host to recover the data.
 - Click **Browse** and select the target directory.
 - Select a conflict resolution option:
 - Rename the recovered file— By default, the recover operation appends a tilde (~) to the beginning of the name of the recovered file *~file_name*. When a file named *~file_name* already exists, the recovered file is renamed *~00_file_name*, and so forth, to *~99_file_name*. When this fails, the recover process does not automatically rename the file and prompts the user is to specify a name for the file.
 - Discard recovered file— Discards the recovered file and keeps the existing file.
 - Replace local file— Replaces the file on the file system with the recovered version.

- Prompt me for an action— Each time the recovery operation encounters a file or folder with the same name in the destination location, the recovery operation prompts you to select a conflict resolution method.
 - To recover the files to a different host, select the hostname from the **Direct recover to** drop down.
 - Click **OK**. The recover status dialog box appears. At any time during the recovery, you can click the **Stop** button to cancel the operation.
10. To monitor the recovery process, on the **Recover progress**, select **Monitor Server**.
- The **NetWorker Monitor** dialog box appears with the following tabs:
- Info— Displays general server information including name, IP, OS type, NetWorker version, Save totals, and Recover totals.
 - Messages— Displays server messages that are logged during the recovery, for example, errors and warnings.
 - Devices— Displays the status for all connected devices.
 - Sessions— Displays Save sessions, Recover sessions, and Browse sessions.
 - Settings— Allows you to adjust the polling interval for server updates.
11. To review the recover log, after the recovery operation completes, select **Recover Log**. The **Console** application appears and displays the contents of the `~/Library/Logs/recover.log` file.

Recovering client files on a different NetWorker server

You can use a NetWorker server, which differs from the original NetWorker server to recover data for a client.

Before you begin

Determine the pool names that were used to write the client data to the media on the original NetWorker server.

About this task

To use a different NetWorker server to recover client data, you must perform the following tasks on the NetWorker server:

- Create a Client resource with the same client ID that the original NetWorker server associated with the client name.
- Create each Pool resource that was used to write the client data to a volume.
- Use the `scanner` command to repopulate the media database and client file indexes with save set information for the client.

Procedure

1. Determine the **Client ID** value of the NetWorker client on the original server:
 - a. On the **Administration** window, click **Protection**.
 - b. In the left pane, click **Clients**.
 - c. In the right pane, right-click the client, and then select **Properties**.
 - d. On the **Globals (1 of 2)** tab, make note of the value in **Client ID** attribute, then click **Cancel** to close the **Properties** window.

2. On the new NetWorker server, create a client:

- a. In the **Name** attribute, type a name for the client.

You can use the same name that was used on the original server, but you cannot use a name that exists for the new server. When a client with the same name exists on the new server, use this format to specify the client name:

`~hostname-#`

where *hostname* is the hostname of the client.

For example, if the client's hostname is jupiter, and a client named jupiter already exists on the new server, type:

`~jupiter-1`

3. On the new NetWorker server, create each Pool resource that was used to write the client data on the original NetWorker server.

 **Note:** Ensure that you create each Pool resource with the same name that you used on the original NetWorker server.

4. Use the `scanner` command to import the save set information into the new NetWorker server.

- To import the save set information into the client file index and media database entries, type the following command:

```
scanner -i -c client_name device_name
```

where *client_name* is the name of the client that appears on the original NetWorker server.

- To import the save set information into the media database only, type the following command: `scanner -m -c client_name device_name` where *client_name* is the name of the client that appears on the original NetWorker server.

 **NOTICE** When you use the `scanner -i` or `scanner -m` to import data before you configure the Client resource on the new server:

- Only the media database contains the client ID and save set information for the imported save sets.
- If the same hostname already exists on the NetWorker server, NetWorker will not use the original hostname to store the save set information because the client ID is different. NetWorker associates the save set information with a hostname in the format *clientname-#*.
- You must create a Client resource with the name *clientname-#* and specify the client ID that you recorded from the original NetWorker server.
- Optionally, after you create the new Client resource, run the `scanner -i` command to store the save set information into the client file index. When you use the `scanner` command, specify the client name as it appears on the original NetWorker server.

Recovering as NMC User with Non-Admin Privileges

NMC user performs recovery operation without any admin or security privileges. Recovery operation is performed with the following privileges:

- Recover remote data
- Recover local data

The NMC user performing the recovery through NMC, REST API and NWUI, does not need to be granted Application administrator or Security administrator privileges. The NMC non-admin user cannot delete the existing recovery tasks, however they can reuse the same recovery task to perform the recovery operation without admin/semi admin privileges. The NMC non-admin user cannot delete or add resources but they can use existing resources. The following are the applications that are supported for recovery operation as non-admin NMC user:

- Recovery through NMC
 - Filesystem
 - Oracle
 - DB2
 - MYSQL
 - Lotus
 - SYB (Skybase)
 - SAPO (SAP-Oracle)
 - VMware
- Recovery through REST API
 - Filesystem
 - VMware
- Recovery through NMUI
 - VMware

Exception: As a non-admin NMC user, user cannot add or delete any resources. Therefore the following recovery operation does not work:

- Creating new remote recovery job
- VMware emergency recovery

Note: If the non-admin user performs remote recovery, the administrator should configure the remote access list attribute in client properties and provide view security privilege to the non-admin user. For non-admin users, a module recovery that requires Lockbox operation, needs Configure NetWorker privileges that the admin should grant.

Preparing the NMC Sever for Recovery

The following sections outline how to prepare the NMC server to take recovery of the data to the original host or a different host, from a point-in-time backup or the last

backup time, through NMC, NWUI, REST-API with only recovery privileges to NMC user.

Creating an NMC Group for Recovery

Use the NMC GUI to create NMC user group resources:

Procedure

1. Connect to the NMC server with a NetWorker Authentication Service administrator account.
2. Click **Setup**.

The **Users and Roles** window appears.

3. In the left navigation pane, select **Users and Roles > Groups**.
4. In the **Groups** window, right-click the **Groups**, and then select **New**.

The **Create Group** window appears.

5. In the **Name** attribute, type a name for the NMC user group.

 **Note:** You cannot modify the name of an existing NMC group.

6. Click **OK**.

Creating an NMC User with recovery privileges

Use the NMC GUI to create NMC user resources.

About this task

Perform the following steps to create an NMC user with recovery privileges:

Procedure

1. Connect to the NMC server with a NetWorker Authentication Service administrator account.
2. Click **Setup**.

The **Users and Roles** window appears.

3. In the left navigation pane, select **Users and Roles > Groups**.
4. In the **Users** window, right click the **Users**, and then select **New**.

The **Create User** window appears.

5. In the **User Name** attribute, type a name for the user.

 **Note:** You cannot modify the name of an existing user.

6. In the **Groups** attribute, select the group created for recovery.
7. Provide password.
8. Click **OK**.

 **Note:** Perform recovery operation as NMC user through NMC GUI, REST API, and NWUI.

Creating a NetWorker User Group for recovery

Use the NMC GUI to create user group resources.

About this task

Perform the following steps to create a NetWorker user group for recovery:

Procedure

1. Connect to the NMC server with the NetWorker Authentication Service administrator account.
2. Click **Enterprise**.
3. Right-click the NetWorker server and select **Launch Application**.
4. On the NetWorker **Administration** window, select **Servers**.
5. In the left navigation pane, select **User Groups**.
6. In the **User Groups** window, right-click and then select **New**.

The **Create User Group** window appears.

7. In the **Name** attribute, type a name for the user group.

 **Note:** You can not modify the name of an existing user.

8. In the **External Roles** attribute, specify the users and groups.
9. In the **Privileges** attribute, select the following privileges to assign to the user group:
 - Recover remote data
 - Recover local data
 - Monitor NetWorker

 **Note:** Privilege Monitor NetWorker is basic for NMC user.

10. Provide the password and click **OK**.

Recover the NMC Server database

The NMC Server database contains management data such as report information. You can recover the NMC Server databases to the original NMC host or to a new NMC host.

Before you can perform a NMC Server database recovery, you must have an NMC Server database backup.

An NMC backup contains the following components:

- NMC database files
- NMC database credential file (`gstd_db.conf`)
- NMC lockbox files
- Legacy authentication configuration files

The "NMC Server management" chapter provides more information about NMC Server database backups.

Prepare for an NMC Server recovery

Before you recover an NMC Server, review the following information.

- If required, install the operating system on the target NMC Server.
 **Note:** To recover an NMC Server from one host to another, both hosts must run on the same operating system.
- If required, install the NetWorker and NMC Server software on the target host. When you are prompted to specify the NetWorker Authentication Service host, specify the NetWorker Authentication Service host that was used by the source NMC Server.
- If you use a License Manager server, install and configure the License Manager software first. If you use the License Manager software and the License Manager server moves to a new host, specify the new License Manager hostname in the **Console** window.
- By default, the recover process overwrites existing NMC files. To recover to the original location, stop the NMC services by typing the following command from a prompt:

On Windows:

```
net stop gstd
```

On Linux:

```
/etc/init.d gstd stop
```

Recovering the NMC Server

Perform the following steps to recover the NMC Server to the original host or a different host, from a point-in-time backup or the last backup time.

Procedure

1. Optional, to recover from an earlier backup, determine the *nsavetime* of the save set.

For example, on the NetWorker Server, type the following command:

```
mminfo -avot -q client=NMC_Server,level=full -r  
client,name,savetime,nsavetime
```

where *NMC_Server* is the hostname of the NMC Server.

Output similar to the following appears:

On Windows:

```
client name date save time  
bu-iddnwserver C:\Program Files\EMC NetWorker\Management  
\nmcdb_stage\pgdata 13/03/2017 1489431765
```

On Linux:

```
client name date save time  
bu-iddnwserver /nsr/nmc/nmcdb_stage 13/03/2017 1489431765
```

The *nsavetime* value appears in the last column.

2. On a target Linux NetWorker Authentication Service, set the `LD_LIBRARY_PATH` environment variable to include the postgres library path.

For example:

```
export LD_LIBRARY_PATH=NMC_Installation_dir/postgres/lib
```

where *NMC_installation_path* is `/opt/lgtnm` by default.

3. Change the directory to the NMC bin directory:

On Windows the bin directory is :

```
C:\Program Files\EMC NetWorker\Management\GST\bin
```

On Linux the bin directory is:

```
/opt/lgtnm
```

4. Follow step 1-9 of **Moving the NMC Server** topic of *NetWorker Administration Guide*.
5. On the target NetWorker Authentication Service, restore the NetWorker Authentication Service backup by typing the following command:

```
recoverpsm -s NetWorker_server -c source_NMC_server /nsr/nmc/  
nmcdb_stage
```

or

```
recoverpsm -f -s NetWorker_server -c source_NMC_server -p  
AES_Passphrase staging_dir -d dir_name
```

 **Note:** If you had set datazone pass phrase during backup, then `-p AES_Passphrase` is required.

where:

- `-f` instructs the recovery operation to delete the database files that currently exist in the database directory. Do not use this option if you want to restore the database files to a different location.
- *NetWorker_server* specifies the name of the NetWorker Server.
- *source_NMC_server* specifies the name of the source NetWorker Authentication Service, when you recover the database to a different NetWorker Authentication Service host.
- *AES_Passphrase* specifies the passphrase that was used during the NMC database backup.
- *staging_dir* specifies the staging directory that was used during the backup of the database on the source NetWorker Authentication Service.
- *dir_name* specifies the directory to relocate the recovered database files. When you use this option, you must manually copy the database files from the destination directory to the database directory defined for the NetWorker Authentication Service. Ensure that you retain the same ownership and permissions on the database files and the credential files after the copy completes.

During a recovery of the NetWorker Authentication Service database, the console GUI is unavailable. Consequently, messages such as mount requests cannot be addressed from the console. Consider the following during a recovery of the NetWorker Authentication Service database:

- Monitor the daemon log files for messages. The use of the NetWorker `nsr_render_log` command can make the `daemon.raw` file more user friendly for interpretation.
 - Use the `nsrwatch` command to view messages and use commands such as `nsrjb` to address those messages.
The *NetWorker Command Reference Guide* provides more information about `nsr_render_log`, `nsrwatch`, `nsrjb` and other NetWorker commands.
6. After the recovery completes, if you stopped the NMC services, start the NMC services, by typing the following command from a prompt:

On Windows:

```
net start gstd
```

On Linux:

```
/etc/init.d gst start
```


CHAPTER 10

Special recoveries on Windows hosts

This chapter contains the following sections:

- [Special windows recoveries Restoring a Windows Domain Controller host.....](#) 620
- [Performing a Windows BMR to physical or virtual computers.....](#) 622
- [Online recovery of Active Directory, DFSR, or Cluster services.....](#) 638

Special windows recoveries Restoring a Windows Domain Controller host

After you recover the file system data on a Windows host, you can recover the AD DS configuration. A Windows host that is assigned a Domain Controller role in a Windows environment has the Active Directory Domain Services (AD DS) software installed.

The Windows Roles and Features save set contains the AD DS backup. To recover the AD DS configuration on a domain controller, perform an authoritative or non-authoritative restore.

Active Directory restore information

Active Directory (AD) is the Windows directory service and the foundation for the Windows Distributed file system (DFS). AD is a component of the Windows system state on Windows 2008 R2 SP1, Windows Server 2012 and Windows Server 2012 R2 domain controllers.

A domain controller is a computer that stores directory data and manages user interactions with a domain, including login, authentication, directory searches, and access to other shared resources.

AD, identified by its writer name NTDS, is backed up as part of the Windows Roles and Features save set, a collection of system components that depend on each other.

These components are backed up together and usually restored together, through a BMR.

Only three of these components lend themselves to being restored separately through an online restore: Active Directory, DFSR, and Cluster Services (Cluster Database).

Selecting a restore method

Consider the following when selecting a restore method:

- The circumstances and characteristics of the failure and the two major categories of failure from an AD perspective are AD data corruption and hardware failure. AD data corruption occurs when the directory contains corrupt data that replicated to all domain controllers. Also, when a large portion of the AD hierarchy that is accidentally changed and this change was replicated to other domain controllers.
- The roles and functions of a failed NetWorker server.
- The types of restore for AD are non-authoritative restore and authoritative restore.

Non-authoritative restore

A non-authoritative restore returns the domain controller to its state at the time of backup, then allows normal replication to overwrite that state with any changes that have occurred after the backup completed.

After restoring the system state, the domain controller queries its replication partners. The replication partners replicate any changes to the restored domain controller, ensuring that the domain controller has an accurate and updated copy of the AD database.

Non-authoritative restore is the default method for restoring AD. This method is used in most situations that result from AD data loss or corruption.

Authoritative restore

An authoritative restore is an extension of the non-authoritative restore process that allows an administrator to recover a domain controller to a specific point in time and mark objects in the AD as authoritative. After you recover objects that are marked authoritative the objects are replicated to all the other domain controllers in the domain. Before you perform an authoritative restore, you must complete the steps for a non-authoritative restore.

Performing a non-authoritative AD restore on Windows Server 2008 R2 SP1, Windows Server 2012, 2012 R2, 2016, and 2019

About this task

To perform a non-authoritative restore of the AD on a Windows Server 2012 or 2012 R2, 2016, and 2019 complete the following tasks.

Procedure

1. Start the target domain controller in Directory Services Restore Mode.
2. Recover objects by using the NetWorker User application or the `recover` command. For example, perform an online AD restore in one of the following ways:

- In the NetWorker User application or NMC, browse to and restore the *WINDOWS ROLES AND FEATURES\NTDS* save set.
- From a command prompt, type the following `recover` command:

```
recover -s NetWorker_server -N "WINDOWS ROLES AND
FEATURES:\NTDS"
```

3. When the restore completes, verify that the operation was successful.
4. To complete the AD restore, restart the domain controller.

Performing an authoritative AD restore on Windows Server 2008 R2 SP1, Windows Server 2012, 2012 R2, 2016, and 2019

Perform the following tasks for an authoritative restore of AD objects.

Procedure

1. Start the target domain controller in Directory Services Restore Mode.
2. To recover the *WINDOWS ROLES AND FEATURES\NTDS* save set, perform the following step:

To use the command prompt for recovery, type the following command:

```
recover -s NetWorker_server -U -N "WINDOWS ROLES AND
FEATURES:\NTDS"
```

3. When the restore completes, verify that the operation was successful.
4. To complete the AD restore, restart the domain controller.

Performing a Windows BMR to physical or virtual computers

This section describes how to use the NetWorker Windows BMR image to perform a Bare Metal Recovery on physical hosts and VMware virtual machines.

Before you perform a BMR, verify that the new host meets the [Requirements for Windows BMR backup and restore](#) on page 437 and ensure that you complete the tasks listed in this section. For more information on Windows Bare Metal Recovery, refer to the section [Windows Bare Metal Recovery](#).

Prerequisites to performing a Windows BMR

If you do not first add the recovering host to a group that has the Recover Local Data privilege, BMR of a NetWorker server fails through the `authc` process. Before you perform a BMR, add the following entries into the users list in `NMC\Server\User Groups`.

For example, to add the recovering host in to the Application Administrators group, add the following entries to the users list in NMC:

```
group=Administrators,host=<recovering_host>
user=administrator,host=<recovering_host>
user=system,host=<recovering_host>
```

where *recovering_host* is the name of the host that you are performing the BMR to.

Gathering configuration information required by a Windows BMR

Before you start a Windows BMR, ensure that you have the following configuration information:

- The driver software for NICs or disk devices, if you perform the Windows BMR to a host with hardware that differs from the source host.
- The network name and IP address of the target host.
- The network name and IP address of the NetWorker server.
- The network name and IP address of the NetWorker storage node, if the target host uses a storage node that is not the NetWorker server.
- The default gateway and the name of the DNS server. If a DNS server is not available, use a local hosts file to resolve hostname of the NetWorker server and storage nodes to the IP address.
- The NetWorker media volumes that contain the backup save sets.

Obtaining the Windows BMR image

About this task

To perform a Windows BMR, use the Windows BMR image available from Online Support website to create a bootable CD/DVD or deploy for a network boot operation. The BMR image contains the Windows PE operating system. WinPE is only available in English, localized versions of the Windows BMR wizard are not available. When you use the image to boot the Windows host, the recovery process starts the NetWorker BMR wizard, which guides you through the recovery process.

You can use the 32-bit, or 64-bit Windows BMR image to recover either an x86, or x64 operating system backup to an x86 or x64 computer.

Note: A BMR treats the AMD and Intel processors as equivalent if they follow the same architecture. For example, you can recover the operating system from the backup of AMD x64 computer to an Intel x64 computer.

Use the following procedure to download the recovery boot image.

Procedure

1. On the Online Support website, search for “NetWorker Wizard ISO”, and then narrow the search results by selecting items that are associated with the NetWorker release number.
2. On the **NetWorker Software Downloads** page:
 - a. Locate the section that is labeled **NetWorker Y.Y - Build xxx**.
 - b. Select the link to download a Windows BMR ISO recovery file.

where:

- Y.Y is the version number of the NetWorker release.
- xxx is the build number of the released version.

Creating a Windows BMR bootable image

Create a Windows BMR bootable CD/DVD or a network boot location from the Windows BMR ISO image, which you downloaded from the Online Support website.

Creating a Windows BMR bootable CD/DVD

Use the ISO image to create a bootable CD/DVD, then configure the host to boot from a CD/DVD.

Procedure

1. Open the CD/DVD creation software, and then select an option to burn an ISO image.
2. Browse to the location of the downloaded NetWorker Windows BMR image, and then complete the steps that are required to create a bootable CD/DVD with the image.

Enabling a protected host to boot from a CD/DVD

Procedure

1. Start the host, and then start the BIOS setup program, by pressing **F2**.

NOTICE If you are restoring either from or to a virtual host such as a VMware virtual machine, you can set up options such as the host boot location within vSphere. The VMware documentation provides specific steps.

2. Select the **boot options** menu, and then ensure that the CD/DVD boot option is at the top of the list of locations from which to boot.
3. Save the changes, and then exit the BIOS program.

Creating a Windows BMR recovery network boot location

About this task

Ensure that you meet the following requirements for using the network boot option:

- Ensure the NetWorker clients that you protect with a Windows BMR backup can start from the network with a Pre-Boot Execution Environment (PXE).
- Configure and make available a Deployment Services server.
- Add the NetWorker Windows BMR boot image to the Deployment Services server so that a client host on the network can start from it.

Enabling a host to boot from the network

About this task

The host should obtain an IP address from the WDS server, and then prompt you to perform a network boot. Typically, a network boot is activated by pressing the **F12** key.

Procedure

1. Start the host, and then start the BIOS setup program.

Typically, this action is performed by pressing the **F2** key.

 **NOTICE** If you are restoring to a virtual host such as a VMware virtual machine, you can set up options such as the host boot location within vSphere. The VMware documentation provides specific steps.

2. Select **BIOS options necessary** so that the network boot option is enabled.
The BIOS documentation provides more information.
3. Save the changes, and then exit the BIOS program.

Perform the BMR

Perform the BMR to a physical host or a virtual host:

- [Performing a Windows BMR to a physical computer](#) on page 624
- [Performing a BMR from a Physical Computer to a Virtual Machine \(P2V\)](#) on page 631

Performing a Windows BMR to a physical computer

About this task

To recover a BMR backup to a physical host, perform the following steps. If the target host uses unified extensible firmware interface (UEFI) volumes, unmount the UEFI volumes before you perform the recovery operation.

Review the following information before you perform a recovery operation to a host that differs from the original:

- Ensure that the hardware configuration of the target host is similar to the original host.
- Ensure that the NetWorker server has a client resource for both the source host and the target host.

- Ensure that the Remote Access attribute of the source client resource contains the account *SYSTEM@target_client*. This attribute enables the recovery process to perform a directed recovery.
- Add *user=system,host=target_client* to the Users attribute of Application Administrators user group.
- Ensure that you have a writable volume available for the media pool being used. After the recover operation recovers all the data, the wizard generates log files in a save set named Offline Restore Logs. The recovery operation performs a backup of the log files to a volume in the media pool.
- Ensure that you enable the NetWorker server to accept manual save operations for the Recovery wizard log file backup.

Procedure

1. Start the target host from the Windows BMR image.

The NetWorker Windows BMR wizard appears.

2. On the **Welcome** screen, click **Next**.

3. If a DNS server is not available on the network, perform the following:

- a. Exit the NetWorker Windows BMR wizard but do not restart the host.

The WinPE command line appears.

- b. Edit the hosts file, for example, `X:\Windows\System32\Drivers\etc\hosts`.

- c. Add the IP address and hostname for the NetWorker server and the NetWorker storage node.

- d. Restart the wizard from the `X:\Program Files\EMC Networker\nsr\wizard` directory.

For example: `X:\Program Files\EMC Networker\nsr\wizard> javaw -jar WinPEWizard.jar`

- e. When the wizard appears, click **Next**.

4. On the **Select Network Interface** screen:

- a. Select the NIC driver.

- b. Click **Next**.

If the driver list does not contain the driver for the NIC on the target host, select **Load Driver**, and then browse to the location that contains the required driver.



NOTICE The selected driver cannot require a restart operation because the recovery process loads the WinPE environment in memory only and changes are not persistent after a restart operation.

5. On the **Configure Hostname and Network** screen, complete the fields:

- a. In the **Hostname** field, type the hostname of the source host.

- b. In the **DNS domain** field, type the name of the domain in which the host resides.

If the host resides in a workgroup instead of a domain, you can leave this field blank.

- c. In the **Configure desired IP Settings** field, choose the tab for the Network Protocol deployed on the network, either **IPv4** or **IPv6**.
- d. In the **TCP/IP Address** settings section, select either **Obtain an IP address automatically (DHCP)** or **Use the following IP Address**.
- e. If you choose **Use the following IP Address**, type the IP address in the IP address field.

If applicable, type the subnet mask in the **Subnet mask** field, and then type the default gateway in **Default gateway** field.

- f. In the **DNS Server** section, select either **Obtain DNS server address automatically** or **Use the following DNS server address**:
 - If you choose **Use the following DNS server address**, type the IP address of the DNS server in the **Preferred DNS server** field.
 - If applicable, type an alternate DNS server address in the **Alternate DNS server** field.

 **NOTICE** If you added the NetWorker server hostname and IP address to the `X:\Windows\System32\Drivers\etc\hosts` file, you can ignore the DNS Server fields.

- g. Click **Next**.

The **Available Disks** screen displays all detected local disks.

6. If the **Windows BMR** wizard fails to detect a disk, perform the following steps:
 - a. Select **Load Driver**.
 - b. Browse to the location that contains the disk driver, and then load the required disk driver.
 - c. To update the list of detected disks, select **Refresh**.
 - d. Click **Next**.

7. On the **Select NetWorker Server** screen, complete the fields:
 - a. In the **Server** field, specify the NetWorker server that performed the backup:
 - Select the NetWorker server from the server list. To update the list of NetWorker servers, click **Search**. The Search function locates only those NetWorker servers on the local subnet.
 - Type the fully qualified domain name (FQDN).

- b. In the **Client** field, ensure that the client name matches the Client resource name on the NetWorker server.

NetWorker automatically populates this field with the values that you specified in the **Hostname** and **DNS Domain** fields on the **Configure Hostname and Network** screen of the wizard. For example, if the client resource on the NetWorker server uses an FQDN, then specify the FQDN of the client in the **Client** field.

To recover the backup to a host that differs from the source host, modify the **Client** field to specify the target hostname.

If you specify a different client, the recovered host uses the same hostname and IP settings as the source computer. If the source computer is running on

the same network, using the same hostname and IP settings can cause hostname and IP address conflicts.

c. Click **Next**.

8. On the **Select Bare Metal Recovery Backup** screen, select the system backup that you want to recover, and then click **Next**.

System backups appear in descending order from most recent to oldest.

9. Review the **Save Sets to Restore** screen, and then click **Next**.

The recovery process reformats critical volumes. The recovery process reformats non-critical volumes only if the disk signature on the target disk differs from the original disk.

For example, to perform a quick format instead of a full format operation if the disk was replaced, select **Perform a quick format of disks**.

 **Note:** A quick format is much faster than full format but does not verify each sector on the volume.

The recovery process does not recover non-critical volume data. [Recovering file system data](#) provides more information.

10. On the **Bare Metal Recovery Summary** screen, select **Options** to display the **Non-Default Recover Options** screen.

11. On the **Non-Default Recover Options** screen:

- a. In the **Additional Options** field, type any required non-default options with their corresponding values.

Non-default options are primarily used for troubleshooting purposes.

- b. To save and close the **Non-Default Recover Options** screen, and then return to the **Bare Metal Recovery Summary** screen, click **OK**.

- c. To begin the recovery process, click **Restore**.

12. On the **Confirmation** screen, select the **I confirm that I want to format the disks and restore the backup** option, and then click **OK**.

 **NOTICE** All data is lost on all volumes that the recovery process reformats.

After the data recovery completes, the wizard writes the recovery log files to volumes in the backup media pool being used. If you do not have a volume available, then the recovery operation appears to be unavailable until media for the media pool becomes available.

 **Note:** You can cancel the log file backup without affecting the recovery operation.

13. After the wizard and log files complete, click either **Reboot** or **Exit**:

- To restart the system when any subsequent application data resources must be performed, click **Reboot**. If you are recovering an Active Directory domain controller, it is recovered in non-authoritative mode by default.
- If you must recover a domain controller in authoritative mode, click **Exit**. The computer returns to the WinPE command prompt. Start into Directory Services Restore Mode (DSRM). See [Performing post recovery tasks for active directory services](#) for more information.

Post-recovery tasks

The following sections provide information about recovering data that was not recovered in the Windows BMR operation.

Using NMM for post-recovery tasks

If the recovered host has applications that are protected with NMM, all application-recovery operations must be performed by using the NMM client interface. The NMM documentation provides information on the post-recovery operations.

About this task

Before reviewing the NMM documentation, review the following information:

- After the recovery has completed and the system is rebooted, check the host's disk and volume configuration. All disks and volumes should appear as they did on the original system. However, if disk signatures do not match the original disks, non-critical disks might be offline or unmounted. Use Microsoft Disk Manager to bring online or mount the disks. After the disks are online, a reboot operation should result in disk drive letter reassignments. If these correct drive letter assignments do not occur, manually assign drive letters to non-critical disks as needed. Non-critical volumes that are accessed by mount points might have similar issues.
- To recover the host, perform additional online recovery of any required user data on non-critical volumes by using the NetWorker User program.
- If a folder is encrypted in Windows, for example, by selecting **Folder Properties > Advanced > Encrypt contents to secure data**, it is recovered as encrypted. However, the encryption attribute is not be set on the folder. You can manually reset the encryption attribute after the recovery operation. This task is a Microsoft limitation.
- Windows BMR supports backup and recovery of files and folders encrypted with Windows Encrypting File System (EFS), and volumes encrypted with BitLocker. After BMR, the EFS or BitLocker services might be running but the EFS encryption attributes on files or folders must be re-enabled and BitLocker volumes must be re-encrypted. For steps to encrypt with EFS and BitLocker, consult Microsoft documentation.

 **NOTICE** You cannot install the NetWorker software on volumes that are encrypted with Microsoft BitLocker.

Using an application backup tool other than NMM

About this task

If you backed up a database application with an application backup tool other than NMM, perform the following post-recovery operations:

- Recover any required file system data by completing the steps in the topic, [Recovering file system data](#).
- Recover the application data by using the application backup tool, such as NetWorker User for SQL Server, NME, or any third-party application backup tool. Refer to the documentation that your application backup tool includes.

Recovering file system data

About this task

Perform an online recovery of any required user data on non-critical volumes. Sometimes, user data on non-critical volumes must be recovered, for instance, when disk hardware was replaced due to a disaster before the Windows BMR operation.

Procedure

1. Manually remount any non-critical volumes as needed.
2. To connect to the NetWorker server that backed up the source client data, start the NetWorker User program by using the `winworkr` command with the `-s` option.

For example: `winworkr -s server_name`

If the `-s` option is not used and there is only one server that is detected, that server is connected automatically. If there are no servers that are detected or if there is more than one server available, the **Change Server** dialog box appears, allowing you to choose the server.

3. To open the **Source Client** dialog box, click **Recover**.
4. Select the source client, and then click **OK**.
5. Select the destination client for the recovered data, and then click **OK**.
6. In the **Recover** screen, select the files to recover.
7. To begin the directed recovery, click **Start**.

Performing post-recovery tasks for Active Directory services

About this task

Perform the offline recovery of the `DISASTER_RECOVERY:\` component save sets if there is a non-authoritative domain controller. If a non-authoritative recovery is wanted, then no additional steps are required. However, if you must perform an authoritative recovery, follow these steps.

Procedure

1. To exit the wizard so that you can start into Directory Services Restore Mode (DSRM), on the **System Recovery Results** screen of the NetWorker Bare Metal Recovery wizard, select **Exit**.

NOTICE Do not select **Reboot** in the wizard. Failure to start into DSRM mode results in a non-authoritative recovery. If you select **Reboot**, perform one of the following:

- On restart, start the system in the WinPE operating system instead of the restored operating system.
- Run the Windows BMR wizard again and ensure that you select **Exit**.

The WinPE command prompt appears.

2. At the WinPE command prompt, type the following `bcdedit` commands.
 - a. To force the system to start into DSRM, add a boot loader entry:

```
bcdedit /copy {default} /d "Directory Service Repair Mode"
```

A message similar to the following appears:

```
The entry was successfully copied to
{00000000-0000-0000-0000-000000000000}
```

The numbers and dashes in the previous message form a Globally Unique Identifier (GUID) that identifies a new entry. In this example, the GUID is for illustration purposes only. The actual GUID that is generated when you run the command is unique.

- b. To set the safeboot option for the bootloader entry in the BCD store, type the following command using the generated GUID:

```
bcdedit /set {GUID_value} safeboot dsrepair
```

where *GUID_value* is the GUID displayed by the previous `bcdedit` command.

- c. To restart the system, exit the WinPE command prompt.

 **Note:** Failure to start into DSRM results in a non-authoritative recovery.

3. (Optional) If you have a `WINDOWS ROLES AND FEATURES:\ Active Directory` subcomponent save set that is newer than the `DISASTER_RECOVERY:\ save` set used in the preceding BMR, you can recover the save set in DSRM through the NetWorker User program.

4. From the WinPE command prompt, run the Windows `ntdsutil` utility.

The `ntdsutil` prompt appears. The `ntdsutil` utility is a command interface similar to the NetWorker recover interface. For help with the `ntdsutil` utility, type:

```
NTDSUTIL: ?
```

5. At the `ntdsutil` prompt, type:

```
NTDSUTIL: activate instance ntds
NTDSUTIL: authoritative restore
```

6. To perform an authoritative recovery of a subtree or individual object, type:

```
NTDSUTIL: restore subtree "distinguished_name"
```

For example:

```
NTDSUTIL: restore subtree
"OU=engineering,DC=Seattle,DC=jupiter,DC=com"
NTDSUTIL: restore subtree
"CN=mars,CN=users,DC=Seattle,DC=jupiter,DC=com"
```

The Microsoft Windows Server Resource Kit documentation on Active Directory provides information.

7. Exit the `ntdsutil` utility by typing `quit` at each successive `ntdsutil` prompt until the command prompt appears.
8. Type the following command at the WinPE command prompt so that the host does not start into DSRM mode on restart.

```
bcdedit /deletevalue safeboot
```

9. Restart the domain controller in normal mode, log in, and then verify that the authoritative changes are replicated to the Active Directory replication partners.

Performing post-recovery tasks for hosts with Windows server roles that use SQL Server Procedure

1. On the target host, rebuild the SQL server by running the following `Setup` command:

```
Setup /QUIET /ACTION=REBUILDDATABASE /
INSTANCENAME=Instance_name /SQLSYSADMINACCOUNTS=domain_name
\administrator
```

The Setup tool is located on the SQL Server installation media and must be run from the command prompt with Windows Administrator privileges. Before you run this command, ensure that the SQL group is offline except for the shared disks.

The following Microsoft article provides more information:

<https://msdn.microsoft.com/en-us/library/ms189302>

2. Bring the SQL server services online.
3. Recover the SQL system databases (master, model, msdb) with NetWorker User for SQL Server, or a third-party application.

Performing post-recovery tasks for a Microsoft Hyper-V virtual machine Use NMM to restore the Hyper-V virtual machines.

Performing a BMR from a Physical Computer to a Virtual Machine (P2V)

This section describes the process of restoring a NetWorker backup of a physical computer to a virtual machine (P2V).

About this task

P2V is supported for physical computers running the following operating systems:

- Windows Server 2008 R2 SP1
- Windows Server 2012
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2019

P2V is supported when restoring to virtual machines created with the following hypervisors:

- Microsoft Hyper-V Server 2008 R2
- Microsoft Hyper-V Server 2012
- Microsoft Hyper-V Server 2012 R2
- VMware ESX 5.1
- VMware ESX 5.5
- VMware ESXi 5

Procedure

1. Perform a backup of the physical computer.
2. On the computer that runs the hypervisor, create a target virtual machine (VM).

- a. Configure the VM to use a virtual network adapter.
- b. On the **VM configuration** page:
 - a. Select the LSI Logic SAS SCSI controller.
 - b. Configure the disks on the VM to match the original physical computer configuration.
 - c. Create the same number of physical disks.
Extra disks can be added after the P2V recovery.
- c. Consider the following:
 - The SCSI disk numbers must match the original disk numbers.
 - The VM disk sizes must match, or exceed, the original disk sizes.

For VMware hypervisors, use either a Windows Server 2008 (32-bit or 64-bit), 2008 R2 (64-bit), or Windows Server 2012 (64-bit) templates as the guest operating system when you create the VM.

3. On the VM, start the WinPE ISO which starts the BMR wizard.
4. On the VM, use the BMR wizard to configure the hostname and network configuration:
 - a. On the **Select NetWorker Server** screen, specify the name of the physical computer as the NetWorker client.
 - b. On the **Save Sets to Restore** screen, review the selected items to restore, and then click **Next**.
 - c. On the **Select Bare Metal Recovery Backup** screen, select the backup to restore. Backups are listed in chronological order with the most recent backup first.
 - d. On the **Summary** screen, if the save set was created with NetWorker 8.1 or earlier, select the **Restore physical computer to virtual machine (P2V)** checkbox.

If the Restore physical computer to virtual machine (P2V) checkbox is not marked, the VM might not start successfully after the restore is complete.

e. To start the restore, select **Restore**.

5. Restart the VM when the P2V BMR is complete.

Performing Post-P2V tasks

The following section provides information about additional tasks that are required after a P2V recovery.

Procedure

1. If you are running VMware, install VMware tools.
2. To remove disabled NIC devices, use **Device Manager**:
 - a. From Device Manager, select the **Show Hidden Devices** option.
 - b. Select the hidden NIC device.
 - c. Select **Uninstall**.

This step is required because the original network adapter is no longer available.

3. To restore network connectivity, configure the virtual network adapter.

Troubleshooting Windows BMR

The following topics provide information to help troubleshoot Windows BMR operations.

Performing a manual uninstall and reconfigure of a NIC on Windows Server 2012 or Windows Server 2012 R2

If the guest operating system is Windows Server 2008 R2 SP1, the P2V BMR retains the NIC settings.

About this task

If the guest operating system is Windows Server 2012 or Windows Server 2012 R2, then Windows performs some Plug-N-Play configuration during the post-BMR restart. This activity disables the original NIC and creates a NIC.

Procedure

1. In the **Device Manager**, select **Display disabled devices > Uninstall the disabled NIC**.
2. Configure the new NIC with the wanted network settings.

Recovering and viewing Windows BMR log files

Windows BMR log files

To help troubleshoot an unsuccessful recovery, the following log files are generated and backed up during the Windows BMR operation:

- `daemon.raw`—This log file is the same as `daemon.log` for monitoring services.
- `Ossr_director.raw`—Contains the recovery workflow of the `DISASTER_RECOVERY:\ save set`. This log also contains any errors that are related to recovering the save set files or Windows ASR writer errors.
- `recover.log`—Contains output from the NetWorker `recover.exe` program. This information is generated during the recovery of each save set. This log also contains messages about errors that are related to critical volume data recovery.
- `WinPE_Wizard.log`—Contains information about the workflow flow that is related to the NetWorker Bare Metal Recovery wizard user interface.
- `winpe_nw_support.raw`—Contains output from the `winpe_nw_support.dll` library. The output provides information about the communication between the NetWorker Bare Metal Recovery wizard and the NetWorker server.
- `winpe_os_support.log`—Contains output information that is related to Microsoft native API calls.

If the Windows BMR fails, you can recover the log files using one of the following options:

- By using FTP on the recovery host.
- By using a directed recovery.
- By copying the log files to a mapped drive.

If the Windows BMR was successful, you can recover the log files directly to the recovered host.

To view log files, you can use either a text editor or the `nsr_render_log` program, depending on the log file format.

Viewing the log files

To view the following log files, use a text editor:

- `recover.log`
- `WinPE_Wizard.log`

To view the following log files, use the `nsr_render_log` program:

- `Ossr_director.raw`
- `winpe_nw_support.raw`

For example, to display the `Ossr_director.raw` file, type the following command at a command prompt:

```
nsr_render_log "C:\logs\Client-bv1\Ossr_director.raw"
```

To direct the `Ossr_director.raw` file to a text file that can be viewed in a text editor, type the following:

```
nsr_render_log "C:\logs\Client-bv1\Ossr_director.raw" > mylog.txt
```

Accessing the log files using FTP

Procedure

1. On the recovery host, access the WinPE command line.

You might have to exit the Windows Bare Metal Recovery wizard to access the WinPE command line. If you exit the wizard, do not restart.

2. Disable the Windows firewall.

For example:

```
wpeutil DisableFirewall
```

By default, the Windows firewall is enabled on WinPE, and this action blocks the FTP port from transferring files.

3. Change to the following directory that contains the log files:

```
X:\Program Files\EMC NetWorker\nsr\logs
```

4. To move the log files to another NetWorker host, use the FTP utility.

Accessing log files using a directed recovery operation

Procedure

1. To connect to the NetWorker server that backed up the source client data, start the NetWorker User program by using the `winworkr` command with the `-s` option:

```
winworkr -s server_name
```

If the `-s` option is not included, and there is only one server that is detected, that server is connected automatically. If there are no servers that are detected or if there is more than one server available, the **Change Server** dialog box appears, enabling you to choose the server.

2. To open the **Source Client** dialog box, click **Recover**.
3. Select the source client, which is the recovered client, and then click **OK**.
4. Select the destination client for the recovered data, and then click **OK**.
5. From the **Options** menu, select **Options**, specify a folder location in which to relocate the recovered log files, and then click **OK**.
6. In the **Recover** window, select the log files to recover.

The log files are typically located in the following directory:

```
X:\Program Files\EMC NetWorker\nsr\logs
```

7. To begin the directed recovery, click **Start**.

[Recovering file system data](#) provides more information about the permissions that are required for directed recoveries.

BMR backup fails when System Reserved Partition is offline

BMR backups may fail with the following error:

```
device is not ready
```

Windows Server 2008 R2 SP1, Windows Server 2012, and Windows Server 2012 R2 has 100 MB reserved as the System Reserved Partition. When backing up the system state, VSS includes the System Reserved Partition (used for BitLocker and Boot files), but the backup fails because the System Reserved Partition is offline. This can occur if the Windows automount capability is disabled. Although there are circumstances where the automount capability must be disabled, it can result in the partition being offline after a restart. Automount must be enabled for a BMR backup to succeed.

To work around this issue, use either of the following solutions:

Solution 1

From the command prompt, run `DISKPART` with the following commands:

```
DISKPART
```

```
List volume
```

```
Select volume <number of 100 MB system partition>
```

```
Online volume (if the volume is offline)
```

If automount is disabled while using third party storage software or if the user manually disabled the automount for the volume, the volumes can go offline.

This Microsoft KB article 2419286, available at <http://support.microsoft.com/kb/2419286>, provides details on preventing volumes from going offline by checking and setting the system automount attribute.

Solution 2

From the Disk Management console:

1. Access Disk Management from the command prompt:

```
C:\>Diskmgmt.msc
```

2. To bring the disk online, assign the drive letter to the 100 MB partition:

- a. Right-click the 100 MB volume, and then select **Change Drive Letter and Paths**.
- b. Assign a new drive letter to the volume.
Assigning the drive letter ensures that the volume are online after a restart.

Wizard cannot locate the NetWorker server or DNS server

If the NetWorker Bare Metal Recovery wizard cannot locate the NetWorker server or the DNS server (if one is being used), consider the following:

- If you are using a local hosts file instead of a DNS server, verify that the hostname and IP address of the NetWorker server was typed correctly.
- If you are using a DNS server, verify that the values typed in the **Configure Hostname and Network** screen were typed correctly.
- Verify that the NetWorker server was correctly specified in the **Select NetWorker Server** screen.

To verify hostname and IP address values, use the ping utility that is in the WinPE environment:

1. Exit the NetWorker Bare Metal Recovery wizard but do not restart the host.
You are returned to the WinPE command line.
2. To locate, and then verify hostnames and IP addresses, use the **ping** utility. For example:

```
ping -a hostname
```

3. Restart the wizard. For example:

```
javaw -jar WinPEWizard.jar
```



Note: After the wizard has been restarted, you can switch between the wizard and the WinPE command line without exiting the wizard.

Multiple NICs cause errors in locating the NetWorker server

An error message similar to the following might appear when you try to recover a host with multiple NICs:

```
Error retrieving the list of Networker servers
```

This message is an indication that the NIC selected by the wizard is not the NIC that was connected to the NetWorker server when the backup was performed and the NIC might not have connectivity to the server. This applies when searching for an available server or specifying a specific server. To resolve the issue, select another NIC.

Network configuration values might not be retained after reboot

Sometimes, a host does not retain its network configuration data after a Windows BMR operation and after the host starts. If the recovered host is experiencing network connectivity issues, confirm that network properties for the local connections are correct. If required, manually update the network configuration data on the host.

VSS backups fail because a critical disk is offline

VSS backups fail if a critical volume is offline during the backup operation. You can remedy the problem by following the steps that are outlined in the Microsoft Knowledgebase (KB) article 980794, which can be found at:

<http://support.microsoft.com/kb/980794>

The patch that is mentioned in this knowledgebase article is most likely on the Windows system if it is up-to-date. In this case, you can create and populate the Registry keys as described in the article.

This issue is most often encountered when backing up a passive node in an MSCS cluster and a critical volume is not on the physical host of the passive node but is instead on the physical host of the active node.

Jobquery fails to establish a connection with large scale jobs

When querying the number of save sets, jobquery fails to establish a connection with the jobsDB when the jobsDB contains more than 3,00,000 records.

The workaround is to run `nsradmin` from the command line with the following parameters:

```
nsradmin -S <jobsdatabase path>
```

8dot3name support disabled after recovery

In a WinPE 5.0 environment, 8dot3 file name support becomes disabled after recovery. This is not an issue from block-based backups.

If you require 8dot3name support, run the following command:

```
fsutil 8dot3name set C: 0
```

The Microsoft knowledgebase article 121007, available at <http://support.microsoft.com/kb/121007>, provides more information.

Additional recovery options

You can specify non-default recovery options on the WinPE command line or in the **Additional Options** field in the NetWorker Bare Metal Recovery wizard.

The following table describes the additional recovery options that can be used with a Windows BMR operation.

Table 98 Additional recovery options

Entry	Result
-D x where <i>x</i> is a number from 1 to 9, with 9 providing the most troubleshoot information and 1 providing the least.	Additional troubleshoot information is in the Windows BMR log files.
-v	Additional information on the progress of the recovery displays in the wizard's System Recovery Status window.
-p	By default, the Windows BMR recovery skips the formatting of non-critical disks. By using the -p option, any existing partitions are deleted and all disks are reformatted on the recovered computer to match the layout of the system image. However, by Microsoft specification, even if the -p option is selected, a non-critical

Table 98 Additional recovery options (continued)

Entry	Result
	volume is not reformatted if the disk signature has not changed since the backup. This option might be useful in situations where a system fails to recover because of disk mismatch errors. In this case, the <code>-p</code> option might resolve those errors. The recovery process does not recover non-critical volume data even if the volume is reformatted. Non-critical volumes can be recovered by using the NetWorker User program after the wizard has completed and the host has been restarted.
<pre>recover -s <NetWorker server> -U -N "WINDOWS ROLES AND FEATURES \Cluster Database"</pre>	When the restored data is meant to override the data on other nodes, it should be restored using the authoritative mode. Once this data is restored to one of the nodes, it is propagated to the other nodes and overwrites any newer data on those nodes. Perform Authoritative restore by using the command on the left. While the recovery is in progress, observe that the status of the groups changes from Online to Pending to Offline in the Failover Cluster Management application. Alternatively, check the Event Viewer, under Application and Services Logs > Failover Clustering > Operational on all nodes that the Cluster Service has stopped and restarted. Recover the shared drive data through <i>winworkr</i> on the cluster node with its current active node. Select source client as the virtual client, and destination client as the current active node.

Restart required after recovery operation

Newly recovered NetWorker client computers running Windows Server 2012 R2 can require an extra restart to restore access to application icons, previously viewable on the desktop.

Online recovery of Active Directory, DFSR, or Cluster services

The `DISASTER RECOVERY:\ save` set includes the `WINDOWS ROLES AND FEATURES` component save set. You can recover the `WINDOWS ROLES AND FEATURES` backup in an online recovery operation, to a host that uses the same Windows operating system instance. NetWorker 8.2 and higher support the online recovery of the following Windows services, which the `WINDOWS ROLES AND FEATURES` component contains:

Active Directory

Solve Desktop provides procedures that describe how to recover this service.

Distributed File System Replication (DFSR)

The topic, *Backing Up and Restoring a Microsoft DFS*, provides more information.

Cluster

Solve Desktop provides procedures that describe how to recover this service.

NetWorker does not support the online recovery of any other Windows service that the `WINDOWS ROLES AND FEATURES save set` contains. Unsupported online recovery of `WINDOWS ROLES AND FEATURES` components results in an inconsistent state of the Windows server.

NOTICE When you perform an online recovery, you cannot mark the `WINDOWS ROLES AND FEATURES save set` and use the *Required Volumes* option. To determine the volume that contains the `WINDOWS ROLES AND FEATURES save set` that you want to restore, mark the `DISASTER RECOVERY:\ save set`, then use the *Required Volumes* option. After you determine the required volumes, unmark the `DISASTER RECOVERY:\ save set` and mark the `WINDOWS ROLES AND FEATURES save set`.

CHAPTER 11

Reporting NetWorker Datazone Activities

This chapter contains the following topics:

• Enterprise data reporting.....	642
• Reporting policy status and backup job status.....	685
• Reporting recover job status.....	705
• Checkpoint-enabled backup reporting.....	706
• SNMP traps.....	707
• NetWorker Notifications.....	717
• Front-end Capacity Estimation.....	729

Enterprise data reporting

NetWorker software automatically collects data on a continual basis from the NetWorker enterprise to facilitate trend analysis, capacity planning, and problem detection.

The NMC server stores the collected information in the Console database for a specified number of days, as described in [Data retention and expiration policies](#) on page 643.

The NetWorker software then integrates and processes this data to produce a number of reports on backup status, backup statistics, events, inactive files, hosts, users, and devices. [Report categories](#) on page 645 provides detailed information about the various types of reports.

The following options are available through the NetWorker Console reporting feature:

- Data collection for the entire enterprise or for specific NetWorker servers.
- Creating of various types of reports.
- User preferences for report data, such as font, size, and whether to use bold. This can be useful in I18N environments.
- Selection of columns to display when viewing reports in a table format, and the order in which to display them.
- The ability to save customized reports for repeated use.
- The ability to determine how long collected data should be retained.
- Only NetWorker administrators can modify these time periods.
- The ability to share reports, or restrict the sharing of reports, with other users by giving them access to the reports.
- The ability to hide shared reports of other users when listing reports.
- The ability to run reports from the command prompt.

Enabling or disabling the gathering of report data

About this task

When you add a host to the enterprise, the **Configuration** wizard enables the **Gather Reporting Data** feature by default. To enable or disable the Gather Reporting Data option after you add a host to the enterprise, perform the following steps.

Procedure

1. From the NMC GUI, click **Enterprise**.
2. In the left navigation pane, expand **Enterprise**, and then right-click the NetWorker server on which to enable the collection of report information.
3. Select **Properties**.
4. In the **Features** section, select **Gather Reporting Data** to enable the feature or clear the option to disable the feature, then click **OK**.

Data retention and expiration policies

The NetWorker Console provides separate expiration policies for retaining different types of data to meet the needs of the environment as described in this table. Only a Console Application Administrator can modify these policies.

Table 99 Data retention policies

Retention policy	Type of data to be retained	Default
Statistical data—Affects all legacy Backup Statistics reports and Policy Statistic reports. The retention policy for statistics data can affect multiple reports.	Backup and cloning statistics.	One year
Recover Statistics — Affects Save Set Data in Recover Statistics reports.	Save set records.	One year
Audit Data— Affects User reports. The retention policy for audit data affects only audit reports.	Reports on all NetWorker tasks (except License Manager tasks) performed by specified users (but only when the NetWorker User Auditing system option is activated).	One year
Completion Data (legacy)— Affects Backup Status reports, except in the save set output. Retention policy for completion data can affect multiple reports.	Savegroup and save set completion data and drive data.	One month
Completion Message (legacy)— Affects Backup Status reports, only in the save set output). The retention policy for completion messages can affect multiple reports.	Messages, such as error messages for failed save sets.	Two weeks

You can view the retention policies for data to which they have access by following the first three steps in [Setting expiration policies for data retention](#) on page 644. These different policies give administrators the flexibility to retain certain types of information for less time than others, as showed in the following example.

 **Note:** Reports not mentioned in the above table have no retention policies.

Example 11 Retention Flexibility

An administrator might want to set the completion message policy to a shorter period than the completion data policy. The precise error messages about what caused a save set backup to stop might not be relevant over a longer period. But it might be useful to save the completion data for a somewhat longer period to help with load balancing and trends.

Example 11 Retention Flexibility (continued)

The longest period (one or more years) might be a suitable selection for save set data. This data is used to generate the NetWorker Backup Statistics reports. These reports can be used to determine historical trends about backups and to help guide capacity planning.

Note: The expiration policies restrict the data that can be retrieved by NetWorker Console. In other words, reports cannot include data that is older than the data retention policy. If, for example, an administrator changed a policy expiration period from 1 year to 1 month and soon afterwards reset it to 1 year, 11 months of data would be lost. Once data is cleared because of the retention policy, you can only retrieve the data by recovering the full database.

Setting expiration policies for data retention

Before you begin

Log in to the NMC server as a Console Security Administrator. The NetWorker Authentication Service administrator account is a Console Security Administrator.

About this task

Perform the following steps to define how long the NMC server stores information about NetWorker server activities in the NMC database.

Procedure

1. From the **NMC GUI** window, click **Reports**.
2. From the **Reports** menu, select **Data Retention**. The **Data Retention** dialog box appears.
3. For each policy, type the number of periods and select a period of time (year, month, week, day).
4. To save the configuration of the data retention policies, click **OK**.

Note: There must be adequate space in the NMC database to hold the data. If the data retention policy settings cause the NMC database to run out of storage space and the NMC processes to stop running. The *NetWorker Installation Guide* provides information about estimating the size of the NMC database.

Restricted report views

NMC users can only view report information about servers to which they have permission to manage.

Since each user can have different access restrictions, different users may see different report results. This applies to customized, private, and shared reports.

For example, a shared Group Summary report entitled “Building C Backups” will show different data for different users if the access permissions for each user includes different NetWorker servers. This applies even if the users run the report at the same time.

On the Configuration tab of each report, the configuration parameters will only display to the user, the allowed NetWorker servers, groups, and clients as sources of report information. The generated report will only contain data from allowed

resources. Users may only run reports for servers to which they are allowed to manage.

 **Note:** If no data is available for a given server, that server will not appear in any lists, regardless of the access permissions for the user.

Report categories

The following table describes the various report categories in the NetWorker software. Each of these categories is discussed in detail in [Preconfigured reports](#) on page 653.

Report categories appear as folders within the Reports folder in the **Reports** window. You can run these reports from the NMC GUI or from a command prompt.

Table 100 Report categories

Category of report	Purpose
Policy Statistics	Provide statistical information about activities and resources in the Data Protection Policies. Include information about the Workflow resources, Client resources, Group resources, and Action results.
Recover Statistics	Provide the history of recovery operations that have been performed by NetWorker servers.
Devices	Provide information about the way devices are being used.
Events	Provide summary and detailed information about NetWorker events.
Hosts	Provide a listing of NetWorker servers in the Enterprise, including information about event and reporting features.
Users	Provide lists of defined NetWorker Console users, logout and login reports, audit reports, and users with restricted views.
Manual saves	Provides save set information about backup operations that are initiated by a user with the <code>save</code> command, and details about clone operations that are initiated by a user with the <code>nsrclone</code> command.

Legacy report categories

The following table describes the various report categories available in NMC, which enables you to report information about activities that occurred on the NetWorker server before an update to NetWorker 19.1, or for NetWorker 8.2.x and earlier servers that the NMC server manages.

Report categories appear as folders within the Legacy Reports folder. You can run these reports from the NMC GUI or from the command prompt.

Table 101 Legacy report categories

Category of report	Purpose
NetWorker Backup Statistics	Provide statistical information about save sets from the media database. Include summaries of size, number of files, and number of save sets that are backed up.

Table 101 Legacy report categories (continued)

Category of report	Purpose
NetWorker Backup Status	Provide status information about group completion and save set backups.
Inactive Files	Manages inactive files on a client or group, and sets the NetWorker software to automatically generate a list of inactive files in an environment.
Data Domain Statistics	Provides deduplication backup statistics for each selected NetWorker client. <i>NetWorker Data Domain Boost Integration Guide</i> provides more information.
NetWorker Clones	Provides the history of automatic and scheduled clone operations.
NetWorker Data Protection Policy	Provides details and summaries for VMware Data Protection Policies. The <i>NetWorker VMware Integration Guide</i> provides more information.
Snapshot Statistics	Provides details and summaries for Snapshot backups.

Report modes and types

All of the reports are listed within the report category folders. These folders are seen in the left pane of the **Reports** window. Each folder contains basic and drill-down reports. [Basic reports](#) on page 652, and [Drill-down reports](#) on page 652 provide detailed information.

Different icons represent the different types of reports:

Table 102 Report icons

Icon	Description
	Basic report
	Shared basic report
	Drill-down report
	Shared drill-down report

Interactive mode

Interactive mode displays a report with dynamic components, which allow you to update the report and display the modified results in real time. The effect of the dynamic components depends on whether a report is viewed as a table or as a chart.

Table view

The table view in interactive mode enables you to:

- Scroll through rows of the table.
- Sort, rearrange, or resize columns in the table, in the same way you sort data in other NMC windows.
- Use the **Choose Table Columns** menu to choose the columns to display, and the order in which to display them.
- Create and view drill-down reports.

The following images provides an example of the Group Summary report in table view.


Figure 74 Group Summary in table view

Server Name	Group	Policy Name	Workflow	Files Count	Save Sets Count	Successful Save	Failed Save Sets	Amount of Data	Target Size (KB)	Deduplication Ratio	Clones
bu-iddnserver	new_group	Server Protection	Server backup	237	3	0	0	1,956	1,956	0%	0
bu-iddnserver	new_group	new_policy	new_workflow	0	1	0	0	0	0	0%	0
bu-iddnserver	NMC server	Server Protection	NMC server backup	2580	5	2	3	46,041	46,041	0%	0
bu-iddnserver	SQL clients	Backup	SQL clients	0	0	0	0	0	0	0%	0
bu-iddnserver	new_group	new_policy	new_workflow	0	1	0	0	0	0	0%	0

Chart view

The chart view in interactive mode displays data in a chart format. You can switch back and forth between different chart formats by selecting a format from the Chart Type list.

The following image provides an example of a Group Summary report in Bar Chart view.

Figure 75 Group Summary in Bar Chart view

Some legacy reports in chart view provide a **Data Selector** option that provides the ability to control the information that appears in the chart. Use the **Data Selector** section to display interesting and useful data groupings in chart format.

For example, in a Group Summary by Server report that is displayed in Bar Chart format, the bar chart displays the amount of data in each group, and the Data Selector lists the "Server" control column, making it possible to see—in one place—a summary of groups across all servers, simply by moving through the list of servers in the Data Selector. This could be useful for finding the group that backed up the most data, or for balancing groups on servers.

You can limit the set of X and Y axes in the report by clearing one or more options from the **Chart Selector** boxes. This does not apply to Drive Utilization reports.

- For Drive Utilization reports, hover over a chart in Save Set view or Drive view to display a tool tip that includes this information:
 - Drive (Drive view only)
 - Save Set Name (Save Set view only)
 - Start Time
 - End Time
 - Client Name
 - Throughput (B/Sec)



Note: The tool tip feature for Drive Utilization reports is available only in interactive mode.

Document mode

Document mode displays data in a static table or chart report that resembles the view in Print Preview as shown by a PDF file viewer.

The following options are available with document mode:

- Orientation (portrait or landscape)
- Table or chart format

- Size (zoom level)

Table view

Document mode reports displayed in a table view contains several columns of information:

- One or more control columns represent qualitative information. For example, server name, save set name, and backup type. The control columns topics generally appear as X-axis data in charts.
- One or more data columns represent quantitative information. For example, Amount of data, number of files, number of save sets, and duration. The data columns topics generally appear as Y-axis data in charts. Each report gives subtotals and totals of all the columns of quantitative data that are shown in the report.
For example, a report on Save Set Details by Client provides a list of clients and the following quantitative information:

- Subtotals of the data columns for each of that client's save sets.
- Totals of all the data columns for each client.
- Totals of the data for all clients in the report.

The report allows you to easily parse the data, visually, on a per-client basis, on a save set-per-client basis, and for all clients in the report.

Chart view

In document mode, NMC displays two graphs for any chart type that displays X-Yaxes. If the top graph contains excessive Y-axis data, NMC may display truncated data in both graphs.

You cannot sort, rearrange, or resize the columns of a tabular report. Also, you cannot choose which columns to display, and the order in which to display the columns. Likewise, you cannot change the chart format while viewing a report. NetWorker software does not maintain any customized changes made while displaying a report in interactive mode (such as sorting or rearranging the columns in a table), except for charts (in Chart Type and Chart Selector). Instead, document mode displays the report in a standard table or chart format, as specified by the internal report definition within NetWorker software.

Unlike interactive mode, which provides you with a set of chart selection parameters that limit the displayed data, a report in document mode displays all the data. As a result, report views in document mode often consist of several screens. For this reason, the viewing choices in document mode include these navigation options to enable you to page through the output:

- First
- Previous
- Next
- Last

Interactive and document mode chart types

These chart types are available in both interactive and document mode:

- Bar chart
- Pie chart
- Plot chart

- Stacking bar chart
- Gantt chart (for Drive Utilization reports only -- more information is provided in the section [Device reports](#) on page 668)

When displaying reports in chart format, the size and appearance of the chart may differ depending on the orientation (portrait or landscape), and the presentation format—that is, whether viewing it in the **Console** window, or in other file formats, such as PDF, HTML, or PostScript. When displaying reports as charts in document mode, or when printing or exporting to HTML or PostScript, the charts are always displayed on a single page, regardless of their size. As a result, some data and labels may not display. To see full report details, view the chart in interactive mode.

The following table shows a simplified version of chart format options.

Table 103 Report chart formats






Format	Example	Description
Bar		Uses bars to illustrate the different types of data. For example, in a bar chart of a NetWorker Backup Statistics Server Summary report, the vertical bars show the amount of data that are backed up by each server. The additional lines show the corresponding numbers of files and save sets that are backed up by each server. The set of axes that are displayed in the report depends on the type of report. To select various elements for display, select or clear the boxes in the Chart Selector.
Plot		Displays data that are graphed as points along X and Y axes. To select various elements for display, select or clear the boxes in the Chart Selector.
Pie		Display data graphically as a percentage of a circular “pie.” When specifying this chart type from the Console window, the Chart Selector includes a radio button that allows the display of only one element, or axis, at a time. If an additional element is selected, it replaces the first. This limitation does not occur when this chart type is specified from the command prompt: • When this chart type is selected from the Console window, all applicable data axes are shown. • When this chart type is specified from the command prompt, only the requested information is included.
Stacking Bar		Displays data in a way that enables you to group and measure the data according to more than one category.

Table 103 Report chart formats (continued)

Format	Example	Description
		For example, use of a stacking bar chart to display a report that measures data according to only a single point of focus would display just a simple bar chart. Stacking bar chart reports generally include by in the name, such as by date or by host.
Gantt		When you view a Drive Utilization report as a chart, NMC automatically displays the data as a Gantt chart, and you cannot change the chart type. The Drive Utilization report is the only report that displays data as a Gantt chart. In Save Set view, the x-axis displays the time, and the y-axis displays save set data. Hovering over the chart in Save Set view displays a tool tip that provides this information. • Save set name • Start time • End time • Client name • Throughput value In Drive view, the x-axis displays the time, and the y-axis displays drive data. Hovering over the chart in Drive View displays a tool tip that provides the following information: • Drive • Start time • End time • Throughput value Chart axis selection

Note: Document mode can display more than one chart in the document. You can insert any or all available Y axes into the report. When you change to document mode, print or export a report, or save a configuration, NMC uses the axis selection that is currently set in the Chart Selector section of the **Configuration** tab. The exceptions to this are stacked bar and pie charts, which display all axes when the `gstclreport` command is used to generate a report.

Stacking bar charts

In interactive mode, movement of the cursor over a section of stacked color causes a pop-up legend to appear. The legend describes the data that are represented by that color. This chart type is inappropriate for complicated data in document mode, since the cursor does not display a legend describing the data that are represented by that

color. Instead, in document mode, select a different chart type (bar, pie, or plot) if the report data is complicated.

When specifying this chart type from the NMC GUI, the Chart Selector includes a radio button that enables the display of only one element, or axis, at a time. If an additional element is selected, it replaces the first. This limitation does not occur when this chart type is specified from the command prompt.

- When you specify this chart type from the NMC GUI, all applicable data axes are shown.
- When you specify this chart type from the command prompt, the `gstclreport` command only displays the requested information.

To appreciate the different ways in which you can use a stacking bar chart, consider these reports:

- NetWorker Backup Statistics Group Summary by Server — Shows statistics that are broken down by savegroup for each server. Different blocks of color are used for the amounts of data that are backed up by each group within the vertical bars that represent the amount of data backed up by servers.
- NetWorker Backup Statistics Server Summary — Shows data from only one focus, a server-centric point of view. If a stacking bar chart is selected to display a NetWorker Backup Statistics Server Summary, the chart would display solid bars of color to represent the servers. However, there would be no blocks of color within the bars, because the report focuses only on the server level. The result would therefore look like a simple bar chart.

Basic reports

The basic reports organize the collected data in a manner that focuses on a specific datazone component, time span, or attribute. For example:

- A Server Summary of Backup Statistics provides backup statistics in a server-centric manner.
- A Monthly Summary of Backup Statistics provides the backup statistics in a date-centric manner.

Select the basic report that best provides the information you need.

Drill-down reports

Drill-down reports present report information in a preset sequence of basic reports. You can save drill-down reports as customized reports in shared mode. You can only use drill-down reports from the NMC GUI. You cannot use drill-down reports from a command prompt.

Select a line of output in a report to generate information about the selected item in the next report in the drill-down sequence.

For example, configure a Policy Summary Over Time category report, and then click **View report**. From the generated Policy Summary report, double-click the output for one of the policies. NMC generates a Monthly Summary report of data for the policy that you selected in the Policy Summary report. In the Monthly Summary report, double-click a month. NMC generates a Daily Summary report of data that is generated on each day of the month that you selected in the Monthly Summary report. In the Daily Summary report, double-click a day. NMC generates a Client Summary report with information about clients for whom data was generated on the day that you selected in the Daily Summary report. In the Client Summary report, double-click one of the clients. NMC generates a Save Set Summary report of all save sets associated with the client that you selected in the Client Summary report, on the

day you selected in the Daily Summary report, in the month that you selected in the Monthly Summary, for the policy you selected in the Policy Summary report.

Note: In document mode for drill-down reports, the print and export commands do not print or export the entire drill-down report, just the basic report that is displayed.

Customized reports

A report that is included with NetWorker software is known as a canned reports, and includes several configuration parameters that allow the tailoring of report data. With customized reports, report versions can be configured—a single time—to fit the needs of the enterprise. These reports can then be saved and rerun whenever necessary, without having to be configured again. This feature saves time, especially with regularly run reports that include complex combinations of parameters. Customized reports can be run either on demand, or according to a preset schedule. The owner of a saved report can also allow it to be shared with all users.

The Hide Other Users Reports option toggles the view of reports between:

- The owner's reports (private and shared).
- The owner's reports, plus all shared custom reports.

[Customizing and saving reports](#) on page 682 and [Sharing a report](#) on page 683 provide more information.

Preconfigured reports

The **Reports** window contains two folders that contain preconfigured reports.

The Reports folder contains preconfigured reports that enable you to query for information about data that is created with a NetWorker 19.1 server. The Legacy Reports folder contains preconfigured reports that enable you to query for information about that created with a NetWorker 8.2.x and earlier server.

Preconfigured reports

The Reports folder contains preconfigured reports that enable you to generate reports about data that was created with a NetWorker 19.1 server :

Policy statistics

The Policy Statistics report category provides you with the ability to create reports that contain details and summary information about Data Protection Policy resources for each selected NetWorker server within the enterprise.

The Policy Statistics report category includes basic and drill-down reports.

Policy reports

NMC provides two types of reports that provide information about Policy resources: Policy Summary reports, and Policy Summary over time reports.

Policy Summary

A basic report that provides information that is gathered from the media database and client file indexes about data that are generated by backup and clone actions in all workflows that are associated with a Policy resource. The reported Information includes the following statistics:

- NetWorker server—Name of the NetWorker server.
- Policy—Name of the Policy resource.

- **File count**—Total number of files.
- **Save Sets Count**—Total number of save sets that are stored in the media database.
- **Amount of data**—Total size of backup data that is stored on media.
- **Target size**— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- **Deduplication ratio**— Deduplication ratio for the data.
- **Clone count**—Total number of clone save sets that are stored in the media database.
- **Clone size**—Total size of cloned data that is stored on media.

Policy Summary over time

Drill-down reports that provide a point-in-time basic report about the data that are generated by backup and clone actions in all workflows that are associated with a Policy resource. You can generate the following types of drill-down reports:

- **Policy Summary**—A basic report that provides a summary of all policies that are associated with the selected NetWorker servers.
- **Monthly Summary**—A summary of monthly activities for the policy that you selected in the Policy Summary report.
- **Daily Summary**—A summary of daily activities for the month that you selected in the Monthly Summary report.
- **Client Summary**—A summary of client information for the day that you selected in the Daily Summary report.
- **Save Set Details**—A summary of information for each save set generated for the client that you selected in the Client Summary report.

Report parameters

The **Parameters** section allows you to define the selection criteria to generate a customized report:

- **NetWorker server**—By default, the report generates information about all the NetWorker servers that are managed by the NMC server. The **Server Name Selected** field provides a list of NetWorker server on which to report information. The **Server Name Available** field provides a list of NetWorker servers for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of NetWorker servers on which to report.
- **Policy**— By default, the report generates information about all policies that are configured on each NetWorker server. The **Policy Name Selected** field provides a list of policies on which to report information. The **Policy Name Available** field provides a list of policies for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Policy resources on which to report.
- **Workflow start and end times**— By default the report generates information about all workflows that started within one day of the current time. Use the **From** and **To** arrows to select a new date range.

Group reports

NMC provides three types of reports that provide information about Group resources: Group Summary reports, Group Details reports, and Group Summary over time reports.

Report parameters

The **Parameters** section allows you to define the selection criteria to generate a customized report:

- **NetWorker server**—By default, the report generates information about all the NetWorker servers that are managed by the NMC server. The **Server Name Selected** field provides a list of NetWorker server on which to report information. The **Server Name Available** field provides a list of NetWorker servers for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of NetWorker servers on which to report.
- **Policy**— By default, the report generates information about all policies that are configured on each NetWorker server. The **Policy Name Selected** field provides a list of policies on which to report information. The **Policy Name Available** field provides a list of policies for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Policy resources on which to report.
- **Group**—By default, the report generates information about all groups that are configured on each NetWorker server. The **Group Name Selected** field provides a list of groups on which to report information. The **Group Name Available** field provides a list of groups for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Group resources on which to report.
- **Workflow start and end times**— By default the report generates information about all workflows that started within one day of the current time. Use the **From** and **To** arrows to select a new date range.

Group Summary

A basic report that provides a list of groups in each policy resource on NetWorker servers that are managed by the NMC server. The report provides the following information:

- **NetWorker server**—Name of the NetWorker server.
- **Group**—Name of the Group resource.
- **Policy Name**—Name of the Policy resource that is associated with the Group resource.
- **Workflow**—Name of the workflow that is associated with the Group resource.
- **File count**—Total number of files.
- **Save Sets Count**—Total number of save sets that are stored in the media database.
- **Amount of data**—Total size of backup data that is stored on media.
- **Target size**— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- **Deduplication ratio**— Deduplication ratio for the data.
- **Clone count**—Total number of clone save sets that are stored in the media database.

- Clone size—Total size of cloned data that is stored on media.

Group Details

A basic report that provides details about all groups on all NetWorker servers that are managed by the NMC server. The report provides the following information:

- NetWorker server—Name of the NetWorker server.
- Group—Name of the Group resource.
- Policy Name—Name of the Policy resource that is associated with the Group resource.
- Workflow—Name of the workflow that is associated with the Group resource.
- File count—Total number of files.
- Save Sets Count—Total number of save sets that are stored in the media database.
- Successful save sets—Total number of backup or clone save sets that are created successfully by the action task.
- Failed save sets—Total number of failed attempts to create backup or clone save sets by the action task.
- Amount of data—Total size of backup data that is stored on media.
- Target size— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- Deduplication ratio— Deduplication ratio for the data.
- Clone count—Total number of clone save sets that are stored in the media database.
- Clone size—Total size of cloned data that is stored on media.
- Successful clones—Total number of clone save sets that are created successfully Group resource.
- Failed Clones—Total number of failed attempts to create a clone save set Group resource.

Group Summary Over Time

Drill-down reports that provides a point-in-time basic report about the data that are generated by all groups that are associated with a Policy resource. You can generate the following types of drill-down reports:

- Group Summary—A basic report that provides a summary of all groups that are associated with the selected NetWorker servers.
- Monthly Summary—A summary of monthly activities for the group that you selected in the Group Summary report.
- Daily Summary—A summary of daily activities for the month that you selected in the Monthly Summary report.
- Client Summary—A summary of client information for the day that you selected in the Daily Summary report.
- Save Set Details—A summary of information for each save set generated for the client that you selected in the Client Summary report.

Workflow reports

NMC provides two types of reports that provide information about Workflow resources: Workflow Summary reports, and Workflow Details reports.

Report parameters

The **Parameters** section allows you to define the selection criteria to generate a customized report:

- **NetWorker server**—By default, the report generates information about all the NetWorker servers that are managed by the NMC server. The **Server Name Selected** field provides a list of NetWorker server on which to report information. The **Server Name Available** field provides a list of NetWorker servers for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of NetWorker servers on which to report.
- **Policy**— By default, the report generates information about all policies that are configured on each NetWorker server. The **Policy Name Selected** field provides a list of policies on which to report information. The **Policy Name Available** field provides a list of policies for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Policy resources on which to report.
- **Workflow**—By default, the report generates information about all workflows that are configured on each NetWorker server. The **Workflow Name Selected** field provides a list of workflows on which to report information. The **Workflow Name Available** field provides a list of workflows for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Workflow resources on which to report.
- **Workflow start and end times**— By default the report generates information about all workflows that started within one day of the current time. Use the **From** and **To** arrows to select a new date range.

Workflow Summary

A basic report that provides a list of groups for the resources that you selected in the **Parameter** section. The report includes the following Information:

- **NetWorker server**—Name of the NetWorker server.
- **Policy Name**—Name of the Policy resource that is associated with the Workflow resource.
- **Number of runs**—Number of times that the Workflow resource has run.
- **Successful**—Number of times that the run of the actions in the workflow have completed successfully.
- **Failed**—Number of times the run of the actions in the workflow run failed.
- **Interrupted**—Number of items that the run of the actions in the workflow were interrupted.
- **Total duration**— Total amount of time that the actions in the workflow have run.
- **File count**—Total number of files.
- **Save Sets Count**—Total number of save sets that are stored in the media database.
- **Amount of data**—Total size of backup data that is stored on media.
- **Target size**— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.

- Deduplication ratio—Deduplication ratio for the data.
- Clone count—Total number of clone save sets that are stored in the media database.
- Clone size—Total size of cloned data that is stored on media.

Workflow Details

A basic report that provides detailed information about the backup and clone data that are generated by all actions that are associated with a Workflow resource. The report includes the following Information:

- NetWorker server—Name of the NetWorker server.
- Policy Name—Name of the Policy resource that is associated with the Workflow resource.
- Workflow start time—Start time of the workflow.
- Total duration—Total amount of time that the actions in the workflow have run.
- Workflow status—Status of the workflow. For example, successful or failed.
- Name of the Group that is associated to the workflow.
- Successful save sets—Total number of backup or clone save sets that are created successfully by the action task.
- Failed save sets—Total number of failed attempts to create backup or clone save sets by the action task.
- File count—Total number of files.
- Save Sets Count—Total number of save sets that are stored in the media database.
- Amount of data—Total size of backup data that is stored on media.
- Target size—Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- Deduplication ratio—Deduplication ratio for the data.
- Clone count—Total number of clone save sets that are stored in the media database.
- Clone size—Total size of cloned data that is stored on media.
- Successful clones—Total number of clone save sets that are created successfully by clone actions in the workflow.
- Failed Clones—Total number of failed attempts to create a clone save set by clone actions in the workflow.

Action reports

NMC provides four types of reports that provide information about Action resources: Action Summary By Group reports, Action Summary By Policy and Workflow reports, Action Details reports, and Action Details By workflow reports.

Action Summary reports

NMC provides two types of summary reports that provide information about Action resources: Action Summary By Group reports, and Action Summary By Policy and Workflow reports.

Action Summary By Group

A basic report that provides a list of actions that are associated with each Group resource for a NetWorker server. The report provides the following information:

- NetWorker server—Name of the NetWorker server.
- Policy Name—Name of the Policy resource that is associated with the Group resource.
- Group—Name of the Group that is associated with the Action resource
- Action—Name of the Action resource.
- File count—Total number of files.
- Save Sets Count—Total number of save sets that are stored in the media database.
- Amount of data—Total size of backup data that is stored on media.
- Target size— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- Deduplication ratio— Deduplication ratio for the data.
- Clone count—Total number of clone save sets that are stored in the media database.
- Clone size—Total size of cloned data that is stored on media.

Action Summary By Policy or Workflow

A basic report that provides a list of actions that are associated with each Group resource for a NetWorker server. The report provides the following information:

- NetWorker server—Name of the NetWorker server.
- Policy Name—Name of the Policy resource that is associated with the Group resource.
- Workflow—The name of the Workflow that is associated with the Action resource.
- Action—Name of the Action resource.
- File count—Total number of files.
- Save Sets Count—Total number of save sets that are stored in the media database.
- Amount of data—Total size of backup data that is stored on media.
- Target size— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- Deduplication ratio— Deduplication ratio for the data.

- **Clone count**—Total number of clone save sets that are stored in the media database.
- **Clone size**—Total size of cloned data that is stored on media.

Report parameters

The **Parameters** section allows you to define the selection criteria to generate a customized report:

- **NetWorker server**—By default, the report generates information about all the NetWorker servers that are managed by the NMC server. The **Server Name Selected** field provides a list of NetWorker server on which to report information. The **Server Name Available** field provides a list of NetWorker servers for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of NetWorker servers on which to report.
- **Policy**— By default, the report generates information about all policies that are configured on each NetWorker server. The **Policy Name Selected** field provides a list of policies on which to report information. The **Policy Name Available** field provides a list of policies for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Policy resources on which to report.
- **Workflow**—Action Summary By Policy or Workflow report only. By default, the report generates information about all workflows that are configured on each NetWorker server. The **Workflow Name Selected** field provides a list of workflows on which to report information. The **Workflow Name Available** field provides a list of workflows for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Workflow resources on which to report.
- **Group**—Action Summary By Group report only. By default, the report generates information about all groups that are configured on each NetWorker server. The **Group Name Selected** field provides a list of groups on which to report information. The **Group Name Available** field provides a list of groups for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Group resources on which to report.
- **Workflow start and end times**— By default the report generates information about all workflows that started within one day of the current time. Use the **From** and **To** arrows to select a new date range.

Action Detail reports

NMC provides two types of detail reports that provide information about Action resources: Action Details reports, and the Action Details By Workflow reports.

Report parameters

The **Parameters** section allows you to define the selection criteria to generate a customized report:

- **NetWorker server**—By default, the report generates information about all the NetWorker servers that are managed by the NMC server. The **Server Name Selected** field provides a list of NetWorker server on which to report information. The **Server Name Available** field provides a list of NetWorker servers for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of NetWorker servers on which to report.
- **Policy**— By default, the report generates information about all policies that are configured on each NetWorker server. The **Policy Name Selected** field provides a list of policies on which to report information. The **Policy Name Available** field provides a list of policies for which you do not want to report information about.

Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Policy resources on which to report.

- **Workflow**—By default, the report generates information about all workflows that are configured on each NetWorker server. The **Workflow Name Selected** field provides a list of workflows on which to report information. The **Workflow Name Available** field provides a list of workflows for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Workflow resources on which to report.
- **Action**—Action Details report only. By default, the report generates information about all actions that are configured on each NetWorker server. The **Action Name Selected** field provides a list of actions on which to report information. The **Action Name Available** field provides a list of actions for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Action resources on which to report.
- **Workflow start and end times**— By default the report generates information about all workflows that started within one day of the current time. Use the **From** and **To** arrows to select a new date range.

Action Details

A basic report that provides detailed information about the backup and clone data generated by the resources that are defined in the Parameters section. The report includes the following Information:

- **NetWorker server**—Name of the NetWorker server.
- **Policy Name**—Name of the Policy resource that is associated with the Workflow resource.
- **Workflow**—Name of the Workflow resource that contains the action.
- **Action**—Name of the Action resource.
- **Action Type**—The action type that is defined for the Action resource. For example, Backup, Clone, or Check Connectivity.
- **Action Start Time**—The time that the task in the Action resource starts.
- **Status**—Status of the task that is performed by the Action resource. For example, succeeded or failed.
- **Group**—Name of the group that is associated with the Action resource.
- **File count**—Total number of files.
- **Save Sets Count**—Total number of save sets that are stored in the media database.
- **Successful save sets**—Total number of backup or clone save sets that are created successfully by the action task.
- **Failed save sets**—Total number of failed attempts to create backup or clone save sets by the action task.
- **Amount of data**—Total size of backup data that is stored on media.
- **Target size**— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- **Deduplication ratio**— Deduplication ratio for the data.
- **Clone count**—Total number of clone save sets that are stored in the media database.

- Clone size—Total size of cloned data that is stored on media.
- Successful clones—Total number of clone save sets that are created successfully Group resource.
- Failed Clones—Total number of failed attempts to create a clone save set Group resource.

Action Details By Workflow

Drill-down reports that provide a point-in-time basic report about the data generated by the resources that are defined in the Parameter section. You can generate the following types of drill-down reports:

- Workflow Summary—A basic report that provides a summary of information about all actions in all workflows that are associated with the selected NetWorker servers.
- Workflow Details—A basic report that provides a summary of all actions in the workflow that you selected in the Workflow Summary report.
- Action Details—A basic report that provides details about each action in the Workflow that you selected in the Workflow Details report.
- Client Summary—A basic report that provides a summary of information about all actions in the client that you selected in the Action Details report.

Client reports

NMC provides three types of reports that provide information about Client resources: Client Summary reports, Client Details reports, and Client Summary by Group reports.

Report parameters

The **Parameters** section allows you to define the selection criteria to generate a customized report:

- **NetWorker server**—By default, the report generates information about all the NetWorker servers that are managed by the NMC server. The **Server Name Selected** field provides a list of NetWorker server on which to report information. The **Server Name Available** field provides a list of NetWorker servers for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of NetWorker servers on which to report.
- **Policy**— By default, the report generates information about all policies that are configured on each NetWorker server. The **Policy Name Selected** field provides a list of policies on which to report information. The **Policy Name Available** field provides a list of policies for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Policy resources on which to report.
- **Group**—Client Summary by Group report only. By default, the report generates information about all groups that are configured on each NetWorker server. The **Group Name Selected** field provides a list of groups on which to report information. The **Group Name Available** field provides a list of groups for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Group resources on which to report.
- **Workflow**—Client Summary report only. By default, the report generates information about all workflows that are configured on each NetWorker server. The **Workflow Name Selected** field provides a list of workflows on which to report information. The **Workflow Name Available** field provides a list of workflows for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Workflow resources on which to report.

- **Workflow start and end times**— By default the report generates information about all workflows that started within one day of the current time. Use the **From** and **To** arrows to select a new date range.

Client Summary

A basic report that provides a list of clients for the resources that you selected in the **Parameter** section. The report includes the following Information:

- NetWorker server—Name of the NetWorker server.
- Client Name—Name of the Client resource.
- File count—Total number of files.
- Save Sets Count—Total number of save sets that are stored in the media database.
- Amount of data—Total size of backup data that is stored on media.
- Target size— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- Deduplication ratio— Deduplication ratio for the data.
- Clone count—Total number of clone save sets that are stored in the media database.
- Clone size—Total size of cloned data that is stored on media.

Client Details

A basic report that provides detailed information about the backup and clone data that are generated for a Client resource. The report includes the following Information:

- NetWorker server—Name of the NetWorker server.
- Client Name—Name of the Client resource.
- Policy Name—Name of the Policy resource that is associated with the Workflow resource.
- Workflow—Name of the Workflow that is associated with the Client resource.
- Group—Name of the Group resource.
- Workflow start time—Start time of the workflow.
- Status—Status of the save set in the media database. For example, succeeded or failed.
- File count—Total number of files.
- Save set size—The original size of the save set, as recorded in the media database.
- Target size— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- Deduplication ratio— Deduplication ratio for the data.

Client Summary By Group

Drill-down reports that provide a point-in-time basic report about the data generated for client in the Group resources that are defined in the **Parameter** section. You can generate the following types of drill -own reports:

- **Group Summary**—A basic report that provides summary information about all groups that are associated with the NetWorker servers selected in the **Parameters** section.
- **Client Summary**—A basic report that provides summary information about all clients that are associated with the group that you selected in the Group Summary report.

Save set reports

NMC provides one basic report, the Save Set Details report. This report provides detailed information about the save sets stored in the media database of a NetWorker server.

Report Parameters

The **Parameters** section allows you to define the selection criteria to generate a customized report:

- **NetWorker server**—By default, the report generates information about all the NetWorker servers that are managed by the NMC server. The **Server Name Selected** field provides a list of NetWorker server on which to report information. The **Server Name Available** field provides a list of NetWorker servers for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of NetWorker servers on which to report.
- **Policy**— By default, the report generates information about all policies that are configured on each NetWorker server. The **Policy Name Selected** field provides a list of policies on which to report information. The **Policy Name Available** field provides a list of policies for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Policy resources on which to report.
- **Client**—By default, the report generates information about all the save sets for each client that is configured on the selected NetWorker servers. The **Client Name Selected** field provides a list of clients on which to report information. The **Client Name Available** field provides a list of clients for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of clients on which to report.
- **Workflow**—By default, the report generates information about all workflows that are configured on each NetWorker server. The **Workflow Name Selected** field provides a list of workflows on which to report information. The **Workflow Name Available** field provides a list of workflows for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Workflow resources on which to report.
- **Save set name**—By default, the report generates information about all save sets for the selected clients on the selected NetWorker servers. The **Save Set Name Selected** field provides a list of save sets on which to report information. The **Save Set Name Available** field provides a list of save sets for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of save sets on which to report.
- **Action type**—By default, the report generates information about all action types for the selected clients on the selected NetWorker servers. The **Action Type Selected** field provides a list of action types on which to report information. The **Action Type Available** field provides a list of action types for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of action types on which to report.

- **Workflow start and end times**— By default the report generates information about all workflows that started within one day of the current time. Use the **From** and **To** arrows to select a new date range.

Save Set Details report

A basic report that provides detailed information about the backup and clone save sets that are stored on a NetWorker server. The report includes the following Information:

- NetWorker server—Name of the NetWorker server.
- Client Name—Name of the Client resource.
- Save Set Name—Name of the save set.
- Save Set ID—The SSID of the save set.
- Clone ID—The cloneid of the save set.
- Action Type—The action type that is defined for the Action resource. For example, Backup, Clone, or Check Connectivity.
- Policy Name—Name of the Policy resource that is associated with the Workflow resource.
- Workflow—Name of the workflow that is associated with the Group resource.
- Group—Name of the Group resource.
- Workflow start time—Start time of the workflow.
- Status—The status of the save set. For example, succeeded or failed.
- File count—Total number of files.
- Save set size—The size of the save set, as recorded in the media database.
- Target size— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- Deduplication ratio— Deduplication ratio for the data.

Manual saves

Users must run the following commands in order to view the save set reports.

- `save -b Default "<C:\Program Files\Java>"`
- `mminfo -q "name=C:\Program Files\Java" -r "volume,client,name,ssid,nfiles,savetime(24),sumsize"`

Monthly and Daily Summary reports

NMC provides Monthly and Daily Summary reports that provide information backup and clone data on a NetWorker server.

Report parameters

The **Parameters** section allows you to define the selection criteria to generate a customized report:

- **NetWorker server**—By default, the report generates information about all the NetWorker servers that are managed by the NMC server. The **Server Name Selected** field provides a list of NetWorker server on which to report information. The **Server Name Available** field provides a list of NetWorker servers for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of NetWorker servers on which to report.

- **Policy**— By default, the report generates information about all policies that are configured on each NetWorker server. The **Policy Name Selected** field provides a list of policies on which to report information. The **Policy Name Available** field provides a list of policies for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Policy resources on which to report.
- **Group**—By default, the report generates information about all groups that are configured on each NetWorker server. The **Group Name Selected** field provides a list of groups on which to report information. The **Group Name Available** field provides a list of groups for which you do not want to report information about. Use the **Add**, **Add All**, **Remove**, and **Remove All** buttons to modify the list of Group resources on which to report.
- **Workflow start and end times**— By default the report generates information about all workflows that started within one day of the current time. Use the **From** and **To** arrows to select a new date range.

Monthly and Daily Summary

The Monthly Summary report provides monthly summary information about groups in the months that are within the range that is specified in the Workflow Start and Workflow End Time attributes. The Daily Summary report provides daily summary information about groups in the days that are within the range that is specified in the Workflow Start and Workflow End Time attributes. The Summary reports provide the following information:

- **Month**—Monthly Summary only. The month in which the report data was created.
- **Date**—Daily Summary only. The day in which the report data was created.
- **File count**—Total number of files.
- **Save Sets Count**—Total number of save sets that are stored in the media database.
- **Amount of data**—Total size of backup data that is stored on media.
- **Target size**— Size of the save set on the target backup or clone device. When the target device is a Data Domain device, the value represents the size of the data after deduplication. When the target device is an AFTD device, the value is the same size as the original save set size.
- **Deduplication ratio**— Deduplication ratio for the data.
- **Clone count**—Total number of clone save sets that are stored in the media database.
- **Clone size**—Total size of cloned data that is stored on media.

NetWorker recovery reports

The recovery reports, available from the **Reports** task pane in the **NMC GUI**, allow you to view the history of recovery operations that have been performed by a NetWorker Server. Also, the NMC Server checks for new recovery operations and stores the recover statistics in the NMC database every 12 hours, and each time a scheduled backup completes.

You can review reports in both chart and table modes. Table mode set is the default mode. You can generate four different types of recover reports:

- Server Summary
- Client Summary

- Recover Details
- Recover Summary Over Time

The NMC Server gathers recover job history every 12 hours and on completion of every scheduled backup action. Recovery reports will not display information about recovery history within 12 hours of when you run the report.

Types of NetWorker recovery reports and configuration

The NetWorker recovery report category includes basic and drill-down reports. The different types of reports that are included within the NetWorker Recover Statistics report category provide recover statistics for each selected NetWorker Server within the enterprise.

The **Configuration** tab allows you to limit the scope of the report that was selected.

The parameters available within the NetWorker Recovery report category are described in this table. The specific parameters available depend on which NetWorker Recovery Statistics report is selected.

Table 104 NetWorker recovery statistics parameters

Parameter	Description	Options
NetWorker Server	Managed hosts within the enterprise.	Selected server names.
Source Client Name	One or more clients whose data is being recovered.	Selected client names.
Target Client	The client where the data is being recovered to.	Selected target client names.
Initiating Client	The client that started the recover.	n/a
User	Name of the user who started the recover.	Selected user names.
Size	The size of the recover.	n/a
Number of files	For file system recoveries, the number of files in the recover.	n/a
Start time/End time	Limits the report to a specified time range. The date/time format available depends on the language locale of the operating system.	Start time of recover/end time of recover.
Completion Status	Final status of the recover.	• Successful • Failed

The parameters available for each report type in the NetWorker Recovery Statistics report category are listed in the user interface.

Recovery Statistics basic reports

Within the NetWorker Recovery Statistics report category, choose any of the basic reports that are listed in the user interface. Once a report is chosen, the **Configuration** tab displays boxes with lists of the selected parameters for that report.

To exclude unwanted parameters from the report, delete them from the list. [Customizing and displaying report output](#) on page 679 provides information on selecting and removing parameters.

Recovery Statistics drill-down report

This drill-down report consists of multiple NetWorker Recovery Statistics basic reports, which are connected in a predetermined sequence. [Drill-down reports](#) on page 652 provides general information about drill-down reports.

The configuration parameters for a drill-down report are the same as the parameters for the top-level report in the report sequence. Thus, if the top layer of the drill-down report is a Server Summary report, the configuration parameters are the same as they would be for the basic report, Server Summary.

When a report is chosen, the **Configuration** tab displays boxes that list the selected parameters for the top-level report.

To exclude unwanted parameters from the report, delete them from the list. [Customizing and displaying report output](#) on page 679 provides information on selecting and removing parameters.

Recover Summary Over Time

Recover Summary Over Time is a drill-down report sequence that allows you to explore the history of recover jobs that were performed by NetWorker Servers over a period.

To generate the Recover Summary Over Time report, you must first specify the same parameters as those in the Server Summary report, which is the first report that is displayed in the sequence.

To drill-down to the client level, perform one of the following, depending on the viewing mode:

- When in table mode, double-click any individual row referencing the desired NetWorker Server.
- When in chart mode, click anywhere in the chart area of the desired NetWorker Server.

The Client Summary report for the selected NetWorker Server appears. Return to the Server Summary report to select another server to explore.

To drill-down to the Recover Details level, perform one of the following, depending on the viewing mode:

- When in table mode, double-click any individual row referencing the desired NetWorker Client.
- When in chart mode, click anywhere in the chart area of the desired NetWorker Client.

The Recover Details report for the selected NetWorker Client appears. Return to the Client Summary report to select another client to explore.

Recovery data retention policy and configuration

The retention policy for the recover statistics that are used to generate these reports can be set with the other retention policies currently defined from the Data Retention page in the Reports task pane. The default retention policy for these statistics is one year.

Device reports

Device reports provide information about the way devices are being used. They show scheduled and manual backup activity on one or more selected devices over time. You

can identify periods of heavy activity or inactivity. Device reports aid NetWorker administrators in performance tuning, and they help identify bottlenecks. For example, if all drives are being used continuously for a long period, at maximum throughput, backup speeds may improve by adding tape drives or moving clients to another backup server.

Types of Device reports and configuration

The Devices report category includes only one report, the Drive Utilization report. This report, which is a drill-down report, supports NetWorker servers running NetWorker software release 7.3 or later. These versions are now unsupported. The report includes backup activity data for all device types, including advanced file type devices and digital data storage devices.

When viewing a Drive Utilization report as a chart, it is automatically displayed as a Gantt chart, where the backup activity level of one or more devices is depicted in relation to time. Unlike with other reports, you cannot choose an alternate chart type.

Placing the cursor over the chart in Save Set view displays a tool tip that provides this information:

- Save set name
- Start time
- End time
- Client name
- Throughput value

Placing the cursor over the chart in Drive View displays a tool tip that provides this information:

- Drive
- Start time
- End time
- Throughput value

Note: One of the activities in the Drive Utilization report is throughput. Since the Drive Utilization Report provides data for backup activities only, throughput values will normally be non-zero. However, zero (0) is considered a valid throughput value.

Event reports

These reports provide summary information about current events on NetWorker and Console servers within the Enterprise. Additional details about a particular event can be displayed, including annotation contents. While the Events window within the NetWorker Console displays the current events of the NetWorker servers, the Event reports provide additional features. The reports enable you to organize, export, and print the event data.

Event reports can include this information:

- Number of events
- Priority of events
- Category of events
- Server name
- Server type

- Event time
- Notes and annotations

 **Note:** When an event has been resolved, it does not remain in the records.

Types of event reports and configuration

The Events report category includes both basic and drill-down reports. The report's Configure tab allows you to limit the scope of the report.

The Event parameters are described in this table. The specific parameters available depend on which Event report is being configured.

 **Note:** Data retention policies do not have any impact on Event reports.

Table 105 Event parameters

Configuration parameter	Description	Options
Server Name	Selects one or more managed hosts.	Selected server names.
Server Type	Selects some or all server types in the enterprise. Only the names of servers that have current events are shown.	Console NetWorker
Priority	Selects only priority events. Priority represents the relative severity of the event.	Warning Waiting Notice Info Emergency Critical Alert
Category	Selects only category events, or all categories. Category refers to the source of the event.	Database Backup Registration Savegroup
Event Time	Selects a time range. This parameter applies only to the Annotation Details report.	Event time (range)

Event basic reports

Within the Events report category, select any of the basic reports that are listed in the user interface. When a report has been chosen, the Configuration tab displays boxes listing the selected parameters for that report.

To exclude unwanted parameters from the report, remove them from the list. [Customizing and displaying report output](#) on page 679 provides information about selecting and removing parameters.

Event drill-down reports

The drill-down reports consist of multiple Event basic reports, which are connected in a predetermined sequence. [Drill-down reports](#) on page 652 provides general information about drill-down reports.

The configuration parameters for a drill-down report are the same as the parameters for the top-level report in the report sequence. Thus, if the top layer of the drill-down report is a Server Summary report, the configuration parameters are the same as they would be for the basic report, Server Summary. When a report has been chosen, the Configuration tab displays boxes listing the selected parameters for the top-level report. To exclude unwanted parameters from the report, remove them from the list. [Customizing and displaying report output](#) on page 679 provides information on selecting and removing parameters.

Host reports

The Hosts report category includes only basic reports. There are two basic reports, as described in this table.

Table 106 Host reports

Report name	Purpose	Configuration parameters	Default
Host List	Provides an overview of servers in the enterprise, including: - Whether the Capture Events feature is enabled for the server. - Whether the Gather Report Data feature is enabled for the server. - Where the server is located in the enterprise path.	None	All servers
Enterprise Inventory	Allows movement through the Enterprise. Limit the report's scope by first viewing one of the lower-level folders within the Enterprise: Start from Enterprise folder.	Enterprise Path	Start from Enterprise folder

Table 106 Host reports (continued)

Report name	Purpose	Configuration parameters	Default
	Start from selected folder.		

[Enterprise](#) on page 756 provides a description of the Enterprise and its folders.

User reports

The Users report category provides information on NetWorker Console user activity. NMC Server Management provides information about NetWorker Console users and creating user accounts.

The Users report category includes only basic reports, no drill-down reports. The Full Name and Description information appears in the User reports only if this information was specified when the user was created.

Preconfigured legacy reports

The Legacy Reports folder provides you with the ability to generate reports about data that was created with a NetWorker 8.2.x and earlier server.

NetWorker backup statistics reports

The different types of reports that are included within the NetWorker Backup Statistics report category provide backup statistics for each selected NetWorker server within the enterprise.

NetWorker Backup Statistics reports may include this information:

- Amount of data that is backed up.
- Number of files that are backed up.
- Number of save sets that are backed up.

Types of NetWorker backup statistics reports and configuration

The NetWorker Backup Statistics report category includes basic and drill-down reports.

The Configure tab allows you to limit the scope of the report that was selected.

The parameters available within the NetWorker Backup Statistics report category are described in this table. The specific parameters available depend on which NetWorker Backup Statistics report is selected.

Table 107 NetWorker backup statistics parameters

Parameter	Description	Options
Server Name	Selects managed hosts within the enterprise.	Selected server names
Group Name	Selects one or more groups.	Selected group names
Client Name	Selects one or more clients.	Selected client names
Save Set Name	Selects one or more save sets.	Selected save set names

Table 107 NetWorker backup statistics parameters (continued)

Parameter	Description	Options
Backup Type	Selects one or more file types.	List of supported file types
Level	Select one or more backup levels.	List of backup levels such as, Full, Incremental, Skip, synthetic full, or Level 1–9
Save Time	Limits the report to a specified time range. The default range is one day for save set details reports. The date/time format available depends on the language locale of the operating system.	Save time (range)

The parameters available for each report type in the NetWorker Backup Statistics report category are listed in the user interface.

Save set data retention policy and configuration

Settings for the save set retention policy impact the data that is available to the NetWorker Backup Statistics reports. If a save set retention policy of six months is specified, NetWorker software cannot query the database for a time range that extends back more than six months. The report cannot display data that has expired because that data has been removed from the database. Thus, even if a save time parameter of one year is specified, the report can display only six months of data if the limit of the save set retention policy is six months.

Backup statistics basic reports

Within the NetWorker Backup Statistics report category, choose any of the basic reports that are listed in the user interface. Once a report is chosen, the Configuration tab displays boxes with lists of the selected parameters for that report. To exclude unwanted parameters from the report, delete them from the list. [Customizing and displaying report output](#) on page 679 provides information on selecting and removing parameters.

 **Note:** These basic reports do not distinguish between regular and deduplication clients.

Backup statistics drill-down reports

Drill-down reports consist of multiple NetWorker Backup Statistics basic reports, which are connected in a predetermined sequence. [Drill-down reports](#) on page 652 provides general information about drill-down reports.

The configuration parameters for a drill-down report are the same as the parameters for the top-level report in the report sequence. Thus, if the top layer of the drill-down report is a Monthly Summary report, the configuration parameters are the same as they would be for the basic report, Monthly Summary.

When a report is chosen, the Configuration tab displays boxes that list the selected parameters for the top-level report. To exclude unwanted parameters from the report, delete them from the list. [Customizing and displaying report output](#) on page 679 provides information on selecting and removing parameters.

NetWorker backup status reports

The NetWorker Backup Status reports consolidate information about the success of scheduled group backups. As with the NetWorker Backup Statistics reports, these reports can provide either an enterprise-wide, or a more focused summary of activity over a specified time range.

The NetWorker Backup Status reports provide the same basic function as selecting Show Details for a group in the **Monitoring** window of the **Administration** window. The NetWorker Backup Status reports, however, allow you to select the scope and level of detail.

The report calculates the amount of time that is taken by each backup group individually. Consequently, if several groups run in parallel, their total combined backup time is greater than the time elapsed between the start of the first group and the completion of the last group. For example:

- Group A starts at 13:00, and completes at 15:00.
- Group B starts at 13:30, and completes at 15:30.

Although the groups both completed within a 2.5-hour period, the total group runtime is counted as 4 hours.

NetWorker Backup Status reports can include this information:

- Total group runs
- Totals of successful, failed, and interrupted group runs
- Success ratio
- Backup duration
- Backup level
- Backup type
- Save type

Backup type and save type information

Backup type is one of the configuration parameters for both NetWorker Backup Statistics and NetWorker Backup Status reports, and it is one of the fields of information that is included in these reports. The backup type indicates whether the files backed up were regular files, bootstrap files, indexes, or a particular database file.

Specialized NetWorker modules (such as NetWorker Module for SAP) are used to back up the various databases. Most of these modules apply a distinct prefix when backing up a save set. This prefix enables NetWorker software to identify the backup type and include it in the reports.

A couple of the Backup Status reports (Save Set Details and Save Set Details by Client) include an additional field of information that is called save type. The save type can be any one of the following:

- Bootstrap
- Index
- Save
- Save (backup command)

Types of NetWorker backup status reports and configuration

The NetWorker Backup Status Report category includes both basic and drill-down reports. The report's Configure tab allows you to limit the scope of the report

selected. The choice of available parameters depends on which report is to be generated.

The parameter options available within the NetWorker Backup Status Report category are described in this table.

Table 108 NetWorker backup status parameters

Parameter	Description	Options
Server Name	Selects one or more NetWorker servers.	Selected server names.
Group Name	Selects one or more savegroups.	Selected group names.
Group Start Time	Limits the report to a specified time range. The default range is one day for save set details reports.	Start and end dates.
Client Name	Selects one or more clients.	Selected client names.
Save Set Name	Selects one or more save sets.	Selected save set names.
Backup Type	Selects one or more file types.	List of supported file types.
Level	Selects one or more backup levels.	• Full • Incremental • Skip • Level 1–9 (Partial list of options)
Status	Selects status.	• Successful • Failed • Interrupted

The parameters available for each report type are listed in the user interface.

Completion data retention and NetWorker backup status

The settings for the completion data policy impact the data that is available to the NetWorker Backup Status reports. The report cannot display data that has expired, because it has been removed from the database.

Thus, even if a one-year time range is specified for the Group Start Time parameter, the report displays only six months if the limit of the completion data policy is six months.

Backup status basic reports

Within the NetWorker Backup Status report category, choose any of the basic reports that are listed in the user interface. When a report has been chosen, the Configuration tab displays boxes listing the selected parameters for that report. To exclude unwanted parameters from the report, remove them from the list. [Customizing and displaying report output](#) on page 679 provides information on selecting and removing parameters.

Backup status drill-down reports

The drill-down reports are composed of multiple NetWorker Backup Status basic reports, which are connected in a predetermined sequence. [Drill-down reports](#) on page 652 provides general information about drill-down reports. When a report has been chosen, the Configuration tab displays boxes with lists of the selected parameters for the top-level report. Thus, if the top layer of the drill-down report is a Daily Summary report, the configuration parameters are the same as they would be for the basic report, Daily Summary.

To exclude unwanted parameters from the report, remove them from the list. [Customizing and displaying report output](#) on page 679 provides information on selecting and removing parameters.

Inactive files

A NetWorker administrator can manage inactive files on a client or group and set the NetWorker software to automatically generate a list of inactive files in an environment. Inactive files are files that have not been accessed or modified other than being backed up regularly. The period of time a file has been inactive is called the Inactivity Threshold.

The inactivity files report is not supported on releases earlier than release 7.4 of the NetWorker servers. These versions are now unsupported.

Client support for this feature will be enabled only on Windows platforms.

The Inactive files report is a drill-down report that lists the inactive files from the latest scheduled backup. The report operates at both the client and group level.

The inactive files report can do the following:

- Generate a report on the percentage of inactive files backed up as part of a group.
- Set the threshold time periods per group so that the percentage of inactive files in that group does not exceed the threshold time period.
- Set alerts so that the NetWorker software sends an alert when the threshold set for a group is exceeded.
- Provide a report that details the percentage of inactive files backed up as part of a group.
- Report the percentage of inactive files per client.

The range limit specification given to configure File Inactivity Threshold and File Inactivity alert threshold attributes can be configured within the following ranges:

- File Inactivity Threshold attribute can be set between 0-365 days.
- File Inactivity Alert Threshold attribute can be set between 0-99.

Group File Details

The Group file Details report provides statistical information about inactive files that are included in a scheduled backup. Data will be provided for every requested NetWorker group at the time of the last backup. Chart mode is the default mode for the report. The data can also be viewed in tabular mode for more detailed information.

When generating the Group Details report, you can specify the following parameters:

- One or more NetWorker servers. Only servers that have the Gather Reporting Data attribute turned on will appear in the selection list.
- One or more NetWorker groups for the selected NetWorker servers.

Client File Details

The Client File Details report provides information about inactive files backed up for selected NetWorker clients. Data will be provided for every requested NetWorker

client at the time of the last backup. Chart mode is the default mode for the report. The data can also be viewed in tabular mode for more detailed information.

When generating the Client File Details report, you can specify the following parameters:

- One or more NetWorker servers. Only servers that have the Gather Reporting Data attribute turned on will appear in the selection list.
- One or more NetWorker groups for the selected NetWorker servers.
- One or more NetWorker clients for the selected NetWorker servers.

Data Domain statistics reports

The Data Domain reports, available from the Reports task pane in the **Console** window, provide Data Domain deduplication backup statistics for each selected NetWorker client.

The *NetWorker Data Domain Boost Integration Guide* provides more information.

NetWorker clone reports

The Clone reports, available from the Reports task pane in the **Console** window, allow you to view the history of automatic and scheduled clone operations that have been performed by NetWorker servers for any server version 7.6 Service Pack 2 and later. These versions are unsupported.

Four different types of clone reports can be generated:

- Server Summary
- Clone Details
- Save Set Details
- Clone Summary Over Time

Be aware that clone reports may not be up-to-date because clone records are gathered by the console server every 12 hours.

Types of NetWorker clone reports and configuration

The NetWorker clone report category includes basic and drill-down reports for each selected NetWorker server within the enterprise. The **Configuration** tab allows you to limit the scope of the report that was selected.

The parameters available for clone reports are described in this table. The specific parameters available depend on which clone report is selected.

Table 109 Clone report parameters

Parameter	Description	Options
NetWorker Server	Select one or more NetWorker servers.	Selected server names.
Client Name	Name of the NetWorker client whose save sets were cloned.	Selected client names.
Clone Name	Name of the scheduled clone resource that is used for cloning.	Selected clone resource.
Save Set	Cloned save set name.	Selected save set names.

Table 109 Clone report parameters (continued)

Parameter	Description	Options
Level	Backup level of the clone.	- Full - Incremental - Skip - Level 1–9 (Partial list of options)
Status	Completion status of the clone.	- Successful - Failed - No save sets found
Type	Type of clone operation.	- Scheduled - Manual
Start/End Time	Limits the report to a specified time range. The default range is one day for save set details reports. The date/time format available depends on the language locale of the operating system.	Start time of clone / End time of clone.

Clone basic reports

Within the Clone report category, choose any of the basic reports that are listed in the user interface. Once a report is chosen, the **Configuration** tab displays boxes with lists of the selected parameters for that report. To exclude unwanted parameters from the report, remove them from the list. [Customizing and displaying report output](#) on page 679 provides information on selecting and removing parameters.

Clone drill-down reports

The Clone Summary over Time drill-down report consists of the basic clone reports, which are connected in a predetermined sequence. [Drill-down reports](#) on page 652 provides general information about drill-down reports.

The configuration parameters for the drill-down report are the same as the parameters for the Server Summary basic clone report.

To generate the Clone Summary Over Time report, first specify the same parameters as those in the Server Summary clone report, which is the first report displayed in the sequence.

To drill-down to the clone detail level, perform one of the following, depending on your viewing mode:

- When in Table mode, double-click any individual row referencing the desired NetWorker server.
- When in Chart mode, click anywhere in the chart area of the desired NetWorker server.

The Clone Details report for the selected NetWorker server appears. Return to the Server Summary report to select another server to explore.

To drill-down to the Save Set Details level, perform one of the following, depending on the viewing mode:

- When in Table mode, double-click any individual row referencing the desired clone resource name.
- When in Chart mode, click anywhere in the chart area of the desired clone resource name.

The Save Set Details report for the selected clone resource appears. Return to the Clone Details report to select another client to explore.

Data Protection Policy reports

The Data Protection policy reports, available from the Reports task pane in the **Console** window, provides details and summaries for Data Protection Policies.

The *NetWorker VMware Integration Guide* provides more information.

Customizing and displaying report output

NMC provides you with configuration parameters for each type of report. Configuration parameters act as filters to limit criteria that are used to generate the information that is provided in a report. By default, each report sets these parameters to include all the information available in the report, the report does not filter any data.

About this task

When you accept the default configuration of the parameters results, NMC generates a report that includes statistics for all backup and clone actions that are initiated in a data protection policy resource within the last day, for all the servers in the enterprise. The statistics reported for each server would include all backup types and levels, and the time range would include all data available. Use the configuration parameters to define the data that is displayed by a report.

Note: An administrative user can restrict the user that have access to certain servers in the enterprise, which can limit the scope of the reports that the user can create and view.

Procedure

1. From the NMC GUI, click **Reports**.
2. Expand a report category folder, and then select an available report type.
The report open on the **Configuration** tab. The possible parameters for that report appear by default in the **Selected** boxes.
3. Define the report criteria:
 - To limit the scope of the report, click any of the parameters in the **Selected** box, then click **Remove** ().
 - To remove all the parameters from the **Selected** box, click **Remove All** ().
Removed parameters appear in the **Available** boxes.
 - To return a single parameter to the **Selected** box, select it from the **Available** box, and then click **Add** ().
 - To return all available parameters to the **Selected** box, click **Add All** ().

4. To display the report, select the **View Report** tab.

Note: If you receive the error `com.sybase.jdbc3.jdbc.SybDriver` when you generate a report, close the NMC GUI, clear the Java Cache on the NMC client, and then generate the report again. The *NetWorker Installation Guide* describes how to clear the Java Cache.

5. Most reports display initially in interactive mode and table format, to modify the report, right-click the **View Report** tab and select one of the following options:

Option	Description
Table	Display the data in Table view.
Chart	Display the data in Chart view
Document	Display the report in Document mode.
Interactive	Display the report in Interactive mode.
Portrait	Display the data in Portrait format.
Landscape	Display the data in Landscape format.

6. To print the report, right-click the **View Report** tab, and select **Print**.
7. To export the report, right-click the **View Report** tab, and select **Export**. In the **Save** dialog box, specify the file name and file location, and then click **Save**.

You can export the report to one of the following formats:

Option	Description
Postscript	For printing. Shows data totals.
PDF	For printing or viewing with a PDF viewer such as Adobe Acrobat. Shows data totals.
HTML	For viewing in a browser. Shows data totals.
CSV	For importing raw data into other programs, such as spreadsheets, that accept the comma separated values (CSV) format. Does not show data totals.

Start date and time formats

NMC includes Workflow Start Time and Workflow End Time parameters for Policy reports, and Start Time and End Time parameters for other reports, including legacy reports.

If a report includes a start date-and-time-range parameter, configure the time range in the following way:

- Specify the end date and time in the **To** box.
- Specify the start date and time in the **From** box.
- Use the arrow beside the time input field to display a calendar and clock selector, which includes adjustment arrows that enable you to set values.

All Policy reports and the Manual Save reports default to a one day time range, where one day represents a 24 hour period before the time on the NMC client host. The Legacy reports do not have a default time range and by default, the report displays the available data in the NMC database at time you generate the report.

Before modifying the time range, consider the following information:

- In US English locales, the default “From” hour is 12:00:00 (midnight/morning) on the “From” date, and the default “To” hour is 11:59:59 (night) on the “To” date. The US English locale is the only one that includes a box for an a.m. or p.m. value.
- In non-US English locales, the default “From” hour is 00:00:00 (midnight/morning) on the “From” date, and the default “To” hour is 23:59:59 (night) on the “To” date.

 **Note:** The Regional and Language Settings on the system determines whether the times appear in 12-hour or 24-hour formats.

Input formats

Date and time input formats in the NetWorker software vary. Some acceptable input formats for a collection of common locales are shown in this table.

Table 110 Date and time input formats for common locales

Language	Date formats	Time formats
US English	• EEEE, MMMM D, YYYY (Monday, March 8, 2009) • MMMM D, YYYY (March 8, 2009) • MMM D, YYYY (Mar 8, 2009) • M/D/YY (3/8/07)	• h:mm:ss a z (11:27:30 P.M. PST) • h:mm:ss a (11:27:30 P.M.) • h:mm a (11:27 A.M.)
UK English	• DD MMMM YYYY 08 March 2009 • DD-MMM-YYYY (08-Mar-2009) • DD/MM/YY (08/03/07)	• HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
French	• EEEE D MMMM YYYY (lundi 8 mars 2009) • D MMMM YYYY (8 mars 2009) • D MMM YYYY (8 mar. 2009) • DD/MM/YY (08/03/07)	• HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
German	• EEEE, D. MMMM YYYY (Montag, 8. März 2009) • D. MMMM YYYY (8. März 2009) • DD.MM.YYYY (08.03.2009DD) • MM.YY (08.03.07)	• HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
Japanese	• YYYY/MM/DD (2009/03/08) • YY/MM/DD (07/03/08)	• HH:mm:ss z (23:27:30 JST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
Simplified Chinese	• YYYY-M-D (2009-3-8)	• HH:mm:ss (23:27:30)

Table 110 Date and time input formats for common locales (continued)

Language	Date formats	Time formats
	YY-M-D (07-03-8)	

Note that in the previous table:

- Formats shown as single digits (M, D, h) may also be entered as double digits. For example, M could be either 7 or 07 for the seventh month.
- In the time-formats column:
 - The **a** character denotes a 12-hour format.
 - The absence of an **a** character denotes a 24-hour format.
 - The **z** character indicates time zone. If the **z** is present, then the output time will contain a time zone.

Relative times can also be entered in the From and To fields. A valid relative time consists of an number followed by a unit of time, for example, *2 months*. Time units can include Hour, Day, Week, Month, and Year.

Remember that these reports are run by using dates that have already occurred. Consequently, even the *To* date is always a past date. The relative time *4 months* would provide report data covering the past 4 months. A report specifying *from 9 months to 1 month* includes data from nine months ago up to one month ago.

Note: For Drive Utilization reports, the time range cannot exceed 8 days. That is, the date entered in the *To* field cannot exceed 8 days from the date entered in the *From* field. If typing a relative time in the *To* field, the value cannot exceed 8 days.

Background processing of reports

When you select the View Report tab, the NMC GUI queries the NMC server. This process happens in the background and may take a while. You can access other areas of the interface while the report data is being processed, the requested report appears when you return to the View tab.

NOTICE Do not request multiple reports simultaneously. Reports run sequentially in the background, and you can browse around in the user interface while a report is running. If you start a new report before an earlier report completes, NMC stops and deletes the earlier report. A report is either complete or deleted. The results are never partial.

Customizing and saving reports

A customized report is a changed copy of a canned report. Canned reports can be changed and then saved under different names. You can preserve the report configuration parameters that are most useful for the enterprise. For NetWorker reporting purposes, the terms customized report and saved report are synonymous.

A customized report can be rerun the same way at a later time, and even by another user. This saves time if the same report information must be generated repeatedly.

Customized reports offer these additional options, available from the right-click menu of a customized report:

- Delete—To delete the report.
- Rename—To rename the report.

- **Save**—To save the report.
- **Save As...**—To resave the report by using a different name.
- **Share**—To add sharing to the report or to remove sharing from the report.

Note: Only the original owner of a customized report or the Console Application Administrator can select these additional options. If the Console Application Administrator removes sharing, the report becomes private again to the original owner, the report's creator.

Since it is a copy, a customized report can be changed again and resaved, or even deleted. Reports can be saved either to preserve particular configurations (such as which servers are polled) or to save the view type (such as pie or bar chart).

Customized reports appear alphabetically in the report hierarchy below the canned report from which they were created. They are stored in the NMC database, which means that users can access them from any host that they use to log in to the NMC GUI and can use the report from a command prompt. [Command line reporting](#) on page 684 provides more information about running reports from the command line.

A customized report stores the following configuration information:

- All options from the report's **Configure** tab.
- Column display preferences for tables.
- Orientation (portrait or landscape).
- Current view type (table or chart). For charts, NMC also saves the current chart type (bar, pie, plot, or stacked bar) and the chart axis selection. [Interactive and document mode chart types](#) on page 649 provides more information about chart axis selection.

Naming reports

When naming a report to save, keep in mind that the set of usable characters is limited in the same way as for hostnames and usernames. Report names may not contain:

- Characters having an ASCII representation number less than ASCII 32 (such as carriage return, bell, newline, escape)
- Comma (,)
- Slash (/) or backslash (\)
- Double quote (") or single quote (')

Note: Report names are not case-sensitive. Also, canned reports cannot be deleted or customized, and then saved under the same name as a report that already exists under the same parent folder or directory.

Saved file ownership and deleted users

When a user saves a report by using the Save As command, that user becomes the owner of the new report. When a Console Application Administrator deletes from the system a user who owns reports, then the Console Application Administrator sees a dialog box that shows all of the reports owned by that user, and can choose either to delete the reports or reset the owner to a different user.

Sharing a report

By default, when you save a customized report, the report is private and appears only in the report hierarchy. The report owner or an NMC user with the Console

Application Administrator role can share the report with other NMC users. Perform the following steps to share a customized report.

About this task

Any user viewing a sharable report may perform these operations on the report:

- Change any runtime parameter of the report (such as configuration or view type).
- Run the report, but not save changes to the report.
- Copy the report by using the Save As command. The user becomes the owner of the new report, and by default, the report is not shared.
- Choose the Hide Other Users' Reports option to toggle the view of reports between only those reports owned by the user (both private and shared), and all shared custom reports.

Perform the following steps to share a report.

Procedure

1. From the NMC GUI window, click **Reports**.
2. Expand the report folder that contains the customized report that you want to share.
3. Right-click the customized report, then select **Share**.

The report is now shared, and is represented in the report hierarchy by a

shared-report icon  or .

Results

Once you enable a report for sharing, all users can see the report in the report folder hierarchy.

 **Note:** The Share option is a toggle. To disable sharing, right-click the shared report and select **Share**.

Command line reporting

Command line reporting offers the following features:

- Allows reports to be run offline, either as needed or by using scheduling software that makes reports available at predetermined times.
- Uses both canned and customized reports, which can be exported in various formats.
- Provides a more advanced feature that requires a fair amount of knowledge about running and scripting from the command prompt of the Console server. This feature should be reserved for advanced users.

 **Note:** Command line reports may only be printed or run to generate exported output. They cannot be saved or shared. Drill-down reports cannot be run from the command line.

The command line reporting program

The command line reporting program is `gstclreport`. It uses the JRE to run. Command line reports must be run on the NMC Console server host.

The options are typical command line options in the form of a hyphen (-) followed by one or two letters and an argument, if applicable. The UNIX man page and the

NetWorker Command Reference Guide provides detailed information about the `gstclreport` command.

System performance

Each time the `gstclreport` command is run, it starts a separate JVM, which can use many system resources. The `gstclreport` command runs a database query and generates report output by using the results. Since this uses both CPU and memory resources on the host computer, it could affect performance of NetWorker software and of the host. Consequently, depending on the system used, it is probably not wise to run more than a few instances of the `gstclreport` command at the same time.

Security

The `gstclreport` command must contact the Console server in order to run a report. The command requires a valid username and password. A user either uses the `-P` option to type the password, or the command checks standard input to see whether the password is there. If a password is not supplied, the program prompts for a password.

On UNIX systems, use of the `-P` option is a security concern, because a user may type the `ps` command and see the commands that were used to start any program that is running.

To solve this problem, use scheduling software that can conceal password input. Alternatively, ensure that the scheduling system sends the password as standard input. For example:

```
echo password | gstclreport
```

A `cron` command can be used to schedule the report, or the command could be placed in a secure script file that is invoked by the `cron` command.

Java runtime environment

Support of command line reporting requires JRE version 1.6 or later to run the `gstclreport` command. The JRE must be installed before installing NetWorker software.

You must also add an environment variable that is named `JAVA_HOME` to the NetWorker server host. Open either the `gstclreport.bat` or `gstclreport.sh` file and follow the instructions at the top of the file to set up the correct environment for command line reporting.

Reporting policy status and backup job status

When you perform a backup, clone or archive actions, NetWorker records the status of the action and job activities. There are three ways to report job activities:

- In the **Monitoring** window for the NetWorker server in NMC. [Monitoring NetWorker Server activities in the Administration window](#) on page 55 describes how to view the action completion status in the Monitoring window.
- Through predefined notifications, that you can define at the policy, workgroup, or action level. [Policy completion and failure notifications](#) on page 686 provides more information.

- By querying the job status. [Querying the job status](#) on page 687 provides more information.

Policy completion and failure notifications

You can configure NetWorker to generate a notification for each action that fails in a policy, or an email that summarizes the status of a policy in which all actions succeed. By default, a UNIX NetWorker server sends an email that provides information about the status of completed actions to the root account of the NetWorker server. A Windows NetWorker server writes information to the `policy_notifications.log` file located in the `NetWorker_install_dir\nsr\logs` directory. [Policy Notifications](#) on page 263 describes how to configure Policy notifications.

Format of the Policy Completion and Policy Failure notifications

Policy notifications are divided into two sections that describe the job activities for a Policy.

This information also appear in the `policy_notifications.log` file.

- **Summary notification report**—Provides a summary of the status of the workflow and actions that are associated with a Policy resource.

For example:

```
---Summary notification report---
Policy name:Server Protection
Workflow name:Server backup, Workflow status:failed,
Workflow start time:Thu Nov 20 21:00:01 GMT-0500 2014,
Duration:
Action name:Server db backup, Action status:failed, Action
start time:Thu Nov 20 21:00:01 GMT-0500 2014, Duration:0
hours 0 minutes 14 seconds
```

- **Action report**—Provides summary and status information about each action that is associated with the Policy resource.

For example:

```
--- Traditional Backup Action report ---
Policy name:Server Protection
Workflow name:NMC server backup
Action name:NMC server backup
Action status:failed
Action start time:Thu Nov 20 21:00:02 GMT-0500 2014
Action duration:0 hours 4 minutes 13 seconds
Total 1 client(s), 0 Succeeded with warning(s), 0 Succeeded,
1 Failed.
---Successful backups---
none
---Failed backups---
bu-iddnwserver2.iddlab.local:C:\Program Files\EMC NetWorker
\Management\nmcd_b_stage, level=full, size 0.000 MB ,
Duration 0 hours 1 minutes 3 seconds, (null) files
```

Customizing the save sets status in the policy notifications

NetWorker reports the status of a save session that completes with warning based on the value defined in **Success threshold** attribute for an action.

Events that might trigger a warning when they occur during a backup include the following conditions:

- The file size increases or decreases
- The mtime of the file changes

To define the success threshold for a save session, select one the following values in the **Specify the Backup Options** screen of the Policy Action Wizard:

- **Warning**—Save sets that complete with warnings are reported as success with warnings.
- **Success**—Save sets that complete with warnings are reported as failed. This is the default value. The number of times NetWorker retries a failed save set is determined by the value defined in the **Retries** attribute, which you specify in the **Specify the Advanced Options** screen of the Policy Action Wizard.

The **Success threshold** attribute also applies to the save sets displayed in the **Monitoring** window.

Querying the job status

When a workflow or action resource runs within a Policy resource, NetWorker stores job information in policy log files and the jobs database (jobsdb) on the NetWorker server host.

The NetWorker software provides two command line programs to query job information in the jobsdb:

- `jobquery` —To locate and retrieve detailed information on a job, including the child jobs of an action.
- `nsrpolicy monitor`—To retrieve summary information about a job.

The man pages or the *NetWorker Command Reference Guide* provides more information on the `jobquery` and `nsrpolicy monitor` commands.

Workflow and action job records

NetWorker represents each Workflow and Action resource with a job record in the jobsdb. Some actions create child jobs to perform the tasks that are associated with the action. NetWorker creates a unique job record for each child job and stores session information about each child job. NetWorker associates each piece of information about a job with an attribute. Each job record is composed of a group of attributes, including the job id attribute. The job id attribute is a numeric value that uniquely identifies the job record. NetWorker groups attributes together by type. A type contains unique attributes and attributes that are common to all job types.

Job record types

To display information about a job record, build queries that are based on the job type. The jobsdb contains the following policy-related job record types:

- **Backup action job** —Job that is created for a traditional or snapshot backup action. A traditional backup action job starts child jobs, for example, the save job and the savefs job to perform action tasks that NetWorker requires to complete an action.
- **Bootstrap save job** — Job that is created for the server database backup action. The bootstrap save job starts child jobs, for example, an index save.
- **Check connectivity job action-** Job that is created for the check connectivity action.
- **Clone job** — Job that is created for a clone action. A clone action job starts child jobs to perform action tasks that complete an action.

- Discover job action — Job that is created for a NAS discover action.
- Generate index action job — Job that is created for a generate index action.
- Probe action job — Job that is created for a probe action.
- Utility job — Is an action that performs a maintenance task, for example, the expire action, the vba-checkpoint-discover action, and the vba-checkpoint-backup action. A job can start a child utility job to perform tasks that the parent job requires to complete an action. For example, the server backup action job starts a child job that runs the `mminfo -B` command.
- Vbasave job — Child job that is created by a VMware backup action job.
- Workflow job — Job that is created for a workflow.

NetWorker clears the information about a job from the `jobsdb` and deletes the associated log files at the interval that is defined by the **Jobsdb retention in hours** attribute in the properties of the NetWorker Server resource. In NetWorker 9.0.1, the default `jobsdb` retention is 72 hours.

Using jobquery

The `jobquery` program provides a CLI similar to the `nsradmin` program. The `jobquery` program contacts the `nsrjobd` process to query job information that is stored in the `jobsdb`. A query is defined by an attribute list that is made up of one or more attribute names with or without values.

About this task

In the query, the attribute name (for example, 'type') is preceded by a period ('.'), and optionally followed by a colon (':') and a comma-separated list of values (for example, "host: mars"; "job state: STARTED, ACTIVE, SESSION ACTIVE"). When a query consists of more than one attribute names, attributes are separated by a semi-colon (;). When an attribute name is specified without values, any resource descriptor that contains this attribute is a match. If an attribute name is followed by one or more values, a resource whose value list matches at least one of the values for the specified attribute satisfies the criteria.

To launch the `jobquery` interface, type:

```
jobquery -s NetWorker_server
```

Where *NetWorker_server* is the hostname of the NetWorker Server. Use the `-s` option when you run the `jobquery` command from a NetWorker host that is not the NetWorker Server.

 **Note:** When you do not use the `-s` option, `jobquery` tries to connect to `nsrjobd` process on the local host. If the `nsrjobd` process is not running on the specified server or the local host, an error is returned.

The `jobquery -s<server>` command connects to the specified NetWorker server and returns jobquery prompt. The data in the job database is queried with the following commands:

- `types` — a command that lists all job types currently known by `nsrjobd` that does not take any argument (for example, `types` return a list indicating Known types: save job, savegroup job, and so on).
- `.` — a command that sets the query criteria and is followed by one or more attribute names, or lists current query criteria when not followed by any attribute. Query criteria may contain several attributes, including job type, host, and job state, with each attribute separated by a semi-colon and each value separated by a comma, as in the following example:

```
jobquery> . type: savegroup job; host: mars; job state: ACTIVE,
COMPLETED
```

This example would return information on all savegroup jobs from the host mars that are either in progress or in completed state.

- **show** — restricts the list of attributes that are returned for each resource descriptor that matches the query. For the above example, specifying the following:

```
show name; job id; completion status; completion severity
```

returns the names, job ids, completion status, and completion severity for all matched completed and active savegroups.

- **print** — runs the query and displays the results. If show list is in effect, each resource descriptor in the result list is restricted to desired attributes.
- **all** — returns all resource descriptors in the jobs database. If show list is in effect, result is restricted to desired attributes.
- **help** — displays help text.
- **quit** — exits jobquery.

Running `jobquery -s NetWorker_server -i input_file` reads input from the file for non-interactive usage. The man pages or the *NetWorker Command Reference Guide* provides detailed information about the jobquery program.

Querying the jobsdb for workflow job records

Each time that you start a workflow, NetWorker creates a single workflow job in the jobsdb. Run the `jobquery` command to display information about the workflow job.

To query the jobsdb for information about workflows in a policy, type the following command in the jobquery interface:

```
. type: workflow job; data protection policy name: policy_name;
workflow name: workflow_name
```

where *policy_name* is the name of the policy that contains the workflow and *workflow_name* is the name of the workflow.

 **Note:** The *policy_name* and *workflow_name* values are case sensitive.

For example, to query the jobsdb for a workflow named SQL Clients in a policy named Backup, type the following commands at the `jobquery` prompt:

```
jobquery>. type: workflow job; data protection policy name: Backup;
workflow name: SQL Clients
jobquery>print
```

Output similar to the following appears:

```
type: workflow job;
activity progress: 1/0/1;
actual exit code: 1;
adhoc job: False;
authtype: ;
automatic: False;
Checkpoint restart ID: ;
Checkpoint restart sequence: ;
command: ;
completion severity: 50;
```

```

completion status: failed;
data protection policy name: Backup;
dependent job id: 0;
end time: 1435107619;
exit code known: True;
host: bu-iddnwserver.iddlab.local;
input flag: False;
job id: 832031;
job log file: \
"C:\Program Files\EMC NetWorker\nsr\logs\policy\Backup\
workflow_SQL clients_832031.raw";
job output: \
"133550 1435107602 1 0 0 3376 4996 0 bu-
iddnwserver.iddlab.local nsrworkflow NSR notice 31 Starting %s
%s' workflow '%s'. 3 11 24 127405:Protection Policy \
0 6 Backup 0 11 SQL clients
123316 1435107602 1 0 0 3376 4996 0 bu-iddnwserver.iddlab.local
nsrworkflow NSR notice 46 Starting action '%s/%s/%s' with
command: '%s'. 4 0 6 Backup 0 11 S\
QL clients 0 6 backup 0 32 savegrp -Z backup:traditional -v
123321 1435107602 1 0 0 3376 4996 0 bu-iddnwserver.iddlab.local
nsrworkflow NSR notice 39 Action '%s/%s/%s's log will be in
%s'. 4 0 6 Backup 0 11 SQL clie\
nts 0 6 backup 23 83 C:\Program Files\EMC NetWorker\nsr\
logs\policy\Backup\SQL clients\backup_832032.raw
123325 1435107619 1 0 0 3376 4996 0 bu-iddnwserver.iddlab.local
nsrworkflow NSR notice 21 Action '%s/%s/%s' %s. 4 0 6 Backup 0
11 SQL clients 0 6 backup 0 6\
failed
133555 1435107619 1 0 0 3376 4996 0 bu-iddnwserver.iddlab.local
nsrworkflow NSR notice 24 Workflow '%s/%s' failed. 2 0 6 Backup
0 11 SQL clients";
job state: COMPLETED;
name: Backup;
ndmp flag: False;
NW Client name/id: ;
override parameters: ;
parent job id: 0;
policy definition changetime: 1434655597016534;
previous job id: 0;
protection groups: SQL clients;
Reason job was terminated: ;
redirect stdio: False;
remote password: ;
remote user: ;
restricted data zone: ;
root parent job id: 0;
savegrp spawned: False;
sibling job id: ;
SSID: ;
start time: 1435107602;
type attributes: ;
type classes: ;
type help: ;
type name: ;
type references: ;

```

```
type table: ;
userid: ;
workflow name: SQL clients;
```

The following table summarizes some of the attributes that appear in workflow job types.

Table 111 Workflow-specific job record attributes

Attribute	Description
Job id	A unique number value that identifies the job.
Parent job id	The job id of the job that started this job. A job may not have a parent job.
Job state	The status of the job. Status values include: CREATED, QUEUED, STARTED, ACTIVE, SESSION ACTIVE, CANCELLED, and COMPLETED.
Job log file	The location and name of the log file that contains detailed information about the job activities.
Job output	The information that is contained in the <i>job log file</i> . Note: Truncation of the content might occur when the file is large, which results in only displaying the last 2 KB of information.
Start time	The time the job started, in seconds since Jan 1, 1970.
End time	The time the job ended, in seconds since Jan 1, 1970.
Completion status	The completion status set by the job. Status values include never started, did not run, succeeded, failed, abandoned, canceled, and communication lost between job and nsrjobd.
Completion severity	The severity level of any error that caused the job to end. Severity levels include: EMERGENCY, ALERT, CRITICAL, SEVERE, ERROR, INTERVENTION, WARNING, NOTICE, and INFORMATION.
Data protection policy name	The name of the policy that contains the workflow.
Workflow name	The name of the Workflow resource.
Override parameters	A list of parameters that were configured in the Workflow resource, when the workflow started. Use override parameters to override the value that is defined for an equivalent action property.

Table 111 Workflow-specific job record attributes (continued)

Attribute	Description
Protection groups	The protection groups that are assigned to the workflow.
Restricted datazone	The datazone to which the resource is assigned.
Policy definition changetime	The last change time of the policy that contains the workflow.
Previous jobid	The job id of the instance of a restarted workflow.

Querying the jobdb for action records

Each time that an action starts, NetWorker creates a job record for the action in the jobdb. Some actions create child actions, for example a backup action creates a save job and a savefs job. Each child action has a unique job record. Use the `jobquery` command to display information about an action job.

To query the jobdb for information about an action job, type the following command in the `jobquery` interface:

```
. type:action_name
```

where *action_name* is the name of the action.

For example, to query the jobdb for a bootstrap save job, type the following commands at the `jobquery` prompt:

```
jobquery>. type: bootstrap save job
jobquery>print
```

Output similar to the following appears:

```
type: bootstrap save job;
actual exit code: 0;
adhoc job: False;
authtype: ;
automatic: False;
canceled input count: 0;
canceled input work items: ;
Checkpoint restart ID: ;
Checkpoint restart sequence: ;
command: nsrdbsave -l 1;
completed output count: 0;
completed output work items: ;
completion severity: 50;
completion status: failed;
data protection policy name: Server Protection;
data size: ;
dependent job id: 0;
end time: 1434895738;
exit code known: True;
failed input count: 2;
failed input work items: bu-iddsql.corp.com,
```



```

bu-iddnwserver.iddlab.local;
file count: ;
filtered input count: 0;
filtered input work items: ;
hard runtime limit: 0;
host: bu-iddnwserver.iddlab.local;
input flag: True;
input job id: ;
job id: 800020;
job log file: \
"C:\Program Files\EMC NetWorker\nsr\logs\policy\Server
Protection\Server\
r backup\Server db backup_800020.raw";
job output: \
"suppressed 799 bytes of output.
140403 1434808808 1 5 0 2284 2280 0 bu-iddnwserver.iddlab.local
nsrdbsave NSR notice 55 Started '%s' job with jobid [%u].
Backup command:\n %s. 3 0 12 inde\
x backup 5 6 800022 0 367 save -q -e \"1 Months\" -b Default -J
bu-iddnwserver.iddlab.local -a \"*policy name=Server Protection
\" -a \"*policy workflow name\
=Server backup\" -a \"*policy action name=Server db backup\" -g
\"Server Protection\" -l full -S -f - -LL -W 78 -N
index:2668af1d-00000004-54528c1a-5452a19b\
-00155000-7396bc56 -x \"C:\\\\Program Files\\\\EMC NetWorker\\\\
\nsr\\\\index\\bu-iddsql.corp.com\"
140402 1434895685 1 5 0 2284 2280 0 bu-iddnwserver.iddlab.local
nsrdbsave NSR notice 35 Completed '%s' job with jobid [%u]. 2 0
12 index backup 5 6 800022
140402 1434895738 1 5 0 2284 2280 0 bu-iddnwserver.iddlab.local
nsrdbsave NSR notice 35 Completed '%s' job with jobid [%u]. 2 0
12 index backup 5 6 800021
112777 1434895738 5 3 13 2200 1572 0 bu-
iddnwserver.iddlab.local nsrd RAP critical 119 Permission
denied, application provided an expired session ticket; us\
er '%s' on '%s', cur time %s, expiration time %s . 4 13 6
SYSTEM 12 27 bu-iddnwserver.iddlab.local 35 10 1434895738 35 10
1434812401
138211 1434895738 3 0 0 2284 2280 0 bu-iddnwserver.iddlab.local
nsrdbsave NSR error 33 Verify that NetWorker is running. 0
140403 1434895738 1 5 0 2284 2280 0 bu-iddnwserver.iddlab.local
nsrdbsave NSR notice 55 Started '%s' job with jobid [%u].
Backup command:\n %s. 3 0 10 mmin\
fo_job 5 6 800040 0 9 mminfo -B 140402 1434895738 1 5 0 2284
2280 0 bu-iddnwserver.iddlab.local nsrdbsave NSR \
notice 35 Completed '%s' job with jobid [%u]. 2 0 10 mminfo_job
5 6 800040 140407 1434895738 1 5 0 2284 2280 0 bu-
iddnwserver.iddlab.local nsrdbsave NSR \
notice 48 See the file '%s' for detail output of each job. 1 0
107 C:\\Program Files\\EMC NetWorker\nsr\logs\policy\Server
Protection\Server backup\Se\
";er db backup_800020_logs
job state: COMPLETED;
level: ;
name: nsrdbsave;
ndmp flag: False;

```

```

number of files: ;
NW Client name/id: ;
override parameters: ;
parallelism: 0;
parent job id: 800019;
policy action name: Server db backup;
previous job id: 0;
Reason job was terminated: ;
redirect stdio: True;
remote password: ;
remote user: SYSTEM;
restricted data zone: ;
root parent job id: 800019;
running input count: 0;
running input work items: ;
savegrp spawned: False;
saveset id: ;
sibling job id: ;
size: ;
soft runtime limit: 0;
SSID: ;
start time: 1434808802;
successful input count: 0;
successful input work items: ;
type attributes: ;
type classes: ;
type help: ;
type name: ;
type references: ;
type table: ;
userid: ;
waiting input count: 0;
waiting input work items: ;
workflow name: Server backup;

```

The following table summarizes some of the attributes that appear in action job types.

Table 112 Action job record attributes

Attribute	Description
Job id	A unique number value that identifies the job.
Parent job id	The job id of the job that started this job. A job may not have a parent job.
Job state	The status of the job. Status values include: CREATED, QUEUED, STARTED, ACTIVE, SESSION ACTIVE, CANCELLED, and COMPLETED.
Job log file	The location and name of the log file that contains detailed information about the job activities.
Job output	The information contained in the <i>job log file</i> .  Note: Truncation of the content might occur when the file is large.

Table 112 Action job record attributes (continued)

Attribute	Description
Start time	The time the job started, in seconds since Jan 1, 1970.
End time	The time the job ended, in seconds since Jan 1, 1970.
Completion status	The completion status set by the job. Status values include never started, did not run, succeeded, failed, abandoned, canceled, and communication lost between job and nsrjobd.
Completion severity	The severity level of any error that caused the job to end. Severity levels include: EMERGENCY, ALERT, CRITICAL, SEVERE, ERROR, INTERVENTION, WARNING, NOTICE, and INFORMATION.
Data protection policy name	The name of the policy that contains the action.
Workflow name	The name of the Workflow resource that contains the action.
Policy action name	The name of the action.
Input job id	The job id of the action that is controlling this action.
Waiting input work items	For the first or head action in a workflow, this is a list of work items for the protection group that is assigned to the workflow that contains the head action. For subsequent actions, this list displays the value in the completed output work items attribute, for the action that precedes this action. When an action starts a work item, the work item value moves from the waiting input work items attribute to the <i>running input work items</i> attribute.
Waiting input count	The number of work items in the <i>waiting input work items</i> attribute.
Filtered input work items	Contains work items that an action has filtered out of the <i>waiting input work items</i> attribute.
Filtered input count	The number of work items in the <i>filtered input work items</i> attribute.
Running input work items	A list of in progress work items that were previously in the <i>waiting input work items</i> attribute. This list does not display in progress work items that were previously in the <i>filtered input work items</i> attribute.
Running input count	The number of work items in the <i>running input work items</i> attribute.

Table 112 Action job record attributes (continued)

Attribute	Description
Successful input work items	A list of input work items that have completed successfully. When an input work item completes successfully, the value moves from the <i>running input work items</i> to the <i>successful input work items</i> attribute.
Successful input count	The number of work items in the <i>successful input work items</i> attribute.
Failed input work items	A list of input work items that have not completed successfully. When an input work item does not complete successfully, the value moves from the <i>running input work items</i> to the <i>failed input work items</i> attribute.
Failed input count	The number of work items in the <i>failed input work items</i> attribute.
Canceled input work items	A list of input work items that were canceled and did not complete. When an input work item is canceled, the value moves from the <i>running input work items</i> to the <i>cancelled input work items</i> attribute.
Canceled input count	The number of work items in the <i>cancelled input work items</i> attribute.
Completed output work items	The list of work items that are produced by this action.
Completed output count	The number of work items in the <i>completed output work items</i> attribute.

Querying session information

Some actions, for example, backup, clone, and expiration actions create child actions to perform the tasks that are required to complete an action. NetWorker creates a session information record for each task that the child job starts. Use the `jobquery` command to view session information for the child action task. To view session information about an action task, the job id of the child job that created the save set is required.

To view session information for a child job, perform the following steps:

1. Review the logs directory to determine the job id of the workflow. [Policy log files](#) provides more information about policy-related log files.
2. Query the jobsdb for the workflow that contains the action.
3. In the output, search for the parent action that started the child job and record the job id of the parent.
4. In the output, search for the child action that contains the job id of the *parent job id* attribute, and record the value in the *job id* attribute of the child action.
5. Query the jobsdb by using the child job id to display the session information that relates to the child task.

Example 12 Viewing session information

In this example, the server backup workflow failed for a host at 7:54 PM on June 26. We want to review session information about the action tasks started by the server database backup and expiration actions.

1. In the C:\Program Files\EMC NetWorker\nsr\logs\policy\Server Protection folder on Windows or the /nsr/logs/policy/Server Protection directory on LINUX, the workflow_Server backup_832001.raw file appears for the workflow. The job id of the server backup workflow is 832049.
2. From a command prompt, start the jobquery program.
3. Display information about the actions started by the server backup workflow .

```
jobquery>print job id: 832049
```

In this example, two action jobs created child jobs. The server database backup action created a child process for the nsrdbsave command and the expiration action created a utility job. Output for the bootstrap save job appears, as follows. The job id for the failed bootstrap backup action is 832050. The following example displays some of the attributes that appear in the print output.

```
type: bootstrap save job;

canceled input count: 0;
canceled input work items: ;
command: nsrdbsave -l 1;
completed output count: 3;
completed output work items: 4253813558/1435241270,
4237036342/1435241270,
4220259140/1435241284;
completion severity: 10;
completion status: succeeded;
data protection policy name: Server Protection;
failed input count: 0;
failed input work items: ;
file count: ;
filtered input count: 0;
filtered input work items: ;
input flag: True;
input job id: ;
job id: 832050;
job log file: \
C:\\Program Files\\EMC NetWorker\\nsr\\logs\\policy\\Server
Protection\\Serve\
backup\\Server db backup_832050.raw";
job output:
job state: COMPLETED;
level: ;
name: nsrdbsave;
override parameters: ;
parent job id: 832049;
policy action name: Server db backup;
```

Example 12 Viewing session information (continued)

```

previous job id: 0;
root parent job id: 832049;
running input count: 0;
running input work items: ;
successful input count: 2;
successful input work items: bu-iddnwserver.iddlab.local,
bu-iddsql.corp.com;
waiting input count: 0;
waiting input work items: ;
workflow name: Server backup;

type: utility job;
canceled input count: 0;
canceled input work items: ;
command: nsrim -MXq;
completed output count: 0;
completed output work items: ;
completion severity: 10;
completion status: succeeded;
data protection policy name: Server Protection;
failed input count: 0;
failed input work items: ;
filtered input count: 0;
filtered input work items: ;
input flag: True;
input job id: 832050;
job id: 832057;
job log file: \
C:\\Program Files\\EMC NetWorker\\nsr\\logs\\policy\\Server
Protection\\Serve\\
backup\\Expiration_832057.raw";
job output: \
88411 1435241295 1 5 0 4996 1560 0 bu-
iddnwserver.iddlab.local nsrim NSR noti\
e 28 Checking for invalid volumes 0
6069 1435241295 1 5 0 4996 1560 0 bu-
iddnwserver.iddlab.local nsrim NSR notic\
21 Processing %d clients 1 1 1 3
6067 1435241295 1 5 0 4996 1560 0 bu-
iddnwserver.iddlab.local nsrim NSR notic\
37 Crosschecking indexes for %d clients. 1 1 1 1
6068 1435241297 0 0 0 4996 1560 0 bu-
iddnwserver.iddlab.local nsrim NSR info \
0 Managing %d volumes. 1 1 1 4
6073 1435241298 0 0 0 4996 1560 0 bu-
iddnwserver.iddlab.local nsrim NSR info \
; Compressing media database. 0
job state: COMPLETED;
name: nsrim;
override parameters: ;
parallelism: 0;
parent job id: 832049;

```

Example 12 Viewing session information (continued)

```

policy action name: Expiration;
previous job id: 0;
root parent job id: 832049;
running input count: 0;
running input work items: ;
successful input count: 0;
successful input work items: ;
waiting input count: 0;
waiting input work items: ;
workflow name: Server backup;

```

4. Display information about the failed bootstrap save job by specifying the job id, obtained from the bootstrap save job output .

```
jobquery>print parent job id: 832050
```

The `jobquery` program displays detailed information about the save jobs that are created by the job. The following example displays some of the attributes that appear in the print output.

```

type: save job;
actual exit code: 0;
adhoc job: False;
authtype: ;
automatic: False;
backup_device: ;
Checkpoint restart ID: ;
Checkpoint restart sequence: ;
command: \
"save -q -e \"1 Months\" -b Default -J bu-
iddnwserver.iddlab.local -a \"*polic\
y name=Server Protection\" -a \"*policy workflow name=Server
backup\" -a \"*po\
licy action name=Server db backup\" -g \"Server Protection\"
-l full -LL -LL -\
S -f - -W 78 -N bootstrap \"C:\\\\Program Files\\\\EMC
NetWorker\\\\nsr\\\\res\
\" \"C:\\\\Program Files\\\\EMC NetWorker\\\\nsr\\\\mm\"
\"C:\\\\Program Files\
\\\\EMC NetWorker\\\\nsr\\\\authc-server\\\\tomcat\\\\data
\"";
completed savetime: 1435241284;
completion severity: 10;
completion status: succeeded;
data class: ;
Data set size: ;
data size: 255;
dedupe sent bytes: ;
dependent job id: 0;
end time: 1435241293;
estimated bytes: ;
exit code known: True;
file count: 185;
Files totals: 156, 114497;

```

Example 12 Viewing session information (continued)

```

first_clone_id_for_ssid: ;
group name: ;
host: bu-iddnwserver.iddlab.local;
Inactive files: 0, 0, 0;
input flag: False;
job id: 832055;
job log file: \
"C:\Program Files\EMC NetWorker\nsr\logs\policy\Server
Protection\Server\
r backup\Server db backup_832050_logs\832055";
job output: ;
job state: COMPLETED;
level: ;
mmdb-avamar-backup-time: ;
mmdb-avamar-client-id: ;
mmdb-avamar-server: ;
name: bootstrap backup;
ndmp flag: False;
New data on De-Dup Node: ;
NW Client name/id: ;
parent job id: 832050;
policy action name: ;
policy name: ;
policy_action_name: ;
policy_name: ;
previous job id: 0;
primary_clone_id: ;
processed bytes: ;
proxy agent name: ;
proxy error code: ;
proxy_hostname: ;
Reason job was terminated: ;
redirect stdout: True;
remote password: ;
remote user: SYSTEM;
restricted data zone: ;
root parent job id: 832049;
savegrp spawned: False;
saveset name: ;
saveset time: 1435241284;
saveset_id: ;
sibling job id: ;
snap session id: ;
SSID: ;
start time: 1435241283;
type attributes: ;
type classes: ;
type help: ;
type name: ;
type references: ;
type table: ;
userid: ;

```


Example 12 Viewing session information (continued)

```
vba_backup: ;
vba_name: ;
vcenter_hostname: ;
vm guest os: ;
vm_name: ;
vm_uuid: ;
```

5. Display session information for the save job by specifying the job id.

```
jobquery>print jobid from session info: 832055
```

The `jobquery` program displays detailed session information about the save job. For example, output similar to the following appears:

```
type: session info;
client name: bu-iddnserver.iddlab.local;
completed: 1;
compression ratio: 0;
current pool: Default;
current read/write total: 254;
device family: disk;
Device path: afd;
device type: adv_file;
extended information: ;
Jobid from session info: 832055;
number of volumes used: 0;
recover file count: 0;
recover file total: 0;
restricted data zone: ;
savegroup name: Server Protection;
saveset id: \
7d52bfb9-00000006-fb8c0b44-558c0b44-00065000-7396bc56;
saveset name: bootstrap;
Session end time: 1435241299;
session id: 18269;
Session mode: 0;
Session start time: 1435241284;
total amount to be read/written: 0;
total volumes needed: 0;
transfer rate: 0;
type attributes: ;
type classes: ;
type help: ;
type name: ;
type references: ;
type table: ;
volume name: bu-iddnserver.iddlab.local.002;
```

Using nsrpolicy monitor

Use the `nsrpolicy monitor` command to query the jobsdb for details and status information about an active or inactive job started by a Data Protection Policy resource.

The `nsrpolicy monitor` command allows you to view information about the last active or inactive job that is associated with a Policy resource in a tabular or non-tabular output. You can display output for all Data Protection Policy resources in a policy, or limit the output by client name, workflow name, or protection group name.

```
nsrpolicy monitor -p policy_name -w workflow_name -c client_name -g
group_name -d -n -j job_id -s networker_server -D debug_level
where:
```

- `-p policy_name`—Specifies the name of the Policy resource. You cannot use this option with `-g group_name`.
- `-w workflow_name`—Specifies the name of the Workflow resource. Requires the `-p policy_name` option or the `-j job_id` option.
- `-c client_name`—Specifies the name of the Client resource. Requires the `-g group_name` option.
- `-g group_name`—Specifies the name of the Protection Group. You cannot use this option when you use the `-p policy_name` option.
- `-d`—Displays detailed information about the job.
- `-n`—Displays the output in non-tabular view.
- `-j job_id`—Displays detailed information about a specific job, which is identified by the jobid. You cannot use this option when you use the `-p policy_name` option.

Displaying job details for a Workflow resource

To retrieve the details about the last active or inactive jobs in a Workflow resource, type the following command:

```
nsrpolicy monitor -p policy_name [-w workflow_name]
```

For example, to provide information about a workflow that is called Default in the Backup Policy, type the following command:

```
nsrpolicy monitor -p Backup -w Default
```

Table 113 Job details for a Workflow

Policy	Workflow	Action	Job Name	Job id	Parent Job id	Job Type	Job Status	Completi on Status	Start Time	Duration
Backup	Default	Backup		32524		Workflow	COMPLE TED	succeede d	5/26/15 16:59:43	00:01:22
Backup	Default	Backup	savegrp	32525	32524	Backup act	COMPLE TED	succeede d	5/26/15 16:59:43	00:01:21

For example, to provide detailed information about the last active or inactive jobs in a workflow that is called Default in the Backup Policy, type:

```
nsrpolicy monitor -p Backup -w Default -d
```

Table 114 Job details for a Workflow continued

Policy	Workflow	Action	Job Name	Job id	Parent Job id	Job Type	Job Status	Completion Status	Start Time	Duration
Backup	Default	Backup		32524		Workflow	COMPLETED	succeeded	5/26/15 16:59:43	00:01:22
Backup	Default	Backup	savegrp	32525	32524	Backup act	COMPLETED	succeeded	5/26/15 16:59:43	00:01:21
			pseudo_save	32527	32525	save job	COMPLETED	succeeded	5/26/15 16:59:50	00:01:14
			C:\Software	32528	32527	save job	COMPLETED	succeeded	5/26/15 17:00:47	00:00:14
			bu-iddnwse	32526	32525	savefs job	COMPLETED	succeeded	5/26/15 16:59:43	00:00:01

For example, to display detailed information about the last active or inactive job in a Workflow resource, in a non-tabular format, type:

```
nsrpolicy monitor -p Backup -w Default -d -n
```

```
Workflow status:
data protection policy name:Backup
workflow name:Default
name:Backup
job id:32524
type:workflow job
job state:COMPLETED
completion status:succeeded
start time: 5/26/15 16:59:43
duration: 00:01:22
Action 1 status:
data protection policy name:Backup
workflow name:Default
policy action name:backup
name:savegrp
job id:32525
parent job id:32524
type:backup action job
job state:COMPLETED
completion status:succeeded
start time: 5/26/15 16:59:43
duration: 00:01:21
```

Displaying job details for a client in a group

To retrieve the information about the last job for a client in a group, type the following command:

```
nsrpolicy monitor -c client_name -g group_name
```

For example:

```
nsrpolicy monitor -c bu-iddnwserver3.iddlab.local -g Default
```

```
Workflow status:
data protection policy name:Backup
workflow name:Default
```

```

name:savegrp
job id:32525
type:backup action job
job state:ACTIVE
completion status:
start time: 5/26/15 16:59:43
duration: unknown

```

Displaying information about a workflow or backup action

To retrieve information about of a specific workflow or job action, type the following command:

```
nsrpolicy monitor -j job_id
```

For example, to view an information about a job with jobid 32524, type the following command:

```
nsrpolicy monitor -j 32524
```

```

activity progress: 1/1/0;
actual exit code: 0;
adhoc job: False;
authtype: ;
automatic: False;
Checkpoint restart ID: ;
Checkpoint restart sequence: ;
command: ;
completion severity: 10;
completion status: succeeded;
data protection policy name: Backup;
dependent job id: 0;
end time: 1432674065;
exit code known: True;
host: bu-iddnwserver3.iddlab.local;
input flag: False;
job id: 32524;
job log file: \
"C:\Program Files\EMC NetWorker\nsr\logs\policy\Backup\
\workflow_Default\
_032524";
job output: \
"133550 1432673983 1 0 0 4100 3352 0 bu-
iddnwserver3.iddlab.local nsrworkflow \
NSR notice 31 Starting %s '%s' workflow '%s'. 3 11 24
127405:Protection Policy\
0 6 Backup 0 7 Default
123316 1432673983 1 0 0 4100 3352 0 bu-
iddnwserver3.iddlab.local nsrworkflow N\
SR notice 46 Starting action '%s/%s/%s' with command: '%s'. 4 0
6 Backup 0 7 D\
efault 0 6 backup 0 32 savegrp -Z backup:traditional -v
123321 1432673983 1 0 0 4100 3352 0 bu-
iddnwserver3.iddlab.local nsrworkflow N\
SR notice 39 Action '%s/%s/%s's log will be in '%s'. 4 0 6
Backup 0 7 Default \
0 6 backup 23 75 C:\Program Files\EMC NetWorker\nsr\logs\
\policy\Backup\

```

```

Default\\backup_032525
123325 1432674065 1 0 0 4100 3352 0 bu-
iddnserver3.iddlab.local nsrworkflow N\
SR notice 21 Action '%s/%s/%s' %s. 4 0 6 Backup 0 7 Default 0 6
backup 0 9 suc\
ceeded
133553 1432674065 1 0 0 4100 3352 0 bu-
iddnserver3.iddlab.local nsrworkflow N\
SR notice 27 Workflow '%s/%s' succeeded. 2 0 6 Backup 0 7
Default";
job state: COMPLETED;
name: Backup;
ndmp flag: False;
NW Client name/id: ;
override parameters: ;
parent job id: 0;
policy definition changetime: 1431525315508563;
previous job id: 0;
protection groups: Default;
Reason job was terminated: ;
redirect stdio: False;
remote password: ;
remote user: ;
restricted data zone: ;
root parent job id: 0;
savegrp spawned: False;
sibling job id: ;
SSID: ;
start time: 1432673983;
type: workflow job;
type attributes: ;
type classes: ;
type help: ;
type name: ;
type references: ;
type table: ;
userid: ;
workflow name: Default;
resource identifier:
223.0.232.10.0.0.0.0.192.87.83.85.172.21.21.102(9);

```

Reporting recover job status

When you perform a recover by using the **NMC Recovery** wizard, NetWorker records the status of the recover operation and job activities. There are two ways to report job activities:

- In the **Recover** window for the NetWorker server in NMC. [Monitoring NetWorker Server activities in the Administration window](#) on page 55 describes how to view the recover status in the **Recover** window.
- By querying the job status by using `nsrrecomp` command on the NetWorker server. [Using nsrrecomp](#) on page 706 provides more information.

Using nsrrecomp

Use the `nsrrecomp` program to query the jobsdb for information about recover jobs and to create a recover completion report. The name specified for the recover job is the name of the saved recover configuration. The `nsrrecomp` program differs from the `jobquery` program because it also queries recover log files and is limited to recover job information only.

Example: Summary report of recover jobs

To generate a summary report of each recover job in the jobsdb, type:

```
nsrrecomp -L
```

Example: Recovery job completion report

To generate a completion report for recover job, type:

```
nsrrecomp -b -l recover_job_name
```

where `-b -l` is optional and used to override the default 2kb limit for job output.

Example: Summary report of the last recovery job

To generate a summary of last recovery job for a Recover resource, type:

```
nsrrecomp -H group_name
```

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrrecomp` program.

Checkpoint-enabled backup reporting

The `daemon.raw` file on the NetWorker Server contains details about groups that are run with checkpoint-enabled clients. When a group backup is completed, policy notifications also report the status of each client backup.

View the policy reports for checkpoint-enabled client backups

There are several things to consider when reviewing the summary notification report and the action report for the backup status of a checkpoint-enabled client.

- When a checkpoint-enabled client backup attempt fails:
 - The action is reported as a failure.
 - The failed save sets are reported in the **Unsuccessful Save Set** status section.
- When a checkpoint-enabled client backup succeeds:
 - The action status is reported as a success.
 - The total number of partial save sets that make up the checkpoint save sets is displayed in the **Save Set Summary** section.
 - The successful save sets are reported in the **Successful Save Set** status section.

Determine the status of a checkpoint-enabled backup

Review the `daemon.raw` file on the NetWorker server to determine the status of a checkpoint-enabled client backup.

```
nsrd info, Savegroup Info: group_name:client_name checkpoint
enabled, mode: mode. (severity 0, message 71193)
```

This message is reported when a backup action starts. This message reports the names of the clients that are checkpoint-enabled, and the mode that was selected at the time of the backup.

```
savegrp test: checkpoint restartable saveset
client_name:save_set created in previous run(s) of the group.
It will be checkpoint restarted. Checkpoint ID cp_id.
```

This message reports that a partial save set is detected for a client in the group and a checkpoint restart occurs for the save set.

```
savegrp group_name checkpoint restartable saveset
client_name:save_set failed and will not be restarted.
```

This message is reported when the backup of a checkpoint-enabled client fails and the backup will not be retried.

Common reasons for this error message include:

- The restart window for the group has been exceeded.
- The maximum number of client retries has been reached.

NOTICE When this message is reported, the failed save set are removed from an AFTD:

```
nsrd info, MeDia Info: save set save_set for client
client_name was aborted and removed from volume
volume_name (severity 0, message 71193)Recovering data.
```

```
savegrp group_name: checkpoint restartable saveset
client_name:save_set completed without interruption.
```

This message reports that the save set for a checkpoint-enabled client successfully completed during the backup action.

SNMP traps

The NetWorker Simple Network Management Protocol (SNMP) Module allows NetWorker servers to send notification messages to SNMP management agents.

You must configure SNMP-enabled network management software to accept traps from the NetWorker server. For detailed information about SNMP management operations refer to your network management documentation.

The NetWorker SNMP Module uses traps to communicate NetWorker event notifications to SNMP management stations. A trap is an unsolicited notification sent from the SNMP agent (the NetWorker server) to the SNMP event manager.

When you configure the SNMP notification in NetWorker, you can define the types of traps that the NetWorker server sends to the SNMP event manager. Typical traps include warnings, critical errors, and other messages from the NetWorker server.

NetWorker 18.2 and later includes the following SNMP trap enhancements:

- SNMP v2c MIB support
 - SNMP trap alert for NetWorker to comply with the MIB format
 - SNMP MIB support for successful or failed notifications for backup, clone, and restore
- Note:** SNMP v2c traps are not supported for client initiated RMAN triggered backup in NMDA Oracle and for client initiated `brbackup` in NMSAP Oracle.
- File system save set with backup and clone workflows with the following payload or attributes:

Table 115 SNMP attributes and descriptions

Attribute	Description
Server	FQDN, IP address, or the hostname of the NetWorker server.
Job Type	Save and Clone.
Job ID	Root parent job ID (based on the action and not on the client).
Status	Completion status (Success or failure).
Client	FQDN, IP address, or the hostname of the client.
Severity	Completion severity code. ▪ JOB_IND_DEBUG = 0 ▪ JOB_IND_INFORMATION = 10 ▪ JOB_IND_NOTICE = 20 ▪ JOB_IND_WARNING = 30 ▪ JOB_IND_INTERVENTION = 40 ▪ JOB_IND_ERROR = 50 ▪ JOB_IND_ERROR = 50 ▪ JOB_IND_SEVERE = 60 ▪ JOB_IND_CRITICAL = 70 ▪ JOB_IND_ALERT = 80 ▪ JOB_IND_EMERGENCY = 90 ▪ JOB_IND_INVALID = -1
Exitcode	The actual exit code.
StartTime	The start time of the action in human readable format.
EndTime	The end time of the action in human readable format.

Note: The NetWorker MIB file is available at the following location:

- On Windows: `<C:\Program Files\EMC NetWorker\nsr\snmp>`
- On Linux: `/opt/nsr/snmp`

Receive SNMP traps on Linux

To receive SNMP traps on Linux, do the following:

Before you begin

Install the following packages:

- net-snmp-agent-libs
- net-snmp-libs
- net-snmp

Procedure

1. Copy the MIB file for NetWorker located at the following location to the Linux trap receiver folder (/usr/share/snmp/mibs):
 - Linux NetWorker server: /opt/nsr/snmp
 - Windows NetWorker server: C:\Program Files\EMC NetWorker\nsr\snmp
2. Add the following to the snmptrapd.conf file:


```
authCommunity log,execute,net public
```
3. On the NetWorker server, specify the SNMP notification command at the policy and action levels:
 - For Linux: /usr/sbin/nsrtrap <IP address of the Linux Trap Receiver>
 - For Windows: "C:\Program Files\EMC NetWorker\nsr\bin\nsrtrap.exe" <IP address of the Windows Trap Receiver>
4. On the Linux trap receiver, execute the following commands to receive SNMP traps:


```
snmptrapd -f -C -c ./snmptrapd.conf -Le -M /usr/share/snmp/mibs -m NETWORKER-MIB
```
5. Initiate the policy or workflow on the NetWorker server to view the traps.

Example 13 Using PuTTY to capture the session output for SNMP traps Saving SNMP traps log file to a predefined location

If you are using PuTTY to capture the session output for SNMP traps, do the following:

1. Launch PuTTY.exe.
2. Navigate to **Session > Logging > All session output** to capture the SNMP traps in a log file.
3. Type the following command:


```
snmptrapd -f -C -c ./snmptrapd.conf -Le -M /usr/share/snmp/mibs -m NETWORKER-MIB
```

To save the SNMP traps log file to a predefined location, type the following command:

Example 13 Using PuTTY to capture the session output for SNMP traps Saving SNMP traps log file to a predefined location (continued)

```
nohup snmptrapd -f -C -c ./snmptrapd.conf -Lf /var/log/
snmptrap.log -M /usr/share/snmp/mibs -m NETWORKER-MIB &
```

The `nohup` option ensures that the command runs even after the session is terminated. Use `&` to run the command in the background.

Receive SNMP traps on Windows

To receive SNMP traps on Linux, do the following:

Before you begin

For SNMP Trap receiver:

- Install the iReasoning MIB browser.
- In the Address field, specify the IP of the NetWorker server.
- Select **Tools > Trap Receiver > Trap Receiver Settings**.
- Select UDP and port 162.

Procedure

1. On the NetWorker server, specify the SNMP notification command at the policy and action levels:
 - For Windows: "C:\Program Files\EMC NetWorker\nsr\bin\nsrtrap.exe" <IP address of the Windows Trap Receiver>
 - For Linux: /usr/sbin/nsrtrap <IP address of the Linux Trap Receiver>
2. Load the NetWorker MIB file located at the following location to the iReasoning MIB Browser using the **Load MIBs** option on the iReasoning MIB Browser.
 - Linux NetWorker server: /opt/nsr/snmp
 - Windows NetWorker server: C:\Program Files\EMC NetWorker\nsr\snmp
3. Navigate to **iReasoning MIB Browser > Trap Receiver**.
4. Initiate the policy or workflow on the NetWorker server to view the traps.

Configuring NetWorker SNMP notifications

The NetWorker software provides notifications to a variety of resources about NetWorker server events. The NetWorker SNMP module is one of those resources.

The module uses the `nsrtrap` program to forward notifications to the SNMP management software. To configure `nsrtrap` to send SNMP notifications to the SNMP server, you must configure a Notification resource on the NetWorker server and the SNMP server to receive the SNMP notifications. When you configure the SNMP notification, you include the IP address or hostname of the SNMP management server, and other `nsrtrap` command line options, for example, the SNMP community and the trap type.

Configuring SNMP notifications in NetWorker

You can create an SNMP notification or modify a preconfigured SNMP notification.

Before you begin

Before you configure the NetWorker SNMP notification, you must first license the NetWorker SNMP module. Contact Dell EMC Licensing for more information.

Procedure

1. On the **NetWorker Administration** window, click **Server**.
2. On the **Server** window, select **Notifications**, and perform one of the following actions:
 - Right-click **SNMP notification request**, and select **Properties**.
 - Right-click **Notifications**, and select **New**.
3. In the **Name** attribute, specify the name of the notification.

 **Note:** You cannot modify the **Name** attribute for an existing notification.

4. Optionally, in the **Comment** field, specify a description of the notification.
5. In the **Event** and **Priority** attributes, select the events and priorities that the notification should communicate to the SNMP server.

 **Note:** You cannot modify the **Event** and **Priority** attributes for an existing notification.

6. In the **Action** attribute, specify the options for the `nsrtrap` command:

```
nsrtrap -c community_name -t trap_type -s specific_trap_type
SNMP_server_name
```

The following table summarizes the available `nsrtrap` options.

Table 116 Command-line options for `nsrtrap`

Option	Description
<code>-c community</code>	Specifies the SNMP community that is authorized to receive traps from the NetWorker server. You configure SNMP communities on the SNMP server. The default setting for this option is Public, which means that the public community can receive traps from the NetWorker server. For security purposes, system administrators often customize the SNMP server to limit the communities that can accept traps. If the SNMP server configuration specifies a community other than Public, specify the community name.
<code>-t trap_type</code>	Optional, sets the type of trap that the NetWorker SNMP Module sends to the SNMP server. The default setting is six, which sets the trap type to “enterprise-specific” and is the correct type for the notifications (error messages) that the NetWorker server sends to the SNMP server. Only modify the trap type if you intend to send a specific trap to the SNMP server and not a NetWorker notification.

Table 116 Command-line options for nsrtrap (continued)

Option	Description
<code>-s <i>specific_trap_type</i></code>	Optional, allows you to customize the identity of the type of trap that the NetWorker server sends. Set this option to any integer value. Use this option along with different SNMP notifications to distinguish different traps from the NetWorker server. For example, you can create separate SNMP notifications for critical messages, warnings, and events or priorities then use the <code>-s</code> option with a unique number to differentiate the various notifications. The Action attribute for each notification appears as follows: • Critical notification: <code>nsrtrap -s 1 <i>SNMP_server_host_name</i></code> • Warning notification: <code>nsrstrap -s 2 <i>SNMP_server_host_name</i></code> • Event or priorities notification: <code>nsrtrap -s 3 <i>SNMP_server_host_name</i></code> Configure the SNMP management software to recognize that NetWorker traps with the specific trap type of 1 are critical messages, trap type 2 are warning messages and trap type 3 are event or priority messages. Additional SNMP notifications can have other settings for the <code>-s</code> option to further differentiate various traps from the NetWorker server.
<code>-v</code>	Sets the output mode to verbose. When you run <code>nsrtrap</code> from the command line in verbose mode, the program displays the community, trap type, specific trap type, and the hostname or IP address.

7. Click **OK**.

Configuring SNMP notifications at the policy level

You can configure SNMP notifications at the policy level.

Before you begin

Before you configure the NetWorker SNMP notification, you must first license the NetWorker SNMP module. Contact Dell EMC Licensing for more information.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Policies**, and then select **Properties**.
The Policy Properties dialog box appears.
3. Under Notifications:
 - a. Select **On Completion** to send a notification on completion of the workflows and actions in the policy.
 - b. Specify the following command to set notifications for Windows or Linux.
 - For Linux: `/usr/sbin/nsrtrap <IP address of the Trap Receiver>`

- **For Windows:** "C:\Program Files\EMC NetWorker\nsr\bin\nsrtrap.exe" <IP address of the Trap Receiver>

4. Click **OK**.

Configuring SNMP notifications at the action level

You can configure SNMP notifications at the action level.

Before you begin

Before you configure the NetWorker SNMP notification, you must first license the NetWorker SNMP module. Contact Dell EMC Licensing for more information.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, expand **Policies**, and then select an existing policy.
3. Select a workflow associated to the selected policy.
4. Select an action.
5. Right-click the action, and select **Properties**.
The Policy Action wizard appears.
6. On the **Specify the Advanced Options** window, under **Notifications**, do the following:
 - a. Select **On Completion** to send a notification on completion of the workflows and actions in the policy.
 - b. Specify the following command to set notifications for Windows or Linux.
 - **For Linux:** /usr/sbin/nsrtrap <IP address of the Trap Receiver>
 - **For Windows:** "C:\Program Files\EMC NetWorker\nsr\bin\nsrtrap.exe" <IP address of the Trap Receiver>
7. Click **OK**.

SNMP traps for restore

Use either NMC or the NetWorker CLI to perform a restore.

Procedure

1. On the NetWorker Administration window, click **Server**.
2. Under **Notifications**, select **Create a new notification using the SNMP notification command**.
The Create Notification window appears.
3. In the **Name** attribute, type a name for the notification.
4. In the **Event** attribute, select **Recover**.
5. In the **Priority** attribute, select the priorities of the corresponding actions.
6. In the **Action** attribute, type a command to run in response to the selected events and priorities.
 - **For Linux NetWorker server:** /usr/sbin/nsrtrap <IP address of the Linux Trap Receiver>

- For Windows NetWorker server: "C:\Program Files\EMC NetWorker\nsr\bin\nsrtrap.exe" <IP address of the Windows Trap Receiver>
7. Initiate restore using either NMC or the NetWorker CLI.
 - NMC: Open NMC and navigate to **Recover > New Recover**.
 - CLI: Go to the NetWorker client on which restore has to be performed and execute the `recover` command.

For more information on receiving SNMP traps, see the following sections:

- [Receive SNMP traps on Linux](#)
- [Receive SNMP traps on Windows](#)

Note:

- Restore workflow configuration is applicable for all file systems as well as NetWorker Modules.
- When recovering save sets of NMDA Oracle and NMSAP Oracle, multiple SNMP v2c traps (one notification per save set) are sent.

View SNMP traps on Linux Trap Receiver

After you configure SNMP notifications at the policy or action level, you can view the SNMP traps.

Procedure

1. Place the **NETWORKER-MIB** file in the `/usr/share/snmp/mibs` directory.
2. Execute the following command from the Linux machine, which is configured as the Trap Receiver:

```
# snmptrapd -f -C -c ./snmptrapd.conf -Le -M /usr/share/snmp/mibs -m NETWORKER-MIB
```

Figure 76 Sample log output

```
[root@pvp05 ~]# snmptrapd -f -C -c ./snmptrapd.conf -Le -M /usr/share/snmp/mibs -m NETWORKER-MIB
NET-SNMP version 5.5
2018-11-01 03:20:17 winn155.laas.com [UDP: [10.63.65.155]:57622->[10.63.65.65]]:
SNMPv2-SMI::mib-2.1.3.0 = Timeticks: (224085671) 25 days, 22:27:34.71  SNMPv2-SMI::snmpModules.1.1.4.1.0 = OID: NETWORKER-MIB::nwserverNotification  SNMPv2-SMI::mib-2.1.1.0 =
STRING: ""--- Traditional Backup Action report ---
Policy name:BC_Policy
Workflow name:WF-BC
Action name:backup
Action status:succeeded
Action start time:10/31/18 06:44:54
Action duration:0 hours 0 minutes 18 seconds
Total 1 client(s), 0 Succeeded with warning(s), 1 Succeeded, 0 Failed.
---Successful Backups---
q-rf-0400.laas.com:/space/VRGDATA/FSR_SA-1_L1/FSR_SA-1_L1_1, level=full, size 17 KB, duration 0 hours 0 minutes 3 seconds, 6 files
---Successful backups with warnings---
none
---Failed backups---
none
2018-11-01 03:20:17 securtest1.laas.com [UDP: [10.63.65.155]:57622->[10.63.65.65]]:
SNMPv2-SMI::mib-2.1.3.0 = Timeticks: (224085671) 25 days, 22:27:34.71  SNMPv2-SMI::snmpModules.1.1.4.1.0 = OID: NETWORKER-MIB::nwserverNotification  SNMPv2-SMI::mib-2.1.1.0 =
STRING: ""---Summary notification report---
Policy name:BC_Policy
Workflow name:WF-BC, Workflow status:succeeded, Workflow start time:10/31/18 06:44:54, Duration:0 hours 0 minutes 19 seconds
Action name:backup, Action status:succeeded, Action start time:10/31/18 06:44:54, Duration:0 hours 0 minutes 18 seconds*
2018-11-01 03:20:17 winn155.laas.com [UDP: [10.63.65.155]:57624->[10.63.65.65]]:
SNMPv2-SMI::mib-2.1.3.0 = Timeticks: (224085671) 25 days, 22:27:34.71  SNMPv2-SMI::snmpModules.1.1.4.1.0 = OID: NETWORKER-MIB::successfulActionNotification  NETWORKER-MIB::nw
over = STRING: "winn155.laas.com"  NETWORKER-MIB::client = STRING: "q-rf-0400.laas.com"  NETWORKER-MIB::type = STRING: "Backup"  NETWORKER-MIB::id = STRING: "
11"  NETWORKER-MIB::status = STRING: "succeeded"  NETWORKER-MIB::severity = STRING: "10"  NETWORKER-MIB::exitcode = STRING: "0"  NETWORKER-MIB::starttime = STRING: "10/31
/18 06:44:54"  NETWORKER-MIB::endtime = STRING: "10/31/18 06:45:12"
```

View SNMP traps on Windows Trap Receiver

After you configure SNMP notifications at the policy or action level, you can view the SNMP traps.

About this task


To view SNMP v2c traps, several third party MIB Browsers are available, such as iReasoning, Power SNMP Manager, SNMP J Manager, and so on. To receive SNMP traps using the iReasoning MIB Browser:

Procedure

1. On the iReasoning MIB Browser, select **Tools > Trap Receiver**.
2. Select **Trap Receiver Settings**, and set the port number as 162 and the protocol as UDP.
3. In the **Address** field, type the NetWorker server IP address.
4. Select **File > Load MIBs** to load the NETWORKER-MIB file

Note: There will be two SNMP notifications captured, the first one is the Legacy Notification and the second one is the OID based SNMP notification, containing the fields: Server, client, jobtype, jobid, status, severity, exitcode, starttime, and endtime.

Figure 77 SNMP trap output



Configuring SNMP management software

You must configure the SNMP management software to recognize and accept traps sent by NetWorker servers.

For specific instructions that describe how to configure the types of acceptable traps in the SNMP management software, refer to the SNMP management software documentation.

NetWorker SMI Network Management Private Enterprise Code

When you configure the SNMP management software to accept traps, you must indicate the specific trap type. Use the Structure of Management Information (SMI) Network Management Private Enterprise Code that applies to the specific network

application that will send traps to the software. The Private Enterprise Code for the NetWorker server is 160. The complete code is .1.3.6.1.4.1.160.

Receiving traps in the SNMP network management software

After you configure the SNMP network management software to accept traps from NetWorker servers, an icon for each NetWorker server appears on the network management console.

You can configure the SNMP network management software in the following ways:

- To indicate that a trap was received. For example, the NetWorker server icon may blink or change color.
- To track pending, alert, and other configured messages.
- To separate traps into event categories, such as Error Events, Status Events, Threshold Events, Configuration Events, Application Alert Events, or All Events. For information on how to set up SNMP trap templates, refer to the network management software documentation.

SNMP v2c trap support for NetWorker Modules

The following table provides information on the SNMP v2c trap support for NetWorker Modules.

Modules	Backup		Clone		Restore	
	GUI	CLI	GUI	CLI	GUI	CLI
File system	18.2	18.2	18.2	18.2	19.1	19.1
Virtualization (vProxy)	19.1	19.1	19.1	19.1	19.1	19.1
Oracle	19.1	19.1 ^a	19.1	19.1 ^b	19.1 ^c	19.1 ^d
SAP Oracle	19.1	19.1 ^e	19.1	19.1 ^f	19.1 ^g	19.1 ^h
BBB	19.1	19.1	19.1	19.1	19.1	19.1
SQL	19.1	19.1	19.1	19.1	19.1	Not Applicable
NMDP	19.1	19.1	19.1	19.1	19.1	19.1
BMR	19.1	19.1	Not Applicable	Not Applicable	19.1	Not Applicable
NSM	19.1	19.1	19.1	19.1	19.1	19.1
NMMEDI	19.1	19.1	19.1	19.1	19.1	19.1

- SNMP v2c traps are not supported for client initiated RMAN scripts.
- SNMP v2c traps are not supported for client initiated RMAN scripts.
- Multiple notifications (one notification per save set) for a single restore.
- Multiple notifications (one notification per save set) for a single restore.
- SNMP v2c traps are not supported for client initiated RMAN scripts.
- SNMP v2c traps are not supported for client initiated RMAN scripts.
- Multiple notifications (one notification per save set) for a single restore.
- Multiple notifications (one notification per save set) for a single restore.

NetWorker Notifications

A notification provides information about events that occur in a NetWorker environment. You can configure the events to be reported and how the NetWorker server reports them to you. Specific programs can be run when an event occurs, including third-party programs. By default, the NetWorker server sends notifications to log files that are located in the `NetWorker_install_dir\logs` directory on Windows and the `/nsr/logs` directory on UNIX.

Preconfigured notifications

NetWorker is preconfigured to provide most of the event notifications that are required to monitor NetWorker events. The following table lists these preconfigured notifications and the associated actions that are performed by the NetWorker server.

Table 117 Preconfigured notifications

Notification	Default action
Bus/Device Reset	Windows: Provides the syntax for the <code>smtpmail</code> command to send an email to the administrator account stating that a bus or device reset has been detected. The action attribute must be modified to replace <i>mailserver</i> with the actual hostname of the mail server. Using smtpmail to email notifications on page 724 describes how to customize the <code>smtpmail</code> command. Linux: Sends an email to the root account stating that a bus or device reset has been detected.
Cleaning cartridge expired	Windows: Reports to the <code>C:\Program Files\EMC NetWorker\nsr\logs\media.log</code> file that a cleaning cartridge has expired. Linux: Sends an email to the root account stating that an expired cleaning cartridge has been detected.
Cleaning cartridge required	Windows: Reports to the <code>C:\Program Files\EMC NetWorker\nsr\logsmedia.log</code> file that a device cleaning is required. Linux: Sends an email to the root account stating that a cleaning cartridge is required.
Client install	Windows: Reports the hostname and NetWorker client software version

Table 117 Preconfigured notifications (continued)

Notification	Default action
	information to the C:\Program Files\EMC NetWorker\nsr\logs\media.log file. Linux: Sends an email to root account: <code>host</code> <code>host_name</code> installed <code>product_version</code>. Where <i>host_name</i> is the name of the NetWorker host, and <i>product_version</i> is the NetWorker client software release and build number.
Device cleaned	Windows: Reports that a device has been cleaned to the C:\Program Files\EMC NetWorker\nsr\logs\media.log file. Linux: Sends an email to the root account stating that a device cleaning operation has completed.
Device cleaning required	Windows: Reports that a device requires cleaning to the C:\Program Files\EMC NetWorker\nsr\logs\media.log file. Linux: Sends an email to the root account stating that a device requires cleaning.
Device disabled	Windows: Reports that a device has been automatically disabled to the C:\Program Files\EMC NetWorker\nsr\logs\media.log file. Linux: Sends an email to the root account stating that NetWorker automatically disabled a device.
Device ordering issue detect	Windows: Provides the syntax for the <code>smtpmail</code> command to send an email to the administrator account with the message Check system device ordering. Moving device on NetWorker_server to service mode. To correct this issue, scan for devices in NMC and re-enable the device. The action attribute must be modified to replace <i>mailserver</i> with the actual hostname of the mail server. Using smtpmail to email notifications on page 724 describes how to customize the <code>smtpmail</code> command.

Table 117 Preconfigured notifications (continued)

Notification	Default action
	Linux: Sends an email to the root account with the message <code>Check system device ordering. Moving device on NetWorker_server to service mode. To correct, scan for devices in NMC and re-enable the device.</code>
Event log	Windows only. Logs notification events that are triggered by events and priorities to the Event Log.
File system full - recovering adv_file space	Launches the <code>nsrim</code> program to remove aborted and expired save sets. Used with advanced file type devices only.
File system full - waiting for adv_file space	Windows: Reports that the advanced file volume is full to the <code>C:\Program Files\EMC NetWorker\logs\media.log</code> file. Linux: Sends an email to the root account stating that an advanced file volume is full.
Inactive Files Alert	Windows: Reports that the space occupied by inactive files exceeds configured threshold to the <code>C:\Program Files\EMC NetWorker\nsr\logs\messages</code> log file. Linux: Sends an email to the root account stating that the space occupied by inactive files exceeds configured threshold.
Index size	Windows: Reports a message that the size of the index will soon exceed the space available to the <code>C:\Program Files\EMC NetWorker\nsr\logs\index.log</code> file. Linux: Sends an email to root with the message <code>Check the size of the client file index because it will soon exceed the space available.</code>
Log default	Windows: Sends data about NetWorker events to the <code>C:\Program Files\EMC NetWorker\nsr\logs\messages</code> log file. Linux: Directs data about the NetWorker events to logger. The logger utility sends the event with a tag of <code>daemon.notice</code> to the Operating system log file defined in the system log configuration file, for example <code>syslog.conf</code> .

Table 117 Preconfigured notifications (continued)

Notification	Default action
NetWorker Daemons Not Running	Windows: Provides the syntax for the <code>smtpmail</code> program to send an email to the administrator account stating that NetWorker daemons are not running on the NetWorker server. The action attribute must be modified to replace <code>mailserver</code> with the actual hostname of the mail server. Using smtpmail to email notifications on page 724 describes how to customize the <code>smtpmail</code> program. Linux: Sends an email to the root account stating that NetWorker daemons are not running on the NetWorker server.
New Virtual Machine	Windows: Reports a message that new virtual machines have been detected to the <code>C:\Program Files\EMC NetWorker\nsr\logs\messages</code> log file. Linux: Sends an email to the root account stating that new virtual machines have been detected.
Registration	Windows: Sends messages about the registration status of the NetWorker products to the <code>C:\Program Files\EMC NetWorker\nsr\logs\messages</code> log file. Linux: Sends an email to root with this message <code>Check the registration status.</code>
Resource File Corruption	Windows: Provides the syntax for the <code>smtpmail</code> program to send an email to the administrator account stating that resource file corruption has been detected on the NetWorker server. The action attribute must be modified to replace <code>mailserver</code> with the actual hostname of the mail server. Using smtpmail to email notifications on page 724 describes how to customize the <code>smtpmail</code> program. Linux: Sends an email to the root account stating that resource file corruption has been detected on the NetWorker server.
Save set marked suspect	Windows: Provides the syntax for the <code>smtpmail</code> program to send an email to the administrator account when a save set has been marked suspect.

Table 117 Preconfigured notifications (continued)

Notification	Default action
	The action attribute must be modified to replace <i>mailserver</i> with the actual hostname of the mail server. Using smtpmail to email notifications on page 724 describes how to customize the <code>smtpmail</code> program. Linux: Sends an email to the root account when a save set has been marked suspect.
SNMP notification request	Sends event notifications to a network management console. This notification occurs when the NetWorker SNMP module has been purchased and enabled. Configuring NetWorker SNMP notifications on page 710 provides details on SNMP notifications
Tape mount request 1	Windows: Requests that media be mounted in a device and displays a pending message in the <code>C:\Program Files\EMC NetWorker\nsr\logs\messages</code> log file. Linux: Sends a request message to the system logger to mount a backup volume, using a local0 facility and an alert level.
Tape mount request 2	Windows: Requests that media be mounted in a device and displays a critical message. Linux: Sends a request message to the system logger to mount a backup volume, using a local0 facility and an alert level.
Tape mount request 3	Windows: Sends a request to mount a backup volume with a priority of Alert, to the <code>C:\Program Files\EMC NetWorker\nsr\logs\media.log</code> file. Linux: Sends an email to the root account requesting that the tape be mounted.
Tape mount request 4	Windows: Provides the syntax for the <code>smtpmail</code> program to send an email to the administrator account that a Tape mount request 4 event has occurred. The action attribute must be modified to replace <i>mailserver</i> with the actual hostname of the mail server. Using smtpmail to email notifications on page 724 describes how to customize the <code>smtpmail</code> program.

Table 117 Preconfigured notifications (continued)

Notification	Default action
	Linux: Sends an email to the root account stating that a Tape mount request 4 event has occurred.
Verify Label failed on unload	Windows: Provides the syntax for the <code>smtpmail</code> program, to send an email to the administrator account stating that a label verification on unload operation has failed. The action attribute must be modified to replace <code>mailserver</code> with the actual hostname of the mail server. Using smtpmail to email notifications on page 724 describes how to customize the <code>smtpmail</code> program. Linux: Sends an email to the root account stating that a label verification on unload operation has failed.
Volume Marked full	Windows: Provides the syntax for the <code>smtpmail</code> program to send an email to the administrator account stating that a volume has been marked full. The action attribute must be modified to replace <code>mailserver</code> with the actual hostname of the mail server. Using smtpmail to email notifications on page 724 describes how to customize the <code>smtpmail</code> program. Linux: Sends an email to the root account stating that a volume has been marked full.
Volume Scan needed	Windows: Sends an event notification to the <code>C:\Program Files\EMC NetWorker\nsr\logs\media.log</code> file with a message that a volume with the Scan needed flag is detected. Linux: Sends an email to the root account with a message that a volume with the Scan needed flag is detected.

Customizing notifications

Notifications require the following three elements:

- Events
- Actions
- Priorities

About Events

An event signals that user intervention is required. For example, if a NetWorker server needs a new tape, the server alerts users to the situation by posting an event to the **Console** window.

NetWorker software generates an event that is based on various factors, including the following scenarios:

- The software or hardware encounters an error that requires user intervention to resolve.
- A NetWorker savegroup has failed.
- Drive ordering or serial number mismatch issues — a description of the problem is provided, along with a corrective action to fix the problem.
- Capacity monitoring — for example, reaching the space threshold on the deduplication node.
- NetWorker software is unable to poll a host it is monitoring for events or for generating reports.
- A license or enabler code that is managed by the License Manager is about to expire.

Some situations do not result in the generation of an event. For example, when a license managed by the NetWorker Console (instead of by the License Manager) approaches its expiration date. In this situation, a message is recorded in the NetWorker logs, but an event is not generated until the expired license causes a backup to fail. Check the **Administration** window from time to time for important messages.

Actions

The Actions attribute defines the action that the NetWorker server takes after an event notification occurs. The following table provides a summary of actions.

Table 118 Actions

Action	Description
eventlog	Windows only, logs the notification message to the event log. Priority determines whether the notification is an error, warning, or information-only message.
nsrlog	Windows only, sends a message about an event to a file. Use option <code>-f</code> to identify a specific file. For example: <pre>nsrlog -f log file path</pre> If no option is specified, then messages go to the <code>/nsr/logs/messages</code> file.  Note: The <i>log file path</i> must not include any space or special characters, and must not be enclosed in quotes.
logger	UNIX only, uses the UNIX syslog facility (<code>/usr/bin/logger</code>) to log information or send messages.
lp	UNIX only, prints the notification.
mail	UNIX only, sends an email to the specified user.

Table 118 Actions (continued)

Action	Description
<code>sendmail</code>	NetWorker Virtual Appliance (NVE), sends an email to a specified user.
<code>smtpmail</code>	Windows only, sends an email to the specified user.
<code>nsrtrap</code>	Sends notifications to an SNMP management console. Use with the following options: • <code>-c</code> community (if not specified, then the default public is used) • <code>-f</code> file (reads message from a file and sends as snmp trap.) • <code>-i</code> version (if not specified, then the default version is SNMPV2) • <code>-s</code> specific (default is NetWorker enterprise assignment, which is 1) • <code>-t</code> trap (default trap is #6 which is the enterprise-specific trap) • <code>-u</code> snmp uptime • <code>-v</code> verbose

Third-party programs can also be used for the action, if the programs support reading from standard input.

For example:

- On UNIX systems, you can use a third-party email program rather than the mail program.
- On Windows systems, you can use a third-party email program rather than the `smtpmail` program to send the information to other locations, such as an email address or pager system.

Only users who belong to the NetWorker server Administrators list, or a member of the Application Administrators user group, can change the Action attribute of an existing notification.

Using smtpmail to email notifications

Use the `smtpmail` program included with the NetWorker software on Windows systems to email an event notification to a list of specified email addresses.

About this task

The `smtpmail` program requires:

- A mail server that allows SMTP relays.
- An active TCP/IP connection. This command does not have dialing capabilities.

The `smtpmail` command reads the message that is sent from standard input.

The message is terminated in one of the following ways:

- An EOF.
- CTRL-Z on console.

- A line consisting of a single period (.).

To use the `smtpmail` program to email event notifications:

Procedure

1. From the **Administration** window, click **Server**.
2. Click **Notifications**.
3. Right-click the notification, then select **Properties**. The **Properties** dialog box appears.
4. In the **Action** attribute, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver
recipient2@mailserver...
```

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
 - `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.
 - `recipient1@mailserver`—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.
5. Click **Ok**.

Priorities

Each NetWorker event has a series of associated messages, and each message has an associated priority. The preconfigured notifications have selected priorities based on the importance of the message being sent. For example, the first time the NetWorker server sends a mount backup volume request, the priority that is assigned to the message is Waiting. The priority of the second request is Alert. The priority of the third request is Critical.

The following table lists the priorities on which notifications are based.

Table 119 Priorities

Priority	Description
Information	Information about the current state of the server.
Notice	Important information.
Warning	A non-fatal error has occurred.
Waiting	The NetWorker server is waiting for an operator to perform a routine task, such as mounting a backup volume.
Alert	A severe condition exists that requires immediate attention.
Critical	The server detected an error that should be fixed.
Emergency	A condition exists that may cause NetWorker to fail unless corrected immediately.

 **Note:** Event priorities are sorted alphabetically, rather than by severity.

Creating a custom notification

About this task

NetWorker also provides preconfigured notifications. [Preconfigured notifications](#) on page 717 provides a complete list of preconfigured notifications.

Procedure

1. From the **Administration** window, click **Server**.
2. Right-click **Notifications**, then select **New**. The **Create Notification** dialog box appears.
3. In the **Name** attribute, type a name for the notification.
4. In the **Event** attribute, select the events to be acted on.
5. In the **Priority** attribute, select the priorities of the corresponding actions.
6. In the **Action** attribute, type a command to run in response to the selected events and priorities.
7. Click **Ok**.

Editing a notification

About this task

 **Note:** You cannot change the name of a notification.

Procedure

1. From the **Administration** window, click **Server**.
 2. Click **Notifications**.
 3. In the right pane, perform one of the following tasks:
 - To modify multiple attributes in a single configuration resource by using the **Notification Properties** window, right-click the staging configuration and select **Properties**.
 - To modify a specific attribute that appears in the resource window, place the mouse in the cell that contains the attribute that you want to change, then right-click. The menu displays an option to edit the attribute. For example, to modify the **Comment** attribute, right-click the resource in the **Comment** cell and select **Edit Comment**.
-  **Note:** To modify a specific attribute for multiple resources, press and hold the **Ctrl** key, select each resource, and then right-click in the cell that contains the attribute that you want to change. The menu displays an option to edit the attribute.
4. Make any required changes, then click **OK**.

Copying a notification

Procedure

1. From the **Administration** window, click **Server**.
2. Click **Notifications**.
3. Right-click the notification to copy, then select **Copy**. The **Create Notification** dialog box appears, containing the same information as the notification that was copied, for **Name** attribute.

4. In the **Name** attribute, type a name for the new notification.
5. Edit any other attributes as appropriate, then click **OK**.

Deleting a custom notification

About this task

 **Note:** You cannot delete preconfigured notifications.

Procedure

1. From the **Administration** window, click **Server**.
2. Click **Notifications**.
3. Right-click the notification to delete, then select **Delete**.
4. When prompted, click **Yes** to confirm the deletion.

Configuring owner notifications

Owner notification is an attribute of the NetWorker Client resource. Use this attribute to send an email to a user with the results of the backup of the individual client.

About this task

For Windows NetWorker servers, use the `smtpmail` program to send the owner notification email. [Using smtpmail to email notifications](#) on page 724 describes how to configure the `smtpmail` program.

For UNIX NetWorker servers, use the `/usr/ucb/mail` program or a third-party mail application to send the owner notification.

Procedure

1. From the **Administration** window, click **Protection**.
2. Select **Clients** in the left navigation pane.
3. Right-click the client, and select **Properties**.
4. Select **Globals (2 of 2)**.
 - For a Windows NetWorker server, use the `smtpmail` program to configure email notifications. [Using smtpmail to email notifications](#) on page 724 describes how to configure `smtpmail`.
 - For a UNIX NetWorker server, use the `/usr/ucb/mail` program:

```
/usr/ucb/mail -s "subject" recipient1@mailserver
recipient2@mailserver...
```

For example:

```
/usr/ucb/mail -s "Backup status for client xyz in group abc"
debbie@mymailhost.com
```

5. Click **OK**.

Results

When the group containing the client completes, the notification is sent to the recipient email address defined in the Owner notification attribute.

For example:

```
-----Original Message-----
From: Super-User [mailto:root@NWserver.corp.com]
```

```

Sent: Thursday, March 22, 2012 12:45 PM
To: debbie@mymailhost.com
Subject: Backup status for client xyz in group abc
cdcsdunnd11c, savefs, "succeeded:full:savefs"
* cdcsdunnd11c:savefs savefs cdcsdunnd11c: succeeded.
cdcsdunnd11c, C:\cmdcons\system32, "<NULL>:full:save"
* cdcsdunnd11c:C:\cmdcons\system32 cdcsdunnd11c:C:\cmdcons
\system32 aborted
* cdcsdunnd11c:C:\cmdcons\system32 Termination request was sent
to job 64006 as requested; Reason given: Aborted

```

Logging event notifications

NetWorker keeps two general notification log files. By default, these files are located in `<NetWorker_install_dir>\logs`:

- The messages log file (Windows only) — The data in the messages log file is generated by `nsrlog`, a program that is part of the NetWorker event notification mechanism. The `nsrlog` program is triggered by a notification, and it prints the message to the messages log file.
- The `daemon.raw` log file — The `nsrd`, `nsrexecd`, and their subordinate processes redirect their output to the `daemon.raw` log file.

To better access and use these event logs in Windows systems, an Event Logging mechanism enables applications to the application event log, and access them from any computer that has the Windows Event Viewer. The Event Viewer enables you to look selectively at the messages that interest you by filtering messages based on the categories that are listed in this table.

Table 120 Event Viewer messages

Event Viewer category	Displayed information
Source	Events from NetWorker software always designate NetWorker as the source.
Category	Mapped from NetWorker notification event type (server, registration, and so on).
Severity	Mapped from NetWorker notification priority: • Critical and Emergency are mapped to Error. • Priorities between Alert and Warning are mapped to Warning. • Notification and Information are mapped to Information.
Event ID	Events from NetWorker software always designate the numeral 1 for the ID.

Breakthrough logging

Breakthrough logging feature in NetWorker helps you to understand the steps involved in the various operations such as Save, Recover and Clone. Each step is logged in the defined order to ensure successful completion of Save, Recover and Clone operations. Breakthrough logging helps the user to review the log associated with each step of operation and determine the step that failed during the execution of that particular operation.

You can check logs for clone and backup at `/nsr/logs/Policy/<policy name>`.

You can check logs for recovery at `/nsr/logs/recover`.

Front-end Capacity Estimation

NetWorker supports an automatic reporting mechanism that communicates with Dell EMC's Usage Intelligence portal. You must install the EMC Secure Remote Services (ESRS) appliance version 3.20.20.08 or later from the [ESRS Virtual Edition Product Page](#), and configure NetWorker to communicate with the appliance. Review the [ESRS v3 Installation Training](#) video for details about how to install the ESRS appliance.

The ESRS RAP resource can be configured to send periodic license, configuration and usage information to Dell EMC as well as track the liveness of NetWorker servers. Several reports are sent, the details of the reports are extracted from command line tools.

The command line tool `nsrcainfo`, generates an estimate of the total data protected in a NetWorker datazone. The capacity estimate uses a simple heuristic where it measures the maximum full backup for each application type and each client in the datazone, this is defined as the client's capacity. The sum of each individual client's capacity provides a capacity estimate for the entire datazone. Configuration information is extracted from the RAP database through the command line tool `nsrdump`. The `nsrdump` tool automatically hides sensitive information like passwords, but can also be configured to hide other information that customers may not wish to share with Dell EMC.

The *NetWorker Command Reference Guide* provides more details on `nsrcainfo`.

Configuring EMC Secure Remote Services (ESRS)

Procedure

1. Install the ESRS appliance v3.20.20.08 or later. You can access ESRS documentation and downloads from the [ESRS Virtual Edition Product Page](#).

 **Note:** This download is separate and is not a part of the NetWorker package.

2. Ensure that NetWorker can reach the host port 9443.
3. In the NetWorker Management Console (NMC) click **Server**, and then select **ESRS**.

Figure 78 ESRS Properties

ESRS Properties - Default ESRS configuration.

General System

Identity

Name: Default ESRS configuration.

Comment: Enter comments here

ESRS hostname: myesrshost

ESRS port: 9443

Username: john.doe@abc.com

Password:

Product key:

ESRS instance id: ELMBET0517TZRF

E-mail contact: Administrator@bu-iddnwserve

Configuration

Send now: ☒

Frequency in weeks: 4

Day of the week: Monday

Hour: 8

Exclude attributes: hostname, datazone pass phrase, password, user id

Exclude resources: NSR lockbox

OK Reset Cancel

4. In the **Identity** area, type information for the fields in the following table.

Table 121 ESRS fields and descriptions

Field	Description
ESRS hostname or IP	Hostname or IP of the ESRS.
ESRS port	TCP port number for the ESRS service. The default port is 9443.
Username	Username is your online support account email. Go to the Support website at https://support.emc.com to check your account. For more details on Support website, see the Preface section of the <i>NetWorker Administration Guide</i> .
Password	Password is your online support account password.
Comment	(Optional) Description of this resource or other explanatory remarks.
E-mail contact	(Optional) Administrator's email address.

During the ESRS Provisioning process, after you use your online support credentials to authenticate, you select your particular site ID from a list of site IDs associated with your login. An email containing a code will be sent to the associated email address. The code that is sent is used to complete the provisioning. The [EMC Secure Remote Services Installation Guide](#) provides information.

The following fields are populated by NetWorker.

Field	Description
Name	Identifier for the ESRS resource.
Product key	After a successful registration, ESRS returns the product key. This key is used for authentication with the ESRS service. (Hidden Attribute) The key automatically gets updated after the ESRS host is successfully registered with the NetWorker Server.
ESRS instance id	The value that is returned from the ESRS host to identify this NetWorker Server instance. The id automatically gets updated after the ESRS host is successfully registered with the NetWorker Server.

5. (Optional) In the **Configuration** area you can change the defaults.

Field	Description
Send now	To send the database immediately select the checkbox. Once the database is sent, NetWorker by default clears the checkbox. Review <code>daemon.raw</code> for errors: - On Windows—Reports are in the <code><NW_installlocation>\nsr\applogs\rh</code> folder. - On UNIX—Reports are in the <code>/nsr/applogs/rh</code> folder.
Frequency in weeks	Select the frequency in weeks that the local database is sent to ESRS. 0 disables scheduled sends. The default frequency is set to 4.
Day of the week	Choose the day of the week that the local database is sent to ESRS. The default day is set to Monday.
Hours	Select the hour of the day that the local database is sent to ESRS. The default hour is set to 8 for 8am.
Exclude attributes	To prevent <code>nsrdump</code> from reporting certain attributes through ESRS, type the list of comma-separated attributes. Excluded attribute names apply to all resources across the database. The default attributes are hostname, datazone pass phrase, password, and user id.
Exclude resources	To prevent <code>nsrdump</code> from reporting certain resources through ESRS, type the list of comma-separated resources. The default excluded resource is set to NSR lockbox.

6. Click **OK**.

If you want to cancel the information you entered click **Reset**.

Troubleshooting ESRS

EMC Secure Remote Services (ESRS) registration fails

Issue

The registration could fail due to:

- Incorrect ESRS hostname, IP, or port
- Incorrect ESRS username or password
- Firewall blocks access to the ESRS port

Fix

Use the correct connection information and open the ESRS port on the firewall.

For incorrect ESRS username or password, go to the Support website at <https://support.emc.com> to check your online support account and verify that you have access to the NetWorker asset. If you still have issues, create a support ticket for ESRS deployment issues that are caused due to username or password issues.

 **Note:** For more details on Support website, see the Preface section of the *NetWorker Administration Guide*.

EMC Secure Remote Services (ESRS) keepalive or file transfer fails

Issue

The failure could be due to invalid credentials:

- ESRS has not been configured and registered
- ESRS instance id has been edited manually
- ESRS product key has been modified

Fix

Configure ESRS and verify that the registration call obtained an ESRS instance id:

- Undo modification of ESRS instance id and/or the product key.
- Alternatively, force the ESRS re-registration by deleting the contents of the ESRS instance id in the NSR ESRS resource.

CHAPTER 12

NetWorker Server Monitoring

This chapter contains the following topics:

- [Enterprise events monitoring](#) 734
- [Monitoring NetWorker Server activities in the Administration window](#) 737
- [Monitoring changes to the NetWorker and NMC Server resources](#) 752
- [Monitoring user access to the NMC server](#) 753
- [Monitoring NetWorker server activities in the log files](#) 753

Enterprise events monitoring

The NetWorker Management Console (NMC) makes the administration of servers more efficient by providing a centralized means of monitoring activity throughout an enterprise.

You can view details of current NetWorker and Data Domain systems. [Managing various servers in the Enterprise](#) on page 757 provides details on adding hosts to be monitored.

Information that can be monitored includes activities and operations that are related to devices and libraries, and events that require user intervention.

An event signals that user intervention is required. For example, if a NetWorker server needs a new tape, the server alerts users to the situation by posting an event to the **Console** window.

NetWorker generates an event that is based on various factors, including the following examples:

- Software or hardware errors that require user intervention to resolve.
- Failed backups.
- Drive ordering or serial number mismatch issues.
A description of the problem is provided, along with a corrective action to fix the problem.
- Capacity monitoring, such as reaching the space threshold on the deduplication node
- Inability to poll a host for event monitoring or report generation.
- Impending expiration of a license or enabler code that is managed by the License Manager.

Some situations do not result in the generation of an event. For example, when a license managed by the NetWorker Console (instead of by the License Manager) approaches its expiration date. In this situation, a message is recorded in the NetWorker logs, but an event is not generated until the expired license causes a backup to fail. Check the **Administration** window from time to time for important messages.

Polling interval for system events

You can set the polling interval for system-level events and activities in the **System Options** dialog box.

Polling interval configuration is available for the following items:

- Events and reporting (in seconds).
- NetWorker activities (in seconds).
- Data Domain events (in seconds).
- NetWorker libraries (in hours).

 **Note:** Event polling for NetWorker libraries can occur a maximum of once per hour.

[Setting system options to improve NMC server performance](#) on page 784 provides information on setting polling intervals.

Enabling or disabling event capture for a host

Enable the **Capture Events** option for a host in the NMC to enable event monitoring for the host. This option is selected by default when you add a host.

Procedure

1. From the NMC GUI, click **Enterprise**.
2. Right-click the host, and select **Properties**.
3. Enable or disable event capture for the host by selecting or clearing the **Capture Events** checkbox.
4. If the host is a Data Domain system, select the **Configure SNMP Monitoring** tab.
 - a. Type `public` in the **SNMP Community String** box.
 - b. Type the value of the SNMP process port that is used by all Data Domain systems that are monitored by the NMC in the **SNMP Process Port** box.
The default port is 162.
 - c. In the **SNMP Traps** list, select the checkbox next to the Data Domain system events that you want to monitor with NetWorker.
5. Click **OK**.

Event viewing

Events appear in the lower right pane of the **Console** window.

The following table describes the information that appears in the columns for each event.

Table 122 NMC event information

Column	Description
Priority	Represents the relative severity of the problem by displaying one of seven icons.
Server Name	Identifies the host that caused the event to be generated.
Server Type	Identifies the type of server to which the event belongs. Server types include but are not limited to NetWorker and Data Domain.
Time	Indicates the day of the week and time that the Console server discovered the problem. The time which an event is reported is always based on the time zone of the Console server. For example: If a backup fails at 11:00 A.M. in New York, a Console server in Los Angeles reports the event as occurring at 8:00 A.M. The time format depends on the current locale setting. Start date and time formats on page 680 provides more information.
Category	Classifies the source of the problem.

Table 122 NMC event information (continued)

Column	Description
Message	Displays the text of the error message that generated the event.
Annotation	Displays an icon when an annotation has been made. An annotation provides a place to record comments that are associated with an event, and can accommodate more information than the Note column. Each annotation can be up to 12 KB. For example, use annotations to log steps that are taken to resolve an event. You can add multiple annotations to a single event, but you cannot edit or delete annotations. To add or view annotations, right-click the event and select Annotation.
Note	Provides an editable field for making brief administrative information that is associated with an event. For example: - Name of the NetWorker administrator or operator that is assigned to the event. - Letters or numbers that allow the sorting of events into a preferred order. To add, edit, or delete a note, double-click the cell in the Note column for the event. When you finish adding, editing, or deleting the note, click outside the cell. The maximum number of characters for a note is 30.

Event priorities

Each event is designated with one of seven possible priorities. When the **Console** window sorts events by priority, it lists the events in alphabetical order, with Emergency between Critical and Information.

The following table provides more information on each type of event priority.

Table 123 Event priorities

Icon	Priority	Description
	Alert	Error condition that is detected by the NetWorker server that should be fixed by a qualified operator.
	Critical	Severe error condition that demands immediate attention.
	Emergency	Condition exists that may cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority.

Table 123 Event priorities (continued)

Icon	Priority	Description
	Information	Information about the current state of the server. This icon represents the lowest priority.
	Notification	Important information.
	Waiting	Indication that the NetWorker server is waiting for an operator to perform a routine task, such as mounting a tape.
	Warning	Non-fatal error has occurred.

Dismissing an event

After you view and act on an event, you can dismiss the event from the **Console** window to prevent other users from acting unnecessarily on events that have already been resolved.

About this task

 **Note:** Dismissing an event makes it disappear from the **Console** window for all NetWorker users.

Procedure

1. From the **Console** window, right-click the event and select **Dismiss**.
A confirmation message appears.
2. Click **Yes**.

Results

There are slight differences in how event dismissals are handled, depending on the source:

- Events from NetWorker software are automatically dismissed in the **Console** window when the problem that triggered the event is resolved.
- Events from device ordering or serial mismatch issues are automatically dismissed in the **Console** window when the problem is resolved via the corrective action provided.

Monitoring NetWorker Server activities in the Administration window

The **Monitoring** window in the NetWorker **Administration** application enables you to monitor the activities of an individual NetWorker Server.

About this task

The **Monitoring** window provides the following types of activity and status information:

- Data protection policies, workflows, and individual actions.

- Cloning, recovering, synthetic full backups, and browsing of client file indexes.
- Operations that are related to devices and jukeboxes.
- Alerts and log messages.

You can also perform some management operations from the **Monitoring** window, for example, starting, stopping, or restarting a data protection policy.


Procedure

1. From the **NMC Console** window, click **Enterprise**.
2. In the **Enterprise** view, right-click the NetWorker Server, and then select **Launch Application**.

The **Administration** window appears.

3. To view the **Monitoring** window, click **Monitoring**.

Figure 79 Monitoring window



About the Monitoring window

On the **Administration** window taskbar, select **Monitoring** to view the details of current NetWorker server activities and status, such as:

- Policies and actions.
- Cloning, recovering, synthetic backups, checkpoint restart backups, and browsing of client file indexes.
- Alerts and log messages, and operations that are related to devices and jukeboxes.

While the **Monitoring** window is used primarily to monitor NetWorker server activities, it can also be used to perform certain operations. These operations include starting, stopping, or restarting a workflow.

The **Monitoring** window includes a docking panel that displays specific types of information. Select the types of information you want to view from the docking panel.

A portion of the **Monitoring** window, which is known as the task monitoring area, is always visible across all windows. A splitter separates the task monitoring area from the rest of the window. You can click and move the splitter to resize the task monitoring area. The arrow icon in the upper right corner of the **Monitoring** window allows you to select which tasks you want to appear in this view.

Smaller windows appear within the **Monitoring** window for each window. Each smaller window, once undocked, is a floating window and can be moved around the page to customize the view. You can select multiple types from the panel to create multiple floating windows that can be viewed simultaneously. The following table describes the various types of information available in the docking panel, and the details each one provides.

Table 124 Monitoring window panel

Window	Information provided
Policies/Actions	The Policies tab provides you with status information about all configure policies and the associated workflows and actions. The Actions tab provides you with status information for all actions. Policies/Actions pane on page 741 provides more information.
Sessions	Allows you to customize whether to display all session types, or only certain session types. The information that is provided depends on which session type you select. For example, if you select Save Sessions , the window lists clients, save sets, groups, backup level, backup start time, duration of the backup, devices, rate, and size. Sessions window on page 57 provides more information.
Alerts	Lists the priority, category, time, and message of any alerts. Alerts pane on page 58 provides more information.
Devices	Lists devices, device status, storage nodes, libraries, volumes, pools, and related messages. Devices pane on page 59 provides more information.
Operations	Lists the status of all library and silo operations, including <code>nsrjb</code> operations that are run from the command prompt. Also lists user input, libraries, origin, operation data, operation start time, duration of the operation, progress messages, and error messages. When displaying Show Details from the Operations window, the length of time that the window is displayed depends on the value that is typed in the Operation Lifespan attribute on the Timers tab of the Properties dialog box for the corresponding library. To access library properties, click Devices in the taskbar. By default, this pane is hidden.
Log	Lists messages that are generated by the NetWorker server, including the priority of each message, the time the message was generated, the source of the message, and the category. Log window on page 62 provides more information.

Customizing the Monitoring window

This section describes how to customize the **Monitoring** window in the **Administration** interface.

Customizing tables

You can customize the organization and display of tabular information in the **Monitoring** window.

Sorting tables

You can change the display of tabular information that appears in the window. You can sort Table grids by column heading, and then by alphabetic or numeric order within those columns.

1. Drag and drop the column heading to its new position.
2. Click the column heading to sort the items into alphabetic and numeric order. An arrow appears in the column heading to indicate the sort order.

Sorting selected rows in a table

Selected rows are sorted to the top of the table. This sorting is particularly useful when you select **Highlight All** from the Find panel to select all rows matching the Find criteria and then moving all selected rows to the top of the table to view the results.

1. From the **Edit** menu, select **Find**, or press **Ctrl + F** to view the **Find** panel.
2. To select the rows, click each row or use the Find criteria.
3. Select **Sort Selected**.

Sorting multiple columns in a table

You can select the column that you want to use as the tertiary sort key, the secondary sort key, and the primary sort key.

1. Click the column that you want to use as the last sort key.
2. Click the column that you want to use as the next-to-last sort key, and so on, until you select the primary column.

Displaying columns in a table

You can select which columns to display in a table.

1. From the **View** menu, select **Choose Table Columns**.
2. Click a column name to select or clear the column and then click **OK**. You can also select the columns to display by right-clicking a table header and selecting **Add Column** from the drop-down.

Displaying panes

You can choose to show or hide panes in the **Monitoring** window.

About this task

Perform the following steps to hide or show a pane in the **Monitoring** window.

Procedure

1. From the **View** menu, select **Show**. A check mark appears beside the panes that appear in the **Monitoring** window.
2. To hide a pane, select a marked pane.

A check mark does not appear beside the pane.

3. To show a pane, select an unmarked pane.

A check mark appears beside the pane.

Policies/Actions pane

The **Policies/Actions** pane provides you with the ability to review status information about policies and actions.

This pane has two tabs:

- **Policies**—Provides a navigation tree that displays all configured policies on the NetWorker Server. Expand each policy to display the workflows that are associated with each policy. Expand each workflow to display each action that is contained in the workflow.
- **Actions**—Provides a list of all Action resources.

Policies pane

The **Monitoring** window in the **NetWorker Administration** window enables you to monitor activities for specific policies, workflows, and actions.

The **Policies/Actions** pane at the top of the **Monitoring** window lists the policies on the NetWorker Server by default. Click the + (plus) sign next to a policy in the list to view the workflows in the policy, and the + (plus) sign next to a workflow to view the actions for a workflow.

The **Policies** pane provides the following information for each item (where applicable):

- Overall status

The following table provides details on the status icons that may appear in the **Policies** pane.

Table 125 Policy status icons

Icon	Status
	Never run
	Running
	Succeeded
	Failed
	Probing
	Interrupted
	Queued
	Cloning

Table 125 Policy status icons (continued)

Icon	Status
	Consolidating (NetWorker Server 8.2.x and lower only)

 **Note:** When the schedule for an action is skip, the status of the action appears as Never Run and the status of the Workflow is Succeeded.

- Most recent start time.
- Duration of the most recent run.
- Next scheduled runtime.
- Name of the assigned save set.
- Device on which the save set is stored.
- Backup level.
- Data transfer rate.
- Size of the save set.
- Messages that resulted from an action.

Right-click an action in the **Policies** pane and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

When you sort the items on the **Policies/Actions** pane by using the **Status** column, NetWorker sorts the items in alphabetical order that is based on the label of the icon.

Consider the following when a policy/action is in a probing state:

- A message is sent when the group starts and finishes the probe operation.
- The results of the probe operation (run backup/do not run backup) are also logged.
- Probes do not affect the final status of the group, and the group status does not indicate the results of the probe.
- If probing indicates that a backup should not run, then the group status reverts to its state before the group running.
- Check the results of the probe in the **Log** window to ensure that the probe indicates that the backup can be taken.

Actions pane

To view a list of all actions, click the **Actions** tab at the bottom of the **Policies** pane. The **Policies** pane becomes the **Actions** pane.

The **Actions** pane provides the following information for each action:

- Overall status

 **Note:** The **Actions** pane displays the same status icons as the **Policies** pane.
- Name
- Assigned policy
- Assigned workflow
- Type
- Date and time of the most recent run

- Duration of the most recent run
- Percent complete, for actions that are in progress
- Next scheduled runtime

Right-click an action in the **Actions** pane and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

Workflow operations

This section describes how to use the **Monitoring** window to start, stop, and restart workflows.

Starting, stopping, and restarting policies

The workflows in a policy can run automatically, based on a schedule. You can also manually start, stop, and restart specific workflows by using the the NMC **NetWorker Administration Monitoring** window.

About this task

You can restart any failed or canceled workflow. Note, however, that the restart must occur within the restart window that you specified for the workflow. Additionally, for a VMware backup, if you cancel a workflow from **NetWorker Administration** and then want to restart the backup, ensure that you restart the workflow from the **NetWorker Administration** window. If a workflow that was started from **NetWorker Administration** is restarted from the **vSphere Web Client**, the backup fails.

Procedure

1. In the **Monitoring** window, select the workflow or actions.
2. Right-click and then select **Start**, **Stop**, or **Restart**.

A confirmation message appears.

 **Note:** You cannot stop, restart, or start individual actions.

3. Click **Yes**.

Viewing workflow backup details

Perform the following steps to view backup details for workflows.

Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Policies** in the docking panel, and expand the Policy that you want to monitor.
3. Right-click the workflow, and then select **Show Details**. The **Workflow Summary** window appears.
4. In the **Workflow runs** pane of the **Workflow Summary** window, select the workflow.
5. Click **Show Messages**. In the **Show Messages** window, select one of the following options:
 - Get Full Log—To display all messages.
 - Print—To print the log.
 - Save—To save the log to a local file.
 - OK—To close the **Show Messages** window.

6. Click **OK** to close the **Workflow Summary** window.

Viewing action backup details

Perform the following steps to view backup details for actions.

Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Actions** in the docking panel.
3. In the **Actions** pane, right-click the action, and then select **Show Details**. The details window for the action appears.
4. Review the information in the **Actions Messages** pane. To display detailed information from the action log file, click **Show Action Logs**, and then select one of the following options:
 - Get Full Log—To display all messages.
 - Print—To print the log.
 - Save—To save the log to a local file.
 - OK—To close the **Show Messages** window.
5. In one of the Actions detail panes, for example, the **Completed successfully** pane, select the action that you want to review.
6. Click **Show Messages**. In the **Show Messages** window, select one of the following options:
 - Get Full Log—To display all messages.
 - Print—To print the log.
 - Save—To save the log to a local file.
 - OK—To close the **Show Messages** window.
7. Click **OK** to close the **Details** window.

Sessions window

Use the **Sessions** window to view the sessions that are running on a NetWorker server. You can change the view of this window to display these sessions:

The **Sessions** pane below the **Policies/Actions** pane provides details on individual save, recover, clone, and synthetic full sessions by client.

To view all sessions or to limit the list of sessions by the session type, click the tabs at the bottom of the **Sessions** pane. Session types include:

- Save
- Recover
- Clone
- Browse
- Synthetic Full/Rehydrated Sessions
- All

To change the displayed session types go to **View > Show**, and select the type of sessions to display. To display all sessions currently running on the NetWorker Server, regardless of type, select **All Sessions**.

You can stop a session (backup, synthetic full backup, clone, and recovery sessions) from the **Monitoring** window, even if the session was started by running the `nsrworkflow` command.

To stop a session, right-click the workflow in the pane, and select **Stop** from the list box.

Alerts pane

The **Alerts** pane displays alerts that are generated by a particular NetWorker server or Data Domain system that has devices that are configured on the NetWorker server. The **Alerts** pane includes priority, category, time, and message information.

An icon represents the priority of the alert. The following table lists and describes each icon.

Table 126 Alerts window icons

Icon	Label	Description
	Alert	Error condition detected by the NetWorker server that should be fixed by a qualified operator.
	Critical	Severe error condition that demands immediate attention.
	Emergency	Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority.
	Information	Information about the current state of the server. This icon represents the lowest priority.
	Notification	Important information.
	Waiting	The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.
	Warning	A non-fatal error has occurred.

When items on the **Alerts** pane are sorted by the **Priority** column, they are sorted in alphabetical order based on the label of the icon.

Removing alerts

Remove individual alert messages from the **Events** tables by removing them from the **Events** table. To delete a message in the **Events** table, right-click the message, and select **Dismiss**.

About this task

 **Note:** The alert message remains in the **Log** window in the NetWorker Administration program.

Devices pane

The **Devices** pane allows you to monitor the status of all devices, including NDMP devices. If the NetWorker server uses shared and logical devices, the window is

adjusted dynamically to present a set of columns appropriate for the current configuration.

The **Devices** pane provides the following information:

- Status of the operation.
- Name of the device.
- Name of the storage node that contains the device.
- For tape devices, the name of the library that contains the device.
- Name of the volume in the device.
- Name of the pool that is associated with the volume.
- Last message generated for the device.
- Whether the operation requires user input.

For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape.

[Entering user input](#) on page 61 provides instructions on how to deal with a user input notification.

If the current server configuration includes a shared device, a **Shared Device Name** column appears on the **Devices** pane. The name of the shared device appears in the **Shared Device Name** column. If other devices for that configuration are not shared devices, then the **Shared Device Name** column is blank for those devices. Only a single device per hardware ID can be active at any particular moment. The information for inactive shared devices is filtered out, and as a result, only one device per hardware ID is presented on the window at any time.

An icon represents the device status. The following table lists and describes each icon.

Table 127 Devices status icons

Icon	Label	Description
	Library device active	The library device is active.
	Library device disabled	The library device is disabled.
	Library device idle	The library device is idle.
	Stand-alone device active	The stand-alone device is active.
	Stand-alone device disabled	The stand-alone device is disabled.
	Stand-alone device idle	The stand-alone device is idle.

When you sort items in the **Devices** pane by the **Status** column, NetWorker sorts the devices in alphabetical order based on the label name of the icon.

Operations window

The **Operations** window displays information about device operations. It provides the following information:

- Status of the operation.

- Name of the library.
 - Whether the operation requires user input.
For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape. [Entering user input](#) on page 61 provides instructions on how to deal with a user input notification.
 - The origin, or source, of the operation.
For example, the interface, nsrjb or the NetWorker server.
 - Time the operation started.
 - Type of operation.
 - Duration of the operation.
 - Status messages from the operation.
 - Any error messages.
-  **NOTICE** Only the last error message of the operation appears in the **Error Messages** column. Move the mouse pointer over the cell containing the last error message to display the entire list of error messages.

The operation status is represented by an icon. The following table lists and describes each of the icons.

Table 128 Operations window icons

Icon	Label	Description
	Failed	The operation failed.
	Queued	The operation is waiting in the queue to run.
	Retry	The operation failed, but may work if you try again.
	Running	The operation is running.
	Successful	The operation completed successfully.
	User Input	The operation requires user input.

When items on the **Operations** window are sorted by the Status column, they are sorted in alphabetical order based on the label of the icon.

Viewing operation details

The **Operation Details** dialog box opens, providing information about the completion of the operation. The **Completion Time** displays the time that the operation finished. The time that it took to complete the operation is the difference between the completion and start times of the operation.

About this task

To save operation details to a file, click **Save** in the **Operation Details** dialog box. When prompted, identify a name and location for the file.

Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.

3. Right-click the operation, then select **Show Details**.

Stopping an operation

Certain operations can be stopped from the **Operations** window.

Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation to stop, then select **Stop**.
4. Click **Yes** to confirm the stop.

Note: Operations that were started from a command line program, such as the `nsrjb` command, cannot be stopped from the **Operations** window. To stop these operations, press `Ctrl-C` from the window where the command was started.

Entering user input

If the system requires user input, select the labeling operation in slow/verbose mode and the **Supply User Input** icon appears.

Procedure

1. Right-click the operation, then select **Supply Input**.
2. Confirm the requirement to supply input.
 - If **Yes**, and input is supplied, the icon in the **User Input** column disappears.

Note: If two users try to respond to the same user input prompt, the input of the first user takes precedence, and the second user receives an error message.

- If **No**, and input is not supplied, the operation will time out and fail.

Log window

To view the most recent notification logs, click the **Log** window from the docking panel in the **Monitoring** window. The **Log** window provides the priority, time, source, category, and message for each log.

Note: If a particular log file is no longer available, check the log file on the NetWorker server. The log files are located in `NetWorker_install_path\logs` directory.

An icon represents the priority of the log entry. The following table lists and describes each icon.

Table 129 Icons in the Log pane

Icon	Label	Description
	Alert	Error condition that is detected by the NetWorker server that should be fixed by a qualified operator.
	Critical	Severe error condition that demands immediate attention.

Table 129 Icons in the Log pane (continued)

Icon	Label	Description
	Emergency	Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority.
	Information	Information about the current state of the server. This icon represents the lowest priority.
	Notification	Important information.
	Waiting	The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.
	Warning	Non-fatal error has occurred.

When you sort items on the **Log** pane by using the **Priority** column, NetWorker sorts the icons in alphabetical order based on the name of the label.

Recover window

The **Recover** window displays information about recover configurations that are created with the NetWorker Management Console (NMC) Recovery wizard.

You can use this window to:

- Start the NMC Recovery wizard to create recover configurations or modify saved recover configurations.
- Identify the status of a recover configuration that is created with the NMC Recovery wizard.
- Start and stop a recover job.


The **Recover** window is divided into five sections:

- **Toolbar**—The toolbar is hidden by default. To display the recovery toolbar, select **View > Show toolbar**
- **Summary**
- **Configured Recovers**
- **Currently Running**

A splitter separates the **Configured Recovers** section from **Currently running** window. You can click and move the splitter to resize these two windows.

The following table shows an example of the **Recover** window.

Figure 80 Recover window



Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. The following table summarizes the function of each toolbar button.

Table 130 Recovery toolbar options

Button	Function
	Starts the NMC Recover wizard to create recover configurations.
	Displays the Properties window for the saved recover configuration that you selected in the Configured Recover window.
	Deletes the saved recover configuration that you selected in the Configured Recover window.
	Displays online help for the Recover window.
	Displays the Find window at the bottom of the Recover window. The Find window allows you to perform keyword searches for messages that appear in the Logs window.
	Start the recover operation for a selected saved recover configuration. This option is only available for a recover configuration that has a Never run, or Failed status.
	Stop in-progress recover operation that you selected in the Currently Running window.

 **Note:** The **Recover** toolbar does not appear by default. To display the **Recover** toolbar, select **View > Show toolbar**.

Recover Summary

The Recover Summary section displays a high-level overview of recover jobs. This section includes the following information:

- Total Recovers—The total number of successful recover jobs.
- Since—The number of successful recover jobs since this date.

Configured Recovers

The **Configured Recovers** window displays a list of saved recover configurations in a tabular format. You can sort the information by column. The **Configured Recovers** table displays the following information for each saved recover configuration:

- Status—The job status of a saved recover configuration.
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB.
- Comment
- OS—The operating system of the source host.
- Recover requestor—The Windows or UNIX account used to create the recover configuration.
- Start Time
- End Time
- Start date

Table 131 Save recover configuration job status

Icon	Description
	The last recover attempt failed.
	The last recover attempt completed successfully.
	The recover job has never run.
	The recover job is scheduled to run in the future.
	The recover job has expired.

Currently running

The **Currently Running** window displays a list of in progress recover jobs in a tabular format. You can sort the information by column. The **Currently Running** table displays the following information for each job:

- Status
- Name
- Source client
- Destination client

- Recovery list
- Recover type—For example, file system or BBB
- Volume
- Comment
- Device
- Size
- Total size
- % complete
- Rate (KB/s)
- Start time
- Duration
- Currently running

Find

The **Find** section appears along the bottom of the **Recover** window, after you select the **Find** button on the **Recover** toolbar. **Find** allows you to search for keywords in the **Configured Recovers** window. The following table summarizes the available find options.

Table 132 Find options

Find option	Description
Find	Highlight the first saved recover configuration that contains the specified keyword.
Prev	Highlight the previous saved recover configuration that contains the specified keyword.
Highlight All	Highlights each saved recover configuration that contains the specified keyword.
Sort Selected	Sorts each highlighted recover configuration in the Configured Recover table so that they appear at the top of the Configured Recover table.
Match case	Make the keyword search case sensitive.

Monitoring changes to the NetWorker and NMC Server resources

NetWorker provides two ways to monitor changes made in to the NetWorker and NMC Server resources:

- **Monitor RAP (resource allocation protocol) attribute in the NetWorker Server resource** — This feature tracks both before and after information related to additions, deletions, or modifications to NetWorker Server resources and their attributes.
- **Security Audit Log feature** — This feature provides the NetWorker Server and the NMC Server with the ability to log specific security audit events related to their operations.

The *NetWorker Security Configuration Guide* describes how to use and configure the Monitor RAP attribute and the Security Audit Log feature.

Disabling or enabling the Monitor RAP Attribute

The Monitor RAP attribute is enabled by default. To change the setting, perform the following steps in the **Console** window.

Procedure

1. From the **Administration** window, select **View > Diagnostic Mode**.
2. Right-click the NetWorker server name in the left pane and select **Properties**.
3. In the **Setup** tab of the **NetWorker Server Properties** dialog box, select the Monitor RAP **Enabled** or the **Disabled** attribute as required.
4. Click **OK**.

Monitoring user access to the NMC server

NMC allows you to determine the last time that a user accessed the NMC user interface, and when the user logged out of the NMC user interface.

Before you begin

Log in to the NMC server as a Console Security Administrator. The NetWorker Authentication Service administrator account is a Console Security Administrator.

Procedure

1. On the toolbar, select **Setup**.
2. In the **User and Roles** navigation pane, select **Users**.
3. In the **Users** window pane, right-click click a column heading and select **Add columns**.
 - To monitor when a user last logged in to the NMC UI, select **Login Time**.
 - To monitor when a user last logged out of the NMC UI, select **Logout Time**.

Monitoring NetWorker server activities in the log files

NetWorker provides plain text and unrendered log files that enable you to monitor NetWorker server activities.

The Troubleshooting chapter provides a summary of the log files on each NetWorker host and how to manage the log files.

CHAPTER 13

NMC Server Management

This chapter contains the following topics:

• Enterprise	756
• Customizing the Console window and views	763
• Using the NMC filters	764
• Connecting to the NMC GUI using an ssh connection	766
• Backing up the NetWorker environment	766
• Using the NMC Configuration Wizard	768
• NMC server authentication	769
• Adding the NMC service account to the Users user group	774
• Enabling two factor authentication for AD and LDAP users	775
• Moving the NMC Server	776
• Migrating NMC users to the authentication service database	777
• Resetting the administrator password	779
• Changing the service port used by the NMC database	781
• Changing database connection credentials	783
• Updating the NMC server IP address/hostname	783
• Setting system options to improve NMC server performance	784
• Displaying international fonts in non-US locale environments	787
• NetWorker License Manager	787
• NMC error messages and corrective actions	788
• Console troubleshooting notes and tips	793

Enterprise

The Enterprise is a visual representation of the NetWorker Console control zone. You can monitor various servers in the enterprise such as the NetWorker and Data Domain servers for events. You can also generate various reports on events, backups, and user activity.

Enterprise components

Enterprise components include hosts and folders.

Hosts

A host, also known as a managed node, is the NetWorker or Data Domain server being monitored. A host terminates a branch in the Enterprise.

Folders

The purpose of folders is to enable the Enterprise to contain multiple levels. Each folder can contain more folders, more hosts, or more of both.

Organizing NetWorker servers

Use the Enterprise to organize the NetWorker servers by some logical or functional criteria.


Examples of organizational criteria include:

- By geography — For example, you can put all the hosts from the same city or country in the same folder.
- By function — For example, you can have the servers that back up web servers in one folder, and the servers that back up mail servers in another folder.
- By administrative divisions within the Enterprise — For example, you can use separate folders for servers that back up Marketing, Sales, or Engineering hosts.

You can create and maintain multiple folders to organize multiple copies of a host in the Enterprise. When you create each folder that is based on different organizational criteria, you can view the organization in different, yet parallel, and complementary ways.

Example: An enterprise that is arranged by geographic location

This figure provides an example of an Enterprise arranged by geographic location. There are three folders, one for each country that manages NetWorker servers: USA, France, and Australia. Each folder contains a number of hosts that correspond to the location of the NetWorker servers. The Australia folder, for instance, contains three host computers that are labeled perth1, perth2, and sydney.

Figure 81 NetWorker servers worldwide

Viewing the enterprise

In the **Console** window, you can view the organization of the NetWorker servers in much the same way as you use a file manager program to view the contents of a file system.

Procedure

1. From the **Console** window, click **Enterprise**.
 - The left pane displays folders and hosts in a tree-like arrangement to illustrate the organization of the NetWorker servers.
 - The right pane displays the contents of the selected folder or host.
2. Select a view option as described in the following table.

Table 133 Viewing the enterprise

To:	Perform the following action:
Show or hide contents of the Enterprise.	Click Enterprise .
Show or hide contents of a folder.	Click the folder.
Show the managed applications that are installed on a host computer.	Click the host.

Managing various servers in the Enterprise

NetWorker Console enables centralized management of NetWorker or Data Domain servers within the Enterprise. Use the **Console** window to add, delete, move, and copy servers.

When you use the NetWorker software to manage many NetWorker servers, you can use a single command `gstmodconf` from a command prompt to efficiently add or

delete multiple hosts. [Adding or deleting multiple servers by using a hostname file](#) on page 761 provides further information.

The server management activities include, but are not limited to, operations that are related to devices and libraries, and events that require user intervention.

Adding a managed host

The **Console** window can display server events and to generate server activity-reports.

About this task

 **Note:** When you configure a Data Domain device with the **New Device** wizard, the wizard adds Data Domain servers as a managed host. The *NetWorker Data Domain Boost Integration Guide* provides more information about Data Domain as a managed host.

Procedure

1. From the **Console** window, click **Enterprise**.
2. In the left pane, right-click **Enterprise**, then select **New > Host**. The **Add New Host** wizard appears.
3. Type a hostname, IP address, DNS name, or WINS name in the **Host Name** attribute, then click **Next**.

 **Note:** Hostnames and aliases cannot exceed 80 characters.

4. Select the server type and click **Next**.
5. Follow the instructions for configuring selected host type, then click **Finish**.

 **Note:** You can also use the **Console Configuration** wizard to add a host.

Deleting a host

You can delete a single host or multiple hosts within a folder.

Procedure

1. From the **Console** window, click **Enterprise**.
2. Right-click the host, then select **Delete**. The **Deleting Host** dialog box appears.
 - To delete multiple hosts, select multiple hosts in the details pane and select **Delete**.
 - If additional copies of the host exist in the Enterprise, use the **Delete all existing copies of the host** option to delete all instances of that same host in a single operation.
3. Click **Yes** to confirm deletion of the host.

Copying a host

You can create multiple copies of a host for a single NetWorker server. For example, you can create one copy of a host in the logical position of the host in the Enterprise, while another copy of the host is in a Hosts-to-Watch folder where you can easily monitor it. In this configuration, you can check the server without browsing through the Enterprise.

Procedure

1. From the **Console** window, click **Enterprise**.
2. Right-click the host, then select **Copy**.
3. Right-click a new location, then select **Paste**.

 **Note:** You can also use the drag-and-drop feature while press and holding the Ctrl key to copy hosts.

Moving a host

To move a host from one location to another in an Enterprise, perform the following steps.

Procedure

1. From the **Console** window, click **Enterprise**.
2. Right-click the host to move, then select **Move**.
3. Right-click a new location, then select **Paste**.

 **Note:** You can also use the drag-and-drop feature while holding down the Ctrl key to move hosts.

Managing folders in the enterprise

The NetWorker software allows you to manage folders within the Enterprise. This means that you can add, rename, delete, and move folders as needed.

You can add new folders directly beneath the Enterprise node or beneath other folders.

Adding a folder**Procedure**

1. From the **Console** window, click **Enterprise**.
2. Right-click the location within the Enterprise where you want the new folder to appear, then select **New > Folder**.

A new folder appears in the Enterprise with the default name Untitled1.

3. Highlight the default name and type a new name to replace it. The name must meet these criteria:
 - Include at least one, but no more than 80 characters.
 - Exclude forward slashes (/).
4. Press **Enter**.

Deleting a folder**Procedure**

1. From the **Console** window, click **Enterprise**.
2. Right-click the folder to delete, then select **Delete**.
 - If hosts exist in the folder, a dialog box prompts you to confirm the deletion of each host. Select **Yes** to continue with the operation, or **No** to cancel it.

- If hosts do not exist in the folder, the NMC server deletes the folder.
- If the folder contains any unique hosts (meaning hosts that do not have copies anywhere else in the Enterprise), an additional dialog box appears to confirm deletion of the unique host.

A separate dialog box with four options appears for each unique host in the folder:

- To delete the specified host, click **Yes**.
 - To delete all hosts and subfolders in the selected folder, without further prompts, click **Yes to All**.
 - To cancel the deletion, click **No**.
 - To cancel any further deletion of hosts in the selected folder, and leave the remaining contents intact, click **Cancel**.
- The NMC server deletes non-unique hosts, and folders containing only non-unique hosts without additional prompting.

 **NOTICE** If there are user group restrictions in place that control which hosts a user can view, the folder might appear empty.

Copying a folder

Procedure

1. From the **Console** window, click **Enterprise**.
2. Right-click the folder to copy, then select **Copy**.
3. Right-click a new location, then select **Paste**. A copy of the folder appears in its new location.

 **NOTICE** You can also use the drag-and-drop feature to copy folders while holding down the Ctrl key.

4. A folder cannot be copied within the same Enterprise level.

Moving a folder

Procedure

1. From the **Console** window, click **Enterprise**.
2. Right-click the folder to move, then select **Move**.
3. Right-click a new location, then select **Paste**. The folder appears in its new location.

 **NOTICE** You can also use the drag-and-drop feature to move folders while holding down the Ctrl key.

Renaming a folder

Procedure

1. From the **Console** window, click **Enterprise**.
2. Right-click the folder, then select **Rename**.
3. Highlight the folder name and type a new name to replace it. The name must meet these criteria:
 - Include at least one, but no more than 80 characters.

- Exclude forward slashes (/).
4. Press **Enter**.

Adding or deleting multiple servers by using a hostname file

For larger enterprises, use the `gstmodconf` command and a hostname file to add or delete multiple NetWorker servers to the Enterprise, with the features Capture Events and Gather Reporting Data enabled. [Using the `gstmodconf` command](#) on page 761 provides more information about the `gstmodconf` command.

Restrictions

Before you use the `gstmodconf` command, review the following restrictions.

If a host already exists anywhere in the Enterprise, either at the base or within a folder, you cannot use the `gstmodconf` command to add copies of the host.

You cannot use this command to add a host to a folder. You can only add a host to the base level of the Enterprise. After you add the host to the Enterprise, use the Console GUI to move the host to a folder. [Moving a host](#) on page 759 provides more information.

When you use the `gstmodconf` command to delete a host, the command only deletes hosts from the base level. The command does not delete hosts that are within folders.

Creating the hostname file

To use the `gstmodconf` command to add or delete multiple hosts simultaneously, specify the hostnames in a hostname text file.

To create a hostname file, use these guidelines.

- Only list one hostname on each line of the file.
- A non-comment line that contains more than one space-separated or tab-separated hostname generates an error.
- To include a comment in the file, start the line with a "#" character.
- Blank lines are treated as comments and ignored, as shown in the following example.

```
Hostname file
#This is a hostname file for XYZ Corporation
apple
banana
grape
kiwi
mango
nectarine
pineapple
strawberry
tangerine
```

Using the `gstmodconf` command

The `gstmodconf` command has this syntax:

```
gstmodconf -i file -f function -s server -k port -l username -P password
```

The *NetWorker Command Reference Guide* or the UNIX man pages provide a complete description of the command and its options.

The following provides an example of how to use `gstmodconf` to add nodes from the file, `xyz_hostlist`. In this example, the NMC server name is *myconsole* and the `xyz_hostlist` file contains the following entries:

```
apple
banan
grape
```

Example: Adding multiple hosts with the `gstmodconf` command

```
gstmodconf -s myconsole -i xyz_hostlist -l Administrator
```

```
Trying 111.22.3.444... connected
processing file 'xyz_hostlist'
adding host 'apple'
successfully added host 'apple'
adding host 'banana'
successfully added host 'banana'
adding host 'grape'
successfully added host 'grape'
//Closing connection
```

You can use the `gstmodconf -s myconsole -i xyz_hostlist -f delete` command to delete multiple Networker hosts.

 **Note:** The `gstmodconf` file on Windows is located in the following folder: `C:\Program Files\EMC NetWorker\Management\GST\bin`. The `gstmodconf` file on Linux is located in the following folder: `/opt/lgtomc/bin/`. The folder location is not in the default environment path.

Error messages generated by the `gstmodconf` command

This section describes two common error messages that can appear when you use the `gstmodconf` command.

The following provides an example of the error that appears when you use the `gstmodconf` command to add a host that exists in the Enterprise:

Example: Trying to add a host that already exists

```
% gstmodconf -s myconsole -i xyz_hostlist
Trying 111.22.3.444... connected
processing file 'xyz_hostlist'
adding host 'apple'
///Error!
{
string object_type = "gterror";
int severity = 16;
int reason = 23;
list msg = {
int level = 1;
string text = 'Host name already exists';
};
// Closing connection...
```

The following output provides an example of the error that appears when you use the `gstmodconf` command but you did not specify the administrator password when the password is not the default value.

Example: Trying to use `gstmodconf` without specifying the password

```
% gstmodconf -s myconsole -i xyz_hostlist
Trying 111.22.3.444... auth failed.
gt_session_connect: clnt_create: Remote system error-Connection
refused.
```

Customizing the Console window and views

This section describes how to customize the **Console** window.

Sorting tables

You can change the display of tabular information that appears in the window. You can sort Table grids by column heading, and then by alphabetic or numeric order within those columns.

1. Drag and drop the column heading to its new position.
2. Click the column heading to sort the items into alphabetic and numeric order. An arrow appears in the column heading to indicate the sort order.

For example: to see all the managed events about servers that were unreachable by the NMC server, perform the following steps.

1. From the **Console** window, select **Events**.
2. Drag the **Message** column until it is over the **Priority** column and drop it.
3. Click the **Message** column heading. A down-arrow appears.

Scan down the list of messages until you find all three servers with the message, `Unable to connect to server`. You can also generate a **Managed Event Details** report to get the same information, and then print, or export it for use in another application.

Sorting selected rows in a table

Selected rows are sorted to the top of the table. This sorting is particularly useful when you select **Highlight All** from the Find panel to select all rows matching the Find criteria and then moving all selected rows to the top of the table to view the results.

1. From the **Edit** menu, select **Find**, or press **Ctrl + F** to view the **Find** panel.
2. To select the rows, click each row or use the Find criteria.
3. Select **Sort Selected**.

Sorting multiple columns in a table

You can select the column that you want to use as the tertiary sort key, the secondary sort key, and the primary sort key.

1. Click the column that you want to use as the last sort key.
2. Click the column that you want to use as the next-to-last sort key, and so on, until you select the primary column.

Displaying columns in a table

You can select which columns to display in a table.

1. From the **View** menu, select **Choose Table Columns**.

2. Click a column name to select or clear the column and then click **OK**. You can also select the columns to display by right-clicking a table header and selecting **Add Column** from the drop-down.

Using the NMC filters


This section describes how to apply filters to view detailed information in the **Administration** window.

From the **Administration** window, you can use filters to search and view details about NetWorker server resources, recover configurations, devices, media, and hosts. Search fields and list boxes display on all NMC windows with filtering capability.

The search fields and list boxes allow you to filter information that appears on a page. By typing a value in the search fields or selecting an option from the list box, the display changes based on the values that you specified in the fields.

For example, in the **Protection > Policies** window, you can search and view details for a policy. By typing **Bronze** in the **Search Name** field, only the policies with the name **Bronze** appear in the list.

Figure 82 Using filters to search and view policies



In this example, the policy with the name **Bronze** displays and the **Send Notification** attribute is set to **On Completion**.

The following table describes how to use filters to search and view details in the **Administration** window.

Table 134 NMC windows with filtering capability

Window	Using filters to search and view details in the Administration window
Protection	To search and view detailed information about the following NetWorker server resources, type a value in the search fields or select an option from the list box, and then press Enter. • Clients • Groups • Policies • Probes • Schedules • VMware View
Recover	To search and view detailed information about recover configurations, type a value in the search fields or select an option from the list box, and then press Enter.

Table 134 NMC windows with filtering capability (continued)

Window	Using filters to search and view details in the Administration window
Devices	To search and view detailed information about devices, type a value in the search fields or select an option from the list box, and then press Enter . <i>Note:</i> Filtering is not available for Staging, VMware Backup Appliances, and VMware Proxies.
Media	To search and view detailed information about the following activities and resources, type a value in the search fields or select an option from the list box, and then press Enter . - Label Templates - Media Pools - Disk Volumes - Tape Volumes - Client Indexes - Save Sets <i>Note:</i> In the Save Sets window, filtering is not available on the Query Save Set tab.
Hosts	To search and view detailed information about hosts, type a value in the search fields or select an option from the list box, and then press Enter . <i>Note:</i> Filtering is not available on the Software Inventory or Software Repository tabs.
Server	To search and view detailed information about the following NetWorker server resources, type a value in the search fields or select an option from the list box, and then press Enter . - Directives - Notifications - Time Policies - User Groups

Toggle the display of items

- To show all available items, select **Filtered: 0 of 0 items**.
- To show only filtered items, select **Showing: 0 of 0 items**.

Hide or show filters

- To hide filters, right-click in the column header and select **Hide Filters (CTRL-H)**.
- To show filters, right-click in the column header and select **Show Filters (CTRL-L)**.

Clear filters

To clear all filters, select **Clear All**. The display returns to the default view.

Connecting to the NMC GUI using an ssh connection

You can use ssh port forwarding to connect to the NMC server and generate reports, from the NMC client.

About this task

Perform the following steps on the NMC client.

Procedure

1. Open an ssh connection from the NMC client to the NMC server with ssh tunnels for ports 9000 and 9001.

For example:

```
ssh -L9000:localhost:9000 -L9001:localhost:9001 -
L5432:localhost:5432 Console_servername -N
```

 **Note:** If you changed the default NMC server ports, specify the correct port numbers.

2. Use `javaws` to connect to the NMC server.

For example:

```
javaws https://localhost:9000/gconsole.jnlp
```

Backing up the NetWorker environment

When you install or upgrade the NetWorker Server, the installation or upgrade process creates a default Server Protection policy that backs up the NetWorker Server and the NMC Server database.

The Server Protection policy includes the following workflows for backing up the NetWorker environment:

- The NMC Server backup workflow performs a backup of the NMC database, which includes NMC Server management data such as report information. The database remains available during the backup. The workflow is scheduled to start a full backup daily at 2:00 p.m. The workflow is assigned to the default NMC Server group, which contains the NMC Server if you specified a NetWorker Server when you configured the NMC Server in the **Console Configuration** wizard.
- The server backup workflow performs a bootstrap backup of the NetWorker Server for disaster recovery purposes. The workflow is scheduled to start at 10:00 a.m. A full backup occurs on the first day of the month, and incremental backups occur the remaining days of the month. The workflow is assigned to the default Server Protection group, which contains a dynamically generated list of the Client resources for the NetWorker Server and the NMC Server.

 **Note:** The Server Protection policy also includes the server maintenance workflow, which performs an expire action to mark expired save sets as recyclable.

You can edit the default policy, workflows, groups, and actions, or create a set of policies for server backup and maintenance.

Configuring an NMC server database backup

The first time that you connect to the NMC GUI, the Console Configuration wizard prompts you to configure an NMC server database backup. If you did not configure the NMC database backup or you want to configure a new NetWorker server to backup the NMC server database, perform the following steps.

Before you begin

Connect to the NMC GUI with an account that has the Console Application Administrators role.

Procedure

1. On the toolbar, select **Setup**.
2. From the **Setup** window, select **Setup > Set Database Backup Server**.
3. In the **NetWorker server** field, specify the hostname of the NetWorker server that will backup the NMC server database.
4. Leave the **Create Client resource and add to the 'Server protection policy'** checkbox selected.
5. In the **Client name** field, specify the hostname of the NMC server.
6. Click **OK**.

Results

When you define an NMC database backup, the wizard performs the following actions on the NetWorker server:

- Creates a Client resource for the NMC server database backup. The **Save set** field for the client contains the path to the database staging directory. By default, the staging directory is in `C:\Program Files\EMC NetWorker\Management\nmldb_stage` on Windows and `/nsr/nmc/nmldb_stage` on Linux.
-  **Note:** The file system that contains the staging directory must have free disk space that is at least equal to the size of the current NMC database.
- Creates a group called NMC server.
- Adds the Client resource to the NMC server group.
- Creates a workflow that is called NMC server backup in the Server Protection policy. The workflow contains the NMC server backup action, which performs a full backup of the NMC server database every day at 2 P.M.
- Adds the NMC server group to the NMC server backup workflow.
-  **Note:** The NMC Server database backup only supports the full and skip backup levels. If you edit the NMC Server backup action and change the levels in the backup schedule to a different level, for example synthetic full, NetWorker performs a full backup of the database.

Changing the staging directory for NMC database backups

To backup the NMC database, the `savepsm` process creates a copy of the NMC database in a staging directory. After the backup operation completes, the `savepsm` process deletes the contents of the staging directory. By default, when you configure an NMC database backup, the configuration process sets the default staging directory to the `NetWorker_installation_directory\nsr\nmc\nmldb_stage` folder on Windows and the `/nsr/nmldb_stage` directory on Linux.

Before you begin

Use NMC to connect to the NetWorker server with a user that is a member of the Application Administrators or Database Administrators user group.

About this task

The size of staging database equals the size of the NMC database. Ensure that the file system on which the `savepsm` process writes the staging database has sufficient free disk space. To change the location of the staging directory, perform the following steps:

Procedure

1. On the **Protection** window, in the left navigation pane, select **Clients**.
 2. On the **Client** window, right-click the client resource for the NMC database backup and select **Modify Client Properties**.
 3. On the **General** tab, modify the **Save set** field and specify the path to the `nmcdb_stage` directory on a file system that has sufficient disk space.
-  **Note:** If the path does not exist, the `savepsm` process creates the directory at the time of the backup.
4. Click **OK**.

Performing a manual backup of the NMC server database

Use the `savepsm` command to perform a manual backup of the NMC server database.

About this task

UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `savepsm` command.

Procedure

1. For Linux hosts, if you did not install NMC server software in the default path `/opt/lgtomc`, then add the `NMC_install_dir/bin` directory to the `LD_LIBRARY_PATH` environment variable.
2. From a command prompt, use the `savepsm` command to backup the NMC database

```
savepsm staging_directory
```

where *staging_directory* is the location that the backup uses to temporarily store a copy of the NMC database for backup.

For example, on windows, type:

```
savepsm e:\nmcdb_stage
```

Using the NMC Configuration Wizard

You can use the **NMC Configuration** wizard to create the account that the NMC server service account in the NetWorker Authentication Service local database,

specify which NetWorker server will back up the NMC database, and add NetWorker servers to the Enterprise.

Before you begin

Connect to the NMC server with a user that has the as Console Application Administrator role.

Procedure

1. From the **NMC** GUI, click **Setup**.
2. From the **Setup** menu, select **Configuration Wizard**.

NMC server authentication

When you use a web browser on a host (NMC client) to connect to the NMC server, the `http` daemon on the NMC server downloads the Java client to the NMC client.

The NMC server relies on the NetWorker Authentication Service to manage and validate users. When you log in to the NMC server, the NMC server contacts the NetWorker Authentication Service on the host that you specified during the NMC installation process to verify the credentials of the user account. When the NetWorker Authentication Service successfully verifies the user, the application issues a time-based, signed, and encrypted SAML token to the requesting process. All the operations that require authentication can use the token to verify the user, until the token expires. The NetWorker Authentication Service maintains a local user database for authentication. NetWorker Authentication Service also supports the use external authentication authorities for authentication. For example, Lightweight Directory Access Protocol (LDAP), Lightweight Directory Access Protocol over SSL (LDAPS), and Microsoft Active Directory server (AD). You can configure the NMC server and the managed NetWorker servers to use LDAP, AD, or the NetWorker Authentication Service local user database to provide user authentication and authorization.

The *NetWorker Security Configuration Guide* describes how to perform the following tasks:

- Manage the NetWorker Authentication Service.
- Configure user authentication on the NMC.
- Configure user authorization to the NMC and NetWorker servers.

Configuring the NMC server to manage additional NetWorker servers

The NMC Server can use only one NetWorker Authentication Service to provide authentication services. When the NMC Server manages more than one NetWorker Server, configure a trust between each NetWorker Server that the NMC Server will manage and NetWorker Server that will provide authentications services to the NMC Server. After you establish each trust, update the user groups on each NetWorker Server to include the users and groups that require access to the NetWorker Server.

Procedure

1. To establish the trust, type the following command on each NetWorker Server that is not local to the NetWorker Authentication Service that NMC uses for authentication:

```
nsrauthtrust -H Authentication_service_host -P
Authentication_service_port_number
```

where:

- The location of the `nsrauthtrust` command differs on Linux and Windows:
 - Linux—`/usr/sbin`
 - Windows—`C:\Program Files\EMC NetWorker\nsr`
- *Authentication_service_host* is the hostname of the NetWorker Server that authenticates the NMC Server host.
- *Authentication_service_port_number* is the port number used by the NetWorker Authentication Service. The default port number is 9090.

For example:

```
nsrauthtrust -H nwserver.corp.com -P 9090
```

2. Grant the NetWorker Authentication Service user groups access to the NetWorker Server, by typing the `nsraddadmin` command:

```
nsraddadmin -H Authentication_service_host -P
Authentication_service_port_number
```

For example:

```
nsraddadmin -H nwserver.corp.com -P 9090
```

The `nsraddadmin` command updates the following user groups:

- Application Administrator—Adds the distinguished name (DN) of the NetWorker Authentication Service Administrators group.
- Security Administrator—Adds the DN of the NetWorker Authentication Service Administrators group.
- Users—Adds the DN of the NetWorker Authentication Service Users group.

After you finish

Add additional users and groups to user groups on each NetWorker server. [Modifying user groups for new NetWorker Authentication Service users](#) provides more information.

Changing the NetWorker Authentication Service hostname and port number

When you install the NMC server software, you specified the hostname of the NetWorker Authentication Service and the port number that the service uses for communication. Perform the following steps to change the host that provides user authentication to the NMC server.

Procedure

1. Connect to the NMC server with an Administrator account on Windows or the root account on UNIX.
2. Stop the EMC `gstd` process:
 - Linux—`/etc/init.d/gstd stop`
 - Windows—Stop the **EMC GST Database Service** service.
3. From a command prompt, to change the NetWorker Authentication Service host that is used by the NMC server, type the `gstauthcfg` command.

The location of the `gstauthcfg` command is not in the path by default and differs on Linux and Windows:

- **Linux**—/opt/lgtonmc/bin
- **Windows**—C:\Program Files\EMC NetWorker\Management\GST\bin

For example:

```
gstauthcfg -c -t -h New_authentication_service_hostname -p port_number
```

 **Note:** The default port number is 9090.

4. Start the EMC `gstd` process:

- **Linux:** /etc/init.d/gstd start
- **Windows:** Start the **EMC GST Database Service** service.

5. To establish the trust, type the following command on each NetWorker Server that is not local to the NetWorker Authentication Service that NMC uses for authentication:

```
nsrauthtrust -H Authentication_service_host -P Authentication_service_port_number
```

where:

- The location of the `nsrauthtrust` command differs on Linux and Windows:
 - **Linux**—/usr/sbin
 - **Windows**—C:\Program Files\EMC NetWorker\nsr
- *Authentication_service_host* is the hostname of the NetWorker Server that authenticates the NMC Server host.
- *Authentication_service_port_number* is the port number used by the NetWorker Authentication Service. The default port number is 9090.

For example:

```
nsrauthtrust -H nwserver.corp.com -P 9090
```

6. Grant the NetWorker Authentication Service user groups access to the NetWorker Server, by typing the `nsraddadmin` command:

```
nsraddadmin -H Authentication_service_host -P Authentication_service_port_number
```

For example:

```
nsraddadmin -H nwserver.corp.com -P 9090
```

The `nsraddadmin` command updates the following user groups:

- **Application Administrator**—Adds the distinguished name (DN) of the NetWorker Authentication Service Administrators group.
 - **Security Administrator**—Adds the DN of the NetWorker Authentication Service Administrators group.
 - **Users**—Adds the DN of the NetWorker Authentication Service Users group.
7. Connect to the NMC server GUI with a user that has the NMC Console Security Administrator role.

8. When prompted to create a service account for the NMC server in the NetWorker Authentication Service database, click **OK**.

 **Note:** If you do not create the service account, the NMC server cannot monitor events or gather reporting data from the managed NetWorker servers.

Modifying user groups for new NetWorker Authentication Service users

Use NMC to add NetWorker Authentication Service users and groups to user groups on a NetWorker server. If you configured the NetWorker Authentication Service to use external LDAP or AD authorities, use NMC to add LDAP or AD users and groups to User Groups on a NetWorker server.

The *NetWorker Security Configuration Guide* provides more information about user groups and how to configure user authorization on a NetWorker server.

Modifying NetWorker user group membership for NMC

Use the External roles field in the User Group resource to manage local database, LDAP, and AD user and group access to the NetWorker server.

Before you begin

Use NMC to connect to the NetWorker server with a user that is a member of the Security Administrators user group on the NetWorker server.

Procedure

1. On the **Administration** window, click **Server**.
2. Click **User Groups**.
3. Right-click the user group, and then select **Properties**.
4. Modify the **External roles** attribute. To add NetWorker Authentication Service local database users or groups, click the + sign, and then select the users or groups. When you add an LDAP or AD user or group, specify the distinguished name (DN).


The following sections provide more information about how to get the dn for the user or group in an AD or LDAP external authentication authority, and how to add the NMC service account.

 **Note:** It is recommended that you specify usernames when your user accounts are a member of a large number of groups.

Example: Adding AD group to the External roles attribute

The following example uses ADSI Edit, a Windows tool that allows you to view information about users and groups in AD directory service. [Microsoft TechNet](#) provides the most up to date information about how to use ADSI Edit.

1. To connect to the AD directory, use ADSI Edit.
2. Navigate to the AD group, right-click the group name, and then select **Properties**.
3. On the **Attribute Editor** window, select **distinguishedName** from the attribute list, and then select **View**.
4. On the **String Attribute Editor** window, with the entire dn highlighted, right-click in the value field, and then select **Copy**. The following figure provides an example of copying the group DN in the ADSI Editor.


Figure 83 Copying the group DN

5. Click **Cancel**, and then close ADSI Editor.
6. Paste the dn value for the group into the **External roles** attribute.

Example: Adding LDAP group to the External Roles attribute

The following example uses LDAP Admin, a third party tool that allows you to view information about users and groups in the LDAP directory service.

1. To connect to the LDAP server, use LDAP Admin.
2. Navigate to the LDAP group, right-click on the group name, and then select **Copy dn to clipboard**. The following figure provides an example of the LDAP Admin window.

Figure 84 Copying the group DN

3. Close the LDAP Admin window.
4. Paste the dn value for the group into the **External roles** attribute.

```
authc_mgmt -u administrator -p "Password1" -e query-ldap-users -D
"query-tenant=IDD" -D
"query-domain=ldapdomain"
```

Adding the NMC service account to the Users user group

When the NMC Server manages multiple NetWorker Servers, the `nsraddadmin -H` command automatically adds a NetWorker Authentication Service group called "Users" to the "Users" user group on each remote NetWorker Server. The NetWorker Authentication Service Users group contains the NMC service account. To monitor operations on a NetWorker Server that is remote to the NMC Server, the NMC service account requires Monitor NetWorker privileges. If the NetWorker "Users" user group does not specify a NetWorker Authentication Service group that contains the NMC service account, NMC cannot monitor remote NetWorker Server operations.

About this task

To add the NMC service account to the "Users" user group on a NetWorker Server, perform the following steps.

Procedure

1. Connect to the NMC server with the NetWorker Authentication Service administrator account.
2. Click **Enterprise**.
3. Right-click the NetWorker Server and select **Launch Application**.



Note: Perform this step and each subsequent step on for each NetWorker Server that is not local to the authentication service that the NMC Server uses to authenticate users.


4. On the **NetWorker Administration** window, select **Servers**.
5. In the left navigation pane, select **User Groups**.
6. Right-click the **Users** user group, and then select **Properties**.
7. Click the plus sign (+) beside the **External roles** attribute.

The **Add Distinguished Names** window appears.

8. In the **Authentication Service Hostname** field, specify the name of the host that the NMC Server uses to authenticate users, and then click **Change**.
9. In the user table, select the service account for the NMC Server and click **OK**.

The following figure provides an example of the Add Distinguished Names window with the service account selected.

Figure 85 Add Distinguished Names window



Add Distinguished Names window

The service account appears in the **External roles** attribute.

10. Click **OK**.

Enabling two factor authentication for AD and LDAP users

You can use NMC to enable two factor authentication for AD and LDAP users.

Procedure

1. On Linux, modify the `/opt/nsr/admin/gst_linux.sh` script and add the following at the beginning of the script:

```
GST_LDAP_USING_2FA="true"
export GST_LDAP_USING_2FA
```

- On Windows, do the following to set the environment variable:
 - Browse to **Control Panel > System and Security > System > Advanced Settings**.
 - On the **General** tab, click **Environment Variables**.
 - In the **System variables** section, click **New**.
 - In the **Variable** name field, type: `GST_LDAP_USING_2FA`, and in the **Variable** value field, type: `True`. Click **OK**.
- Log in to NMC as an Administrator.
- From the **NMCConsole** window, click **Setup**.
- In the left pane, select **Users**.
 - Right-click the service account (for example, `svc_nmc_*`), and then select **Properties**.
 - On the **Groups** window, select **Administrators** and click **OK** to add the service account user as a part of this group.
- Restart the EMC GST service:
 - Linux: `/etc/init.d/gstd start`
 - Windows: Start the **EMC GST** service

7. Configure AD and LDAP users. The *NetWorker Security Configuration Guide* provides information on configuring AD and LDAP users.
8. Connect to the NMC server with an LDAP or AD user.

Moving the NMC Server

You can move an NMC Server from one host to another only if both hosts use the same operating system.

Before you begin

- Perform a level full backup of the NMC database on the source NMC Server. [Performing a manual backup of the NMC server database](#) on page 768 provides more details.
- On the target NMC host, install the NetWorker and NMC Server software. When prompted to specify the NetWorker Authentication Service host, specify the same NetWorker Authentication Service host as the source NMC Server. The *NetWorker Installation Guide* provides more information.
- If you use a License Manager server, then install and configure the License Manager software first. If you use the License Manager software and the License Manager server moves to a new host, then specify the new License Manager hostname in the **Console** window.

Procedure

1. Connect to the NMC GUI on the target NMC Server.
2. In the NMC GUI, connect to the NetWorker Server that performed the NMC database backup.
3. On the **Administration** window, select **Protection**.
4. In the left navigation pane, select **Clients**.
5. Create a Client resource for the target NMC host. [Create a Client resource with the Client Properties dialog box](#) on page 484 describes how to create a Client resource.
6. Edit the Client resource for the source NMC Server. On the **Globals (2 of 2)** tab in the **Remote Access** attribute specify the administrator account of the target NMC Server.

For example, `administrator@target_NMC_server`
where `target_NMC_server` is the hostname of the target NMC Server.

7. Stop the NMC Server service on the source NMC Server.
8. Stop the NMC Server service on the target NMC Server.
9. For Linux hosts, if you did not install NMC Server software in the default path `/opt/lgtomc`, then add the `NMC_install_dir/bin` directory to the `LD_LIBRARY_PATH` environment variable.
10. Run the `recoverpsm` command on the recovery host:

```
recoverpsm -f -s NetWorker_server -c source_NMC_server -p
AES_Passphrase staging_dir
```

where:

- `NetWorker_server` is the name of the NetWorker Server.

- *source_NMC_server* is the name of the source NMC Server.
- *AES_Passphrase* is the passphrase that was specified for the NMC database backup.
- *staging_dir* is the staging directory specified during the backup of the database on the source NMC Server.

Note: When you perform an NMC database backup, the backup operation performs a backup of the database from the staging directory. As a result, the save set name is name of the staging directory. Use the `mminfo` command on the NetWorker Server to determine the name of the staging directory.

The *NetWorker Command Reference Guide* or the UNIX man pages provide a complete description of the `recoverpsm` command line options.

11. If the source NMC Server managed NetWorker 8.2.x and earlier servers that use LDAP authentication, then recover the LDAP configuration authority files. Use the `recover` command, the NetWorker User program, or the NMC **Recovery** wizard to recover all the files in the `console_install_dir/cst` directory. Recover these files to the `console_install_dir/cst` directory on the target NMC Server.
12. Start the NMC Server service on the target NMC Server and connect to the NMC GUI.

After you finish

If the target NMC Server uses a different NetWorker Server to provide authentication services than the NetWorker Server that the source NMC Server used, then you must use the `gstauthcfg` command on the NMC Server to update the NetWorker Authentication Service host, and then run the `nsrauthtrust` commands on each NetWorker Server that is managed by the NMC Server.

When the source NMC Server uses a different NetWorker Server for authentication and you do not establish a trust, the following behavior occurs:

- The NMC **Events** window displays `Unable to connect to the server` error messages for each managed NetWorker Server.
- When you try to connect to the NetWorker Server, a message similar to the following appears: `Unable to connect to the server: Unable to set user privileges based on user token for SYSTEM: Failed to validate security token.`

[Changing the Authentication service hostname and port number](#) provides more information.

Migrating NMC users to the authentication service database

If you did not migrate the NMC users to the authentication service database when the login process prompted you to during the login process after the NMC server after an update, you can perform the migration later.

Before you begin

Log in to the NMC server as a Console Security Administrator. The NetWorker Authentication Service administrator account is a Console Security Administrator.

Procedure

1. Click **Setup**.
2. From the **Setup** menu, select **Migrate Users**.
3. On the **Migrate Users** page, select the users that you want to migrate.

 **Note:** By default all users are selected for migration. The migration deletes unselected user accounts.

4. For each user, perform the following steps:
 - a. In the **Password** field, specify an initial password.
 Ensure the password complies with the following minimum requirements:
 - Nine characters long
 - One uppercase letter
 - One lowercase letter
 - One special character
 - One numeric character
 - b. Leave the default selection for **Password Change Required**, which ensures that when the user connects to the NMC Server for the first time, that the log in process prompts the user to change their password.
 - c. In the **Groups** field, if the user will manage user accounts, select the Administrators group.

Updating the NetWorker User Group resources for migrated NMC users

The NetWorker server uses the membership in the External Roles field of the user group resources to determine the privileges that are assigned to the NetWorker Authentication Service local database users. After the log in process migrates NMC users into the NetWorker Authentication Service local database, update the User Group resources on each managed NetWorker server, to provide the migrated NMC users with the privileges to each NetWorker server.

About this task

Perform the following steps while logged in to the NMC server with the Administrator account.

Procedure

1. In the NMC GUI, create an NMC group that contains the local database users. This group allows you to quickly add multiple users that require the same privileges to one or more user groups:
 - a. On the NMC GUI, click **Setup**.
 - b. On the **User and Roles** navigation pane, right-click **Groups** and select **New**.
 - c. In the **Name** field, specify a unique name for the group.
 In the **Local Users** section, select all the user accounts to add to this group, and then click **OK**.
2. In the **Administration** window, perform the following steps:
 - a. On the toolbar, select **Server**.

- b. On the left navigation pane, expand **User Groups**.
- c. Right-click the user group to which the NMC users require membership, and select **Properties**.
- d. In the **Configuration** section, click the **Add (+)** button beside the **External Roles** attribute.
- e. Select each local database user or group that requires the privileges that are assigned to the user group, and then click **OK**.

To select multiple successive users or groups, hold the **Ctrl** key while you select the first and last user or group. To select multiple individual users or groups in any order, hold the **Shift** key while you select each user or group.

For more information on External Roles and User Groups, see *NetWorker Security Configuration Guide*

Results

The distinguished name (dn) for each selected user and group appears in the **External Roles** field.

Resetting the administrator password

To reset the administrator password, create a JSON file on the NetWorker server that contains the new password in a Base64 encoded format.

Procedure

1. To determine the Base64 password value for the new password, use Base64 encoding utilities:
 - On Windows, perform the following steps:
 - a. Create a text file and specify the password value in clear text, on one line.
For example, create a password file that is called *mypassword_in.txt* with the password value *"1.Password"*.
 - b. To create a Base64 encoded password for the password value that is defined in the *mypassword_in.txt* file, use the *certutil.exe* utility.
For example:

```
certutil.exe -encode mypassword_in.txt mypassword_out.txt
```

where *mypassword_out.txt* is the name of the output file that contains the Base64 encoded password.

Output similar to the following appears:

```
Input Length = 10
Output Length = 74
CertUtil: -encode command completed successfully.
```

The contents of the *mypassword_out.txt* file contains the following encoded text for the password value *"1.Password"*:

```
-----BEGIN CERTIFICATE-----
MS5QYXNzd29yZA==
-----END CERTIFICATE-----
```

where the Base64 encoded password is *MS5QYXNzd29yZA==*.

- To create the Base64 encoded password on Linux, use the `base64` utility. For example, to create the Base64 encoded password for a password value of `"1.Password"`, type:

```
echo -n "1.Password" | base64
```

The command displays the encoded text for the password value

```
"1.Password": MS5QYXNzd29yZA==
```

2. Use a text editor to open the `authc-local-config.json.template` file, which is located in the `C:\Program Files\EMC NetWorker\nsr\authc-server\scripts` folder on Windows and the `/opt/nsr/authc-server/scripts` directory on Linux.
3. In the template file, perform the following steps:
 - a. Replace the `your_username` variable with the name of the administrator account for which you want to reset the password.
 - b. Replace the `your_encoded_password` variable with the base64 encoded password value.

For example, to reset the password for the user account administrator with a password of `"1.Password"`, the modified file appears as follows:

```
{
  "local_users": [
    {
      "user name": "administrator",
      "password": "MS5QYXNzd29yZA=="
    }
  ]
}
```

4. Rename the `authc-local-config.json.template` file to `authc-local-config.json`.
5. Copy the `authc-local-config.json` file to the Tomcat `conf` folder.

By default, the `conf` folder is `/nsr/authc/conf` on Linux and `C:\Program Files\EMC NetWorker\authc-server\tomcat\conf` on Windows.

6. Change privileges on the `authc-local-config.json` file:

```
chmod 755 /nsr/authc/conf/authc-local-config.json
```

If you do not change the privileges, the `authc-server.log` displays an error indicating that you do not have the necessary permissions to open the file.

7. On the NetWorker server, stop, and then start the services:

- For Windows, type the following commands from a command prompt:

```
net stop nsrexecd
```

```
net start nsrd
```

 **Note:** If the NetWorker server is also the NMC server, start the NMC server service. Type the following commands: `net start gstd`

- For Linux, type the following commands:

```
/etc/init.d/networker stop
```

```
/etc/init.d/networker start
```

When the NetWorker Authentication Service starts, the startup process checks for the `authc-local-config.json`. If the file exists and the password adheres to the minimum password policy requirements defined for a password, the NetWorker Authentication Service resets the password. Review the `authc-server.log` file for errors.

By default, the `authc-server.log` file is located in `/nsr/authc/logs` on Linux and `C:\Program Files\EMC NetWorker\authc\tomcat\logs` on Windows.

 **Note:** The startup process automatically deletes the `authc-local-config.json` file to ensure that the password is not reset the next time that you restart the NetWorker Authentication Service.

8. To confirm that you can connect to the NetWorker Authentication Service with the new password, use the `authc_mgmt` command.

For example:

```
authc_mgmt -u administrator -p "1.Password" -e find-all-users
```

The query returns 2 records.

```
User Id User Name
1000 administrator
1001 svc_nmc_bu-iddnwserver2
```

Changing the service port used by the NMC database

The installation process prompts you to specify the NMC database port. By default, the NetWorker Management Console database uses port 5432 for TCP/IP communications. You can change the port after the installation process completes.

Changing the service port used by the NMC database on Linux

About this task

Perform the following steps to change the service port that is used by NMC.

Procedure

1. Stop the NMC daemons, by typing the command below, based on the initialization system running on your Linux machine:
 - `sysvinit—/etc/init.d/gst stop`
 - `systemd—systemctl stop gst`

2. Edit the `/opt/lgtonmc/etc/gstd.conf` file to add or change the following line:

```
db_svc_port=port_number
```

For example:

```
db_svc_port=2639
```

3. Run the `/opt/lgtonmc/bin/gstconfig` command to update the port value in the NetWorker NMC server configuration file.
4. Edit the `postgresql.conf` file to add or change the following line:

```
port=port_number
```

For example:

`port=2639`

Note: By default the `postgresql.conf` file is located in the `/nsr/nmc/nmcdb/pgdata` directory.

5. Close the terminal or command prompt window.
6. Start the NMC daemons, by typing the command below, based on the initialization system running on your Linux machine:
 - `sysvinit—/etc/init.d/gst start`
 - `systemd—systemctl start gst`

This action also starts the `postgres` and `httpd` processes.

NOTICE If `/etc/init.d/gst` file is missing for `sysvinit` systems or `gst` file is not enabled for `systemd` systems, run the script - `/opt/lgtonmc/bin/nmc_config`

Multiple `Postgres` processes appear. Two or more `httpd` processes appear. By default, these `httpd` processes run as `nsrnmcc`.

7. Confirm that the daemons have started, by typing the following command: `ps -ef | grep lgtonmc`.

Output similar to the following appears when the daemons have started:

```
nsrnmcc 7190 1 0 Nov23 ? 00:00:06 /opt/lgtonmc/bin/gstd
nsrnmcc 7196 1 0 Nov23 ? 00:00:00 /opt/lgtonmc/apache/bin/
httpd -f /opt/lgtonmc/apache/conf/httpd.conf
nsrnmcc 7197 7196 0 Nov23 ? 00:00:00 /opt/lgtonmc/
apache/bin/httpd -f /opt/lgtonmc/apache/conf/httpd.conf
nsrnmcc 7212 1 0 Nov23 ? 00:00:00 /opt/lgtonmc/
postgres/bin/postgres -D /nsr/nmc/nmcdb/pgdata
root 18176 18141 0 02:47 pts/0 00:00:00 grep lgtonmc
```

Changing the service port used by the NMC database on Windows

About this task

Perform the following steps to change the service port that is used by NMC.

Procedure

1. Stop the **EMC GSTD Service** service.
2. Edit the `gstd.conf` file to add or change the following line:

`db_svc_port=port_number`

For example:

`db_svc_port=2639`

Note: By default the `gstd.conf` file is located in the `C:\Program Files\EMC NetWorker\Management\GST\etc` directory.

3. Edit the `postgresql.conf` file to add or change the following line:

`port=port_number`

For example:

`port=2639`

 **Note:** By default the `postgresql.conf` file is located in the `C:\Program Files\EMC NetWorker\Management\nmcd\pgdata` directory.

4. Use the `regedit` command to update the port number in the registry.
 - a. Browse to `\\HKEY_LOCAL_MACHINE\SOFTWARE\ODBC\ODBC.INI\lgto_gst_pgsql`.
 - b. Edit the **Port** registry key.
 - c. In the **Value Data** field, specify the new port number.
 - d. Click **OK**.
5. Start the EMC GST Service.

Changing database connection credentials

When the NMC server starts for the first time, it automatically generates the login credentials that are used to log in to the NetWorker Console database. The NMC server stores this information internally and the user does not need to know the required credentials. However, it may be necessary to force the NMC server to change the database connection credentials.

Procedure

1. Stop the GST Service.
2. Set the environment variable `GST_RESET_DBPWD` to any value.
For Windows system, set this value as a System Variable, then restart the system after you set the variable.
3. Restart the GST Service.
4. Delete the `GST_RESET_DBPWD` environment variable. On Windows system, restart the computer after you delete the variable.

Updating the NMC server IP address/hostname

If you modify the IP address or hostname of the NMC server or if you add or remove protocols such as IPv6, you must update the NMC server configuration.

About this task

Perform the following steps with the root account on Linux hosts or the Administrator account on Windows hosts.

Procedure

1. Stop the `gstd` service:
 - On Linux:
 - `sysvinit—/etc/init.d/gst stop`
 - `systemd—systemctl stop gst`
 - On Windows: Stop the **EMC GSTD Service** service.

2. Edit the `gstd.conf` file and update the IP address that is defined for the line *string authssvc_hostname*.
3. Browse to the NetWorker `bin` directory then run the platform-specific commands:
 - On Windows, run `gstconfig` in the `NMC_install_dir\GST\bin` folder.
 - On Linux, as root, run the `gstconfig` command in the `/opt/lgtomc/bin` directory.
4. Start the `gstd` service:
 - On Linux:
 - `sysvinit—/etc/init.d/gst start`
 - `systemd—systemctl start gst`
 - On Windows: Start the **EMC GST Service** service.
5. For NMC server hostname changes only, delete the Client resource that you created to perform NMC server database backups, then create a new client resource.

Setting system options to improve NMC server performance

The NMC server includes several options that enable users to fine-tune the performance of the NMC server.

About this task

To set system options, log in to the NMC server as a Console administrator.

Procedure

1. From the **Console** window, click **Setup**.
2. From the **Setup** menu, select **System Options**.
3. Set a value, or enable or disable the appropriate system option. The following table provides a description of the available system options.

 **NOTICE** Do not adjust these system options without careful consideration. A mistake in setting system options can seriously degrade performance.

Table 135 NMC server system options

System option	Description
Log-on banner	Default Value: Warning: Authorized user only Defines the log-on banner displayed in the NMC server login window.
Debug level	Default value: 0 Range: 1-20 Defines the level of debug information to log in the <code>gstd.raw</code> file. Increase this value to troubleshoot only.

Table 135 NMC server system options (continued)

System option	Description
Polling interval for events and reporting (seconds)	Default value: 20 Range: 2-unlimited Defines how frequently the NMC server contacts the managed NetWorker servers for event and report updates.
Polling interval for NetWorker activities (seconds)	Default value: 10 Range: 2-unlimited Defines the frequency in which the NMC server contacts the managed NetWorker servers for activity updates.
Polling thread factor	Default value: 5 Range: 0-10 Defines how many server threads to create when polling the NetWorker server for NetWorker activities, events, and reporting. The higher the number the higher the number of threads created. It is not a one-to-one relationship.
Maximum number of log messages	Default value: 32 Range: 32-512 Defines the number of log messages that display in the Console Log window.
NetWorker user auditing	Default value: enabled When enabled, the NMC server collects auditing information. For example, NetWorker server configuration changes performed from the Console GUI. The NMC server database stores the auditing information. To view audit information browse to Reports > Users > User Audit Report. When disabled, the NMC server does not collect auditing information.
User authentication for NetWorker	Default value: enabled Defines how the Console user accesses a managed NetWorker server. - When enabled, the Console username determines the Console user access. Individual User Authentication on page 786 provides detailed information. - When disabled, the user id of the gstd process owner determines the Console user access.
RPC ping via UDP when connecting to NetWorker	Default value: disabled Before the NMC server connects to a managed NetWorker server, the NMC server confirms that the NetWorker server daemons are running.

Table 135 NMC server system options (continued)

System option	Description
	- When enabled, the NMC server uses the UDP protocol to confirm that the NetWorker server is up and running. - When disabled, the NMC server uses the TCP protocol to confirm that the NetWorker server is up and running.

4. Click **OK**.

Individual User Authentication

Console security administrators restrict or grant Console user access to NetWorker servers based on the Console username when you enable the User Authentication for NetWorker system option, after a subsequent restart of the NMC server service. The NMC server software enables this system option by default.

Requests to NetWorker servers through the Administration window always come from the NMC server, regardless of any system option settings.

When you enable the User Authentication for NetWorker system option:

- Access requests to a NetWorker server appear to be coming from users on the NMC server, rather than from the `gstd` process owner on the NMC server.
- A NetWorker 8.2.x and earlier server allows requests only from users who belong to the Administrators list of the NetWorker server. You must include the username of the Console daemon process owner in the NetWorker Administrators list on NetWorker 8.2.x and earlier servers to which the Console users have access. The *NetWorker Installation Guide* describes how to add the Console daemon process owner to the NetWorker Administrators list by using the `nsraddadmin` command.

 **NOTICE** You must specify the username of the root or system user on the NMC server, regardless of whether you use individual user authentication.

Impact on network connections

When you enable individual user authentication, the NMC Server software might require more network connections. Additional network connections might firewall port requirements. The *NetWorker Security Configuration Guide* provides information about firewalls.

When you set the User Authentication for NetWorker system option, the NetWorker Authentication Service software creates a separate network connection the NetWorker Authentication Service to a NetWorker Server for each Console user that has an Administration window open to that server.

When you do not set the user authentication for NetWorker system option, there is only one network connection from the NetWorker Authentication Service to the managed NetWorker Server.

Displaying international fonts in non-US locale environments

To use or view data from a localized NetWorker Server, ensure that the appropriate font is available to the NMC Server.

The *NetWorker Installation Guide* describes how to display international fonts on an NMC Server that operates in English mode.

NetWorker License Manager

The NetWorker License Manager (LLM) software provides centralized license management, which enables you to maintain all licenses in the Enterprise from a single host if using the traditional licensing model.

Note: NetWorker 19.1 requires the use of the Dell EMC Licensing Solution, which deploys an Dell EMC Licensing Server. You do not require the NetWorker License Manager and it is recommended that you skip the NetWorker License Manager software installation during the NetWorker 19.1 install. When upgrading to NetWorker 19.1, you can back up the NetWorker License Manager by following the procedure outlined in the section "Backing up the NetWorker License Manager" in the *NetWorker Licensing Guide*.

With the NetWorker License Manager, you can move NetWorker software from one host to another, or change the IP address on an existing NetWorker Server without having to reauthorize the software. You can install the NetWorker License Manager program as an option during the NetWorker software installation.

The latest *NetWorker License Manager server Installation and Administration Guide* provides more information on how to install and use the NetWorker License Manager.

Entering an enabler code

Procedure

1. From the **Console** window, click **Setup**.
2. Right-click **Licensing**, then select **New**. The **Create** dialog box appears.
3. In the **Enabler Code** attribute, type the enabler code and leave the other attributes blank.
4. Click **OK**.

Deleting an enabler code

Procedure

1. From the **Console** window, click **Setup** and then click **Licensing**.
2. Right-click the license to delete, then select **Delete**.
3. Click **Yes** to confirm the deletion.

Entering an authorization code

Procedure

1. Log in as a Console Application Administrator.
2. From the **Console** window, click **Setup** and then click **Licensing**.
3. Right-click the license to be authorized, then select **Properties**. The **Properties** dialog box appears.
4. In the **Auth Code** attribute, enter the authorization code for the product (the authorization code assigned to the specified permanent enabler or update enabler code).
5. Click **OK**. The license is now permanently enabled.

Changing the License Manager server

You can change the License Manager server that manages NetWorker Console licenses at any time.

About this task

Procedure

1. Log in as a Console Application Administrator.
2. From the **Console** window, click **Setup**.
3. Right-click **Licensing**, then select **Change LLM Server**. The **Change LLM Server** dialog box appears.
4. In the **LLM Server** attribute, type the hostname of the appropriate server and click **OK**.

NMC error messages and corrective actions

The following table provides a list of NMC error messages or symptoms and corrective actions to take.

Table 136 Error messages or symptoms

Error message or symptom	Possible cause	Corrective action
If the Console server fails to load and instead displays a Save As... dialog box.	JavaScript is not enabled on the host. The security level in Internet Explorer is set to High, which disables JavaScript, which is needed to launch the product, or JavaScript has been disabled by some other means.	In Internet Explorer, ensure that the security level is lower than high, which disables JavaScript, or enable Active Scripting.
The NetWorker Server does not accept the authorization code.	A temporary enabler code has already expired.	Log out, then stop and restart the NMC Server services.

Table 136 Error messages or symptoms (continued)

Error message or symptom	Possible cause	Corrective action
An application window is unresponsive.	Insufficient disk space on the file system where the NMC database is installed.	- Ensure that the NMC Server is running. If it is not, close all application windows and check the <code>gstd</code> log file for errors. - Back up and move the Console database, if required. - On a Windows system, run InstallShield with the Repair option to move the database to a different drive.
	Application ran out of memory.	Close all instances of the application and restart it.
	Another dialog box is open in the NMC window or Administration window.	Close any open dialog boxes or error messages.
Connection refused: no further information. or Problem contacting server <code>server_name</code> :	NMC Server is in the process of crashing or has already crashed.	Check to see if the NMC Server is running. - If it is running, stop and restart the NMC server. - If it is not, close all application windows and check the <code>gstd</code> log file for errors.
	Console server has been started within the previous few minutes.	Wait a couple of minutes and retry.
Failed to bind to port <code>port_number</code> message appears in the <code>gstd.raw</code> log file.	Another process is using the <code>gstd</code> service port (default 9001) or the port is in a timeout (TIME_WAIT/FIN_WAIT) state.	Close any running NMC GUIs or any processes that may be using the <code>gstd</code> service port. Wait until the timeout period passes so that the operating system can free up the port. The timeout period may differ between operating systems.
Database fetch operation failed messages appears in the <code>gstd.raw</code> log file.	The NMC database is corrupt.	Recover the database.
Display problem: In Internet Explorer:	The <code>gstd</code> service is not running on the NMC server.	Restart the NMC server.

Table 136 Error messages or symptoms (continued)

Error message or symptom	Possible cause	Corrective action
The page cannot be displayed.	Browser is not pointing to the correct URL.	Check the install log file to determine the HTTP port that is used by the NMC Server.
	Network connection is down.	Ping the NMC Server to confirm the network connection. If it is available, contact the system administrator.
Enabler code not accepted.	Temporary enabler code has expired.	Close the NMC Server and log in again. Repeat the procedure of typing the enabler code. If the enabler code is still not accepted, log out, then stop and restart the NMC Server.
Database delete operation failed: Reference object does not exist.	Another user has already deleted that user or folder.	None
Database store operation failed: An object with pathname <i>"pathname"</i> already exists.	- Another user is trying to add a folder to the same location in the Enterprise simultaneously. - An object was added with the same name as an existing object.	- Wait a few moments and try again. - Check whether there is an existing object with the same name.
Invalid Object ID.	Another user deleted that host.	None
Could not contact License Manager on <i>hostname</i> . - or - Program not registered.	License Manager hostname has not been assigned or License Manager is not running or installed.	If you are using the License Manager and a hostname has not been assigned: Select the Software Administration task. Click Licensing. Click Software Administration on the menu bar. Click Change LLM Server. Type the new License Manager hostname. Click OK.

Table 136 Error messages or symptoms (continued)

Error message or symptom	Possible cause	Corrective action
		If License Manager is installed, but not running, start it. The <i>NetWorker License Manager server Installation and Administration Guide</i> provides details.
	NetWorker Client was stopped, but the License Manager was not stopped, and then the NetWorker Client was restarted. Although both services are now running, NetWorker Client must be started before License Manager is started. If the services are not started in the correct order, an error condition occurs.	Stop the NetWorker software. Stop License Manager, if it is running. Restart License Manager. Restart the NetWorker software.
License allocation failed.	Temporary license for NetWorker software is expired.	Enter enabler codes and register the product.
License managed event indicates that license is expiring/expired even though it has been authorized.	License has been authorized within the last 24 hours.	None needed. To remove the managed event from the display, dismiss the event or it is deleted within 24 hours.
Logging of troubleshoot messages has stopped. alloc /opt: File system full.	Disk space on the /opt file system is nearly full.	Allocate more disk space.
Event disappears from the Events window.	Another user dismissed it, or the problem that was causing the event no longer exists.	None
Dialog box: "Java Web Start – Download Error" with the message, "Unable to launch NetWorker Console".	Java Web Start preferences are set to something that is incompatible with the rest of the environment. (For example, a proxy server has been set up that stops Java Web Start from downloading the Console client software	Check the Preference settings in the Java Web Start Application Manager for compatibility with the environment. Change any settings that prohibit the download of the Console client software.

Table 136 Error messages or symptoms (continued)

Error message or symptom	Possible cause	Corrective action
	from the Console web server.) This error message may also occur if the Console is being launched on a localized operating system and the Java Web Start cache path contains non-English characters.	In the proxy server example, go to the General tab of the Preferences dialog box and select None, for Proxies. If the Java Web Start cache path contains non-English characters, change the path to contain no non-English characters.
gstd.log file error: internal error: could not end transaction	When you move the system time ahead, the NMC Server starts a time out event and closes database client connection for the gstd process.	None
GC overhead limit exceeds	This error message appears when you are performing tasks in NMC and there is insufficient Java memory, or heap.	Increase the Java heap size to 1400MB. 1. Start the Java Control Panel application: <code>javaws -viewer</code>. 2. Close the Java Cache window. 3. On the Java tab, click View. 4. On the Java Runtime Environment Settings window, double-click in the Runtime Parameters cell for the Java version that you use with NMC. 5. In the Runtime Parameters field, specify a heap size of 1400 MB: <code>-Xms1400m</code> 6. Click OK. 7. Click OK to close the Java Control Panel. 8. Close the NMC and NetWorker Administration windows and reconnect to the NetWorker.

Console troubleshooting notes and tips

This section provides general troubleshooting tips for the NMC server.

Troubleshooting an NMC server that is not responding

If the NMC server is not responding, answer the following questions:

- Is a potentially long-running process such as a device operation (label or inventory, for example) currently running?
Any process that you start on the Console server locks the user interface until that process completes. To perform multiple, long-running operations simultaneously (that is, to administer multiple NetWorker servers), open a separate web browser instance of the NMC server to run each operation.
- Are all of the following processes running?
 - GST server (**gstd**)
 - Database server **dbsrv12**
 - Web server (**httpd**)
- Is the `ntpd` command synchronizing at midnight?
In some cases, when you have a cron job that performs an `ntpd` synchronization at exactly midnight, the NMC server may lose connection to the database. To resolve this issue, modify the time that the cron job performs the `ntpd` synchronization or have `ntp` run as a service and synchronize continuously.

Determining if the Console server is running on a Windows system

On a Windows computer:

1. From the **Start** menu, select **Control Panel > Administrative Tools > Services**.
2. Verify that **EMC GST Service** is running.

Determining if the Console server is running on a Linux system

Use the `ps` command to confirm that the process, which the NMC server requires, are running.

- For the `gst` server process, type:

```
/usr/bin/ps -ef | grep gstd
```

If the `gstd` process is running, output similar to the following appears:

```
nsrnmcc 7190 1 0 Nov23 ? 00:00:06 /opt/lgtomc/bin/gstd
```

- For the database server process, type:

```
/usr/bin/ps -ef | grep postgres
```

If the database server is running, output similar to the following appears:

```
nsrnmcc 7212 1 0 Nov23 ? 00:00:00 /opt/lgtomc/postgres/bin/postgres -D /nsr/nmc/nmcdb/pgdata
nsrnmcc 7213 7212 0 Nov23 ? 00:00:00 postgres: logger process
```

```

nsrnmcc 7215 7212 0 Nov23 ? 00:00:00 postgres: checkpointer
process
nsrnmcc 7216 7212 0 Nov23 ? 00:00:00 postgres: writer process
nsrnmcc 7217 7212 0 Nov23 ? 00:00:00 postgres: wal writer
process
nsrnmcc 7218 7212 0 Nov23 ? 00:00:00 postgres: autovacuum
launcher process
nsrnmcc 7219 7212 0 Nov23 ? 00:00:03 postgres: stats
collector process
nsrnmcc 7231 7212 0 Nov23 ? 00:00:00 postgres: lgto_gst
lgto_gst 127.0.0.1(44296) idle

```

- For the web server process, type:

```
/usr/bin/ps -ef | grep httpd
```

If the web process is running, output similar to the following appears:

```

nsrnmcc 7196 1 0 Nov23 ? 00:00:00 /opt/lgtonmc/apache/bin/
httpd -f /opt/lgtonmc/apache/conf/httpd.conf
nsrnmcc 7197 7196 0 Nov23 ? 00:00:00 /opt/lgtonmc/apache/bin/
httpd -f /opt/lgtonmc/apache/conf/httpd.conf

```

Unable to connect to host: Please check Security setting and daemon logs on the NetWorker client and Console server for more information

This message can appear when you perform **Client Configuration** wizard tasks, **Device Configuration** wizard tasks, or when you browse save sets simultaneously while you add or modify a Client resource.

Check for one of the following scenarios when you receive this error.

- Verify that the SSL key matches between the NMC Server and the NetWorker client host. The SSL key is in the NSR Peer Information attribute, which is located in each host's nsrladb database. A mismatch can occur when the nsrladb on one host is corrupted.

To resolve this issue, delete the Console Server's NSR Peer Information from the NetWorker Client's nsrladb, and delete the NetWorker Client's NSR Peer Information from the Console Server's nsrladb as following:

- To delete the Console Server's NSR Peer Information from the NetWorker Client's nsrladb, on the client host, type:

```

nsradmin -p nsrexec
nsradmin> print type:NSR peer information

```

Note: Identify the Console Server's NSR Peer Information, and delete it.

```

nsradmin> delete type: NSR peer information;name:<Console
Server name>
Delete? Yes

```

- To delete the NetWorker Client's NSR Peer Information from the Console Server's nsrladb, on the Console Server host, type:

```

nsradmin -p nsrexec
nsradmin> print type:NSR peer information

```

Note: Identify the NetWorker Client's NSR Peer Information, and delete it.

```
nsradmin> delete type: NSR peer information;name:<Client
name>
Delete? Yes
```

After the deletion is complete, it is not mandatory to restart the NetWorker or Console services.

- The Client cannot resolve hostname of NMC Server or NW Server. Sometimes, NMC can resolve the client hostname, but, client cannot resolve the NMC or NetWorker Server hostname.
To resolve this issue, ping the NetWorker Server and NMC server from the Client. If the ping fails, DNS is not resolving the hostname issue and add the hostname to the client hosts file.
- Ensure NetWorker users have at least the "Operate NetWorker" privilege to launch the Client Wizard. To resolve this issue, add the user to the `user_group` in the NetWorker Server.
- The NetWorker Server may not be present in the client's servers file. To resolve this issue, add the NetWorker Server to the client's servers file.
- The NMC Server, NetWorker Server, and NetWorker client hosts must only use `nsrauth` authentication.

Username/password validation fails when you use the NMC New Device wizard to configure an AFTD if storage node is UNIX

When you use the NMC **New Device** wizard to configure an AFTD, the username/password validation for browsing the file system may fail for a UNIX storage node .

This failure can occur in the following situations:

- The system is missing the Pluggable Authentication Modules (PAM) library.
- The rule in the `pam.conf` file (`/etc/pam.conf`) for OTHER service is set to deny.

The operating system documentation provides more information about how to install the PAM library and how to modify the `pam.conf` file.

Querying large numbers of save sets in the NetWorker user interface may cause a Java heap space error

When you query a large number of save sets in the NetWorker user interface, the query may fail with a Java heap space error.

To resolve this issue, increase the Java heap size that the NMC application uses.

1. On the NMC server host, open the `Console_install_dir\web\gconsole.jnlp` file in a text editor.
2. Increase the `default max-heap-size` value from 700MB to 1400MB.

For example:

```
<resources>
<j2se version="1.5+" initial-heap-size="64M"
max-heap-size="1400M"/>
```

 **Note:** To provide meaningful query results and to reduce the chance of encountering this error, narrow the save set search criteria by specifying selection parameters.

NMC user interface exits unexpectedly

If the NMC guided user interface (GUI) loses its connection to the `gstd` service because the `gstd` service was shut down or failed, then the GUI displays a warning and exits after 10 seconds. This is normal behavior. [NMC error messages and corrective actions](#) on page 788 provides more troubleshooting information.

Label and Mount devices page is not displayed in NMC device configuration wizard

The display of the pool selection page is toggled based on the below criteria:

- The page is displayed when the user selects new folders that do not contain any volume information.
- The page is displayed when using existing folders containing volume information for another device (on different datazone), but not a device in the current NetWorker server.
- This page is not displayed when the user selects folders containing volume information for another device (on the same datazone).

CHAPTER 14

NetWorker Server Management

This chapter contains the following topics:

• Setting up the server	798
• Viewing the migration log file	799
• Hostname changes	799
• Managing the NSR task resource for nsrclientfix	800
• Parallelism and multiplexing	801
• Managing server access	805
• Resource databases	806
• Indexes	807
• Internationalization	816
• Creating a Server Backup action	817
• Creating an expire action	820

Setting up the server

When you set up the NetWorker server, enter the NetWorker product serial number that appears on the Enabler Certificate that you received from Dell EMC Licensing.

Procedure

1. From the **Administration** window, click **Protection**.
2. Select the server name.
3. From the **File** menu, select **Properties**.
4. In the **Properties** dialog box, configure the appropriate attributes.
5. Click the **System Summary** tab and enter the product serial number for the server, as well as any other required information.
6. Click **Ok**.

License the NetWorker Server

The *NetWorker Licensing Guide* describes how to license the NetWorker Server.

Setting the Job inactivity timeout

Use the Job inactivity timeout attribute to specify the maximum time, in minutes that the NetWorker server should wait for a response from a job before the server considers the job inactive and terminates the job.

About this task

The job inactivity timeout applies to all actions defined in all workflows in a policy. The inactivity timeout value assigned to an action, only applies to the action to which you defined the timeout value.

Procedure

1. On the **Administration** window, click **Server**.
2. In the left pane of the **Server** window, right-click the NetWorker server.
3. From the **File** menu, select **Properties**.
4. Select the **Configuration** tab.
5. In the **Job inactivity timeout** attribute, specify the timeout value in minutes.
6. Click **Ok**.

Modifying the retention period for jobs in the jobs database

By default, the NetWorker server retains information about jobs in the Jobs database for 72 hours. During this time, all details such as the status of workflows run will be available for viewing.

About this task

if required, you can change the jobs database retention to a longer period. Note, however, that as the retention period grows and data is preserved for a longer period of time, performance impacts may be observed. Perform the following steps to modify the amount of time NetWorker retains jobs information in the Jobs database:

Procedure

1. On the **Administration** window, click **Server**.
2. In the left pane of the **Server** window, right-click the NetWorker server.
3. From the **File** menu, select **Properties**.
4. Select the **Configuration** tab.
5. In the **Jobsdb retention in hours**, specify a retention time value in hours.

 **Note:** After the expiration of jobsdb, any expired workflows will display a status of `never run`.

6. Click **OK**.

Viewing the migration log file

When you update the NetWorker Server from version 8.2.x and earlier to version 19.1, the migration process creates log files that provide information about the resources and attribute migration results.

About this task

When you connect to the NetWorker Server for the first time after an update, a window appears that provides you with the option to view the main migration log window. The NetWorker Server does not remove the log files. Perform the following steps to view the main migration log file at a later time:

Procedure

1. Connect to the NetWorker Server from the NMC GUI.
2. From the **File** menu, select **Open Migration Log File**.

Hostname changes

NetWorker considers each unique client name as a separate client. NetWorker assigns each unique client name in the datazone a unique identifier called client ID. NetWorker stores the client ID for each client in the media database.

The NetWorker software has a built-in mechanism to prevent the `nsrkd` daemon from starting on the NetWorker server if the startup process detects a change in the name of the NetWorker server. For example, when you change the hostname of the NetWorker server or modify the aliases order in the `hosts` file.

A message similar to the following appears in the `daemon.raw` file:

```
NetWorker is unable to continue its startup sequence due to a server
hostname change to hostname. Please verify that the server's hostname
and its aliases are properly represented in the local host database
(eg. /etc/hosts) and DNS.
```

This mechanism prevents the NetWorker software from assigning a new client ID to the NetWorker server, which is based on the new hostname. To resolve this issue, verify the hostname resolution of the NetWorker server. The "Networking and Connectivity" chapter provides more information.

If the startup process did not detect the hostname change, NetWorker assigns the NetWorker server a new client ID, which can impact NetWorker operations. Use the `nsrclientfix` command to analyze the media database and identify client ID inconsistencies. To resolve client ID issues, use the `nsrclientfix` command to

merge information about multiple clients in the media database and resource database into one client resource with the original client ID. The following KB articles on the Online Support website provide more information about using the `nsrclientfix` command:

- For NetWorker Server client ID issues: 000185727
- For NetWorker Client client ID issues: 000193911

 **Note:** KB article 000196727 describes how to rename a NetWorker server.

Managing the NSR task resource for `nsrclientfix`

By default NetWorker uses an NSR Task resource that is named `DefaultNsrclientfixTask`. The resource runs the `nsrtask` command daily but only runs the `nsrclientfix` command on the days defined by the resource schedule. By default, NetWorker runs the `nsrclientfix` command every Sunday at 7:00 P.M. and reports client ID issues in the `daemon.raw` file.

About this task

When the `DefaultNsrclientfixTask` task detects a client ID issue, an error message similar to the following appears in the `daemon.raw` file:

```
nsrd NSR Index Warning: Detected error with client id(s): hostname
```

You can use the `nsradmin` program to modify the schedule of the scan.

Procedure

1. On the NetWorker server, start the `nsradmin` program from a command prompt.
2. At the `nsradmin` prompt, set the current query to select the NSR task resource named `DefaultNsrclientfixTask`:

```
print type:nsr task;name:DefaultNsrclientfixtask
```

Output similar to the following appears:

```
type: NSR task;
name: DefaultNsrclientfixTask;
comment: Periodic execution of nsrclientfix Task;
action: "NSR client fix:DefaultNsrclientfix";
autostart: Enabled;
start time: "7:00";
interval: "24:00";
period: Week;
plan: "
exec skip skip skip skip skip skip";
last start: "Thu Oct 30 15:13:04 2014";
last end: "Thu Oct 30 15:13:05 2014";
last message: Successful;
job id: ;
last job: 32086;
status: idle;
```

3. Use the `update` command to modify the following attributes:
 - Autostart— Acceptable options are Start now, enabled and disabled.

- **Start time**—Specify a new start time in the format "**HH:MM**".
- **Period**—Specifies when the plan cycle repeats. Acceptable options are week and month.
- **Interval**—Specifies how often to run the task. Specify a 24 hour clock value in the format "**HH:MM**".
- **Plan**—When you set the period to weekly, the plan attribute defines which days of the week the NetWorker server runs the `nsrclientfix` command. When you set the period to monthly, the plan attribute defines which days in a 30 day period the NetWorker server runs the `nsrclientfix` command. Acceptable values are `exec` and `skip`.

 **Note:** The **action** attribute specifies the name of the NSR Client Fix resource, which contains the `nsrclientfix` command.

For example, to specify that the task should run every day of the week at 1:00 P.M. except for Sunday, type the following command:

```
update: start time: "13:00"; plan: skip exec exec exec exec exec
exec
```

The *Command Reference Guide* provides more information about the `nsrtask` and NSR client fix resources.

Parallelism and multiplexing

Parallelism is a general term within the NetWorker software for a number of configurable options that allow you to adjust the volume of data that a system processes, which can improve the performance of servers, storage nodes, and devices. Multiplexing is the ability to write multiple save streams simultaneously to the same storage device.

This section identifies attributes related to parallelism and multiplexing and describes how they work together to optimize your NetWorker environment.

Parallelism

You can use several attributes in various NetWorker resources to adjust the volume of data that a host processes to improve overall performance.

The following attributes are related to parallelism:

- Client parallelism
- Server parallelism
- Action parallelism
- Max active devices
- Media library parallelism

These attributes are described in detail in the following sections.

Client parallelism and parallel save streams

Client parallelism defines the number of data streams that a client can use simultaneously during backup.

Data streams include backup data streams, savefs processes, and probe jobs.

The default value is different for the NetWorker server than it is for all other client resources:

- For the NetWorker server client resource, the default value is 12. This higher default value enables the server to complete a larger number of index backups during a Server backup action.
- For all other clients, the default value is 4.

To define client parallelism, use the **Parallelism** attribute of the Client resource. You can find the parallelism attribute on the **Globals(1 of 2)** tab of the Client property dialog box, in the **NetWorker Administration** window.

The *NetWorker Network Data Management Protocol (NDMP) User Guide* provides more information about recommended parallelism settings for NDMP clients.

To avoid disk contention for clients other than the NetWorker server, specify a value that is the same as or fewer than the number of physical disks on the client that are included in the backup.

For a Windows client with the ALL keyword save set attribute, the backup includes the local disks, for example C: and D: drives as well as the System State and System DB. In this example, you can keep the default parallelism setting of 4. If you define multiple save sets on the same disk, for example, C:\users, C:\system, C:\docs and so on, a higher client parallelism results in multiple save streams attempting to access the disk at the same time.

The *NetWorker Performance Optimization Planning Guide* provides more information about recommended client parallelism values and performance benefits.

Enabling the parallel save streams (PSS) feature for a Client resource allows you to back up each save set for the client by using multiple parallel save streams to one or more destination backup devices. You can use PSS to perform the scheduled file level backup of file systems, and block based backups.

You can use PSS for clients with supported UNIX, Linux, and Windows operating systems. Supported save sets for PSS include the Save Set ALL, and individual save points including Disaster_Recovery, deduplicated, and CSV volumes (Windows only). Checkpoint restart is not supported when you use PSS.

When you enable PSS, you can specify the maximum number of save streams that a client can send simultaneously for one or more save set backups concurrently running by using the **Parallelism** attribute in the **Client Properties** dialog. The default value for the **Parallelism** attribute is different for the NetWorker Server than it is for all other Client resources:

- For the NetWorker Server Client resource, the default value is 12. This higher default value enables the server to complete a larger number of index backups during a file system backup of the server or other index backups.
- For all other clients, the default value is 4.

Enabling PSS results in significant performance improvements due to save set aggregation, where the NetWorker Server starts a single save process per client with all client save sets that are passed to the single process for various processing optimizations, such as minimal Windows VSS snapshots and support for the following:

- Four parallel streams are started per save set, subject to any client parallelism limitations that might prevent all save sets from starting simultaneously.
- The ability to modify the number of parallel streams per save set by defining the new `PSS:streams_per_ss` environment variable save operations attribute in the properties of a Client resource. For example, setting `PSS:streams_per_ss=2,*` splits all save sets into two parallel save streams, whereas

PSS:streams_per_ss=3,/data1, 5,/data2 splits /data1 into three parallel save streams and /data2 into five parallel save streams.

- Automatic stream reclaiming, which dynamically increases the number of active streams for an already running save set backup to maximize utilization of limited client parallelism conditions.

Note: It is recommended that you set the client parallelism value to be a multiple of the *PSS:streams_per_ss* parameter default value 4 or its largest defined value when configured. For example, a multiple of 4 is 8, 12, or 16.

If the client parallelism is less than the *PSS:streams_per_ss* default 4 or the lowest configured value, the backup fails displaying an error message.

The *PSS:streams_per_ss* values range from 1 to 8. If you specify an invalid value, the backup proceeds with the default value 4, and a warning message displays stating that the entire *PSS:streams_per_ss* parameter is ignored.

The *NetWorker Performance Optimization Planning Guide* provides complete details on PSS requirements and performance benefits.

Server parallelism

To define the server parallelism for a NetWorker server, use the Parallelism attribute of the Server resource. The Parallelism attribute appears in the NetWorker **Administrator** window on the **General** tab of the **Server property** dialog box.

Server parallelism defines the number of simultaneous data streams that the NetWorker server allows.

Data streams include backup data streams, savefs processes, and probe jobs.

The default and the maximum server parallelism values vary depending on the edition of NetWorker software. Each storage node that you enable and connect to the NetWorker server can increase the parallelism maximum value. The maximum parallelism value for any NetWorker server and storage node combination can vary. The *NetWorker Release Notes* provides more information.

Optimally, configure the NetWorker server to process enough data streams to keep all the backup devices in the datazone writing at their maximum speed. When you tune the server parallelism setting, along with other settings discussed in this section, you can maximize the speed that NetWorker writes the data to backup devices.

Action parallelism

Action parallelism defines the maximum number of simultaneous data streams that can occur on all clients in a group that is associated with the workflow that contains action.

Data streams include backup data streams, savefs processes, and probe jobs.

To define the parallelism for an action, modify the Parallelism attribute on the **Specify the Advanced Options** page in the **Action** wizard. For a Backup action, the default parallelism value is 100 and maximum value is 1000. For a clone action, the default parallelism value is 0. For all other action types, the default value is 0, or unlimited.

Max active devices

In a DDS environment, use the Max active devices attribute, on the **General** tab of the Storage Node resource to define the maximum number of active devices for a storage node.

This attribute sets the maximum number of devices that NetWorker may use from the storage node in a DDS configuration. In large environments with media libraries that have a large number of devices, storage nodes might not have the ability to optimize all the drives in the library. The Max active devices attribute allows you to limit the number of devices that the storage node uses at a specified time, which allows the storage node to have access to all the devices in the library, but does not limit the storage node to the number of devices it can fully optimize.

Media Library parallelism

To define the media library parallelism, use the **Max parallelism** attribute on the **Configuration** tab of the Library resource .

Media library parallelism allows you to define the maximum number of available devices for inventory and label operations.

It is recommended that you set the **Max parallelism** attribute of the Library resource to one less than the number of devices within the library, which allows you to reserve on device for recovery operations.

To improve the efficiency of library operations that operate on multiple volumes, use multiple devices in parallel for these operations. However, you may want to restrict the number of devices that NetWorker uses for inventorying and labeling operations, to ensure that some devices are available for other library operations.

Multiplexing

Multiplexing is the ability to write multiple data streams simultaneously to the same storage device. It is often more efficient for the NetWorker server to multiplex multiple save sets to the same device. There are also times when limiting the number of data streams to a particular device improves performance of the NetWorker environment.

Use the Target sessions, Max sessions, and Pool parallelism attributes to increase or limit the number of data streams that NetWorker writes to a device.

Target sessions

Use the Target sessions attribute on the Configuration tab of the Device resource to define the optimal number of backup sessions to assign to an active device.

Target sessions is not a hard limit; to set a hard limit for the number of sessions that a particular device can accept, use the Max sessions attribute.

The Target sessions attribute aids in load balancing devices by determining when the NetWorker software should write save streams to a device.

When a save session starts, the following actions occur:

- If a device is already receiving the number of backup sessions determined by the target sessions value, the NetWorker server uses the next underutilized device for the backups.
- If all available devices are receiving the number of backup sessions determined by their target sessions value, the NetWorker server overrides the set value and uses the device with the least activity for the next backup session.

Because it is often more efficient for the NetWorker server to multiplex multiple save sets to the same device, rather than write each save set to a separate device, the NetWorker server attempts to assign to each device a number of save sets, up to the value of target sessions, before assigning a save set to another device.

NOTICE When the NetWorker software assesses how many devices need to be involved in multiple save streams assignments with the same storage node, the device with the lowest target session value is used as a reference.

Max sessions

The Max sessions attribute on the Configuration tab of the Device resource defines the maximum number of save sessions for a device. The max sessions value is never less than the target sessions value. It is recommended to use the default values for Max sessions as lowering these values can impact performance.

Pool parallelism

The Max parallelism attribute on the Configuration tab of the Pool resource defines the parallelism for a pool.

Pool parallelism determines the maximum number of simultaneous save streams for each device that belong to a NetWorker pool. The default value for this attribute is 0, which means that the attribute has no effect on other parallelism settings.

You can use pool parallelism to increase recovery times. For example, you can create a pool to back up business critical data and use this attribute to restrict the number of save sets that NetWorker writes in parallel to the media in the pool. As a result, recovery speed increases for data on that media.

However, when you set the Max parallelism attribute to 1, a prolonged delay between the backup of save sets may occur. To resolve this issue, increase the Max parallelism attribute for the pool resource.

Note: For AFTD and DD Boost devices, the Max nsrmmmd count setting for a device affects the Max parallelism attribute. For example, consider an AFTD device (AFTD_1) that has a Max sessions attribute of 20 and a Max nsrmmmd count of 4. Now suppose a backup pool with a Pool parallelism attribute of 1 selects AFTD_1. The total number of save sessions that NetWorker can initiate for AFTD_1 is 4, one for each nsrmmmd process. Tape and FTD devices can only spawn one nsrmmmd process at a time, so if the previous example used a tape device, then the total number of save sessions would be 1.

Managing server access

User privileges define the NetWorker operations and tasks that NMC, AD, and LDAP users can perform on a NetWorker Server.

The *NetWorker Security Configuration Guide* describes how to restrict access to the NetWorker Server and NetWorker operations, including the following information:

- How to restrict administrator access to the NetWorker Server.
- How to modify the privileges assigned to NMC, LDAP, and AD users and groups.
- How to Restrict server and client initiated backup and recover operations.

Resource databases

Information about the NetWorker Server resides in series of files in the following directories:

```
NetWorker_install_path\res\nsrd\00
.
NetWorker_install_path\res\nsrd\09
```

NetWorker stores each resource in a separate numbered file. As you create resources, for example, a new Client, Group, or Pool resources, the NetWorker Server adds files to the directories.

A Client resource database (`nsrexec`) also exists on each NetWorker host and contains configuration information about each NetWorker host. The `nsrexec` database resides in a series of files in the following directories:

```
NetWorker_install_path\res\nsrladb\00
.
NetWorker_install_path\res\nsrladb\09
```

The *NetWorker Security Configuration Guide* provides more information about the Client resource database.

Viewing resources in the resource database

You can view and modify NetWorker resources through the NMC **Administration** window.

NetWorker also provides a command line tool, `nsradmin`, to modify resource databases.

For example:

- To access the NetWorker server resource database, type:

```
nsradmin -s server_name
```

- To access the client resource database, use the following command :

```
nsradmin -p nsrexec
```

- To access the Package Manager database, use the following command:

```
nsradmin -p nsrcpd
```

Repairing resource database corruption

A power outage, operating system failure, or manual edits the database with a text editor can cause NetWorker resource database file corruption.

If the NetWorker server cannot read the resource files when the NetWorker services start, a message similar to following appears in the `daemon.raw` file.

```
nsrd: WARNING: NSR configuration database detected invalid
resource ...\00019803aa14713c89456b41
nsrd: Invalid resource saved at ...\00019803aa14713c89456b41
```

The NetWorker server removes any corrupt resource files from the `nsrdb` directory structure and places them in the `dbg` directory. NetWorker creates the `dbg` directory only after resource database file corruption has occurred. To correct this issue, open

the corrupt file with a text editor and review the file contents, then re-create the resource. You can delete the corrupted resource file.

 **Note:** If you do not know the cause of the resource file corruption, contact Technical Support assistance.

Indexes

The NetWorker server tracks the files it backs up in two databases, which are stored on the local file system of the server:

- The client file index tracks the files that belong to a save set. There is one client file index for each client.
- The media database tracks:
 - Volume name
 - Backup dates of the save sets on the volume
 - File systems in each save set

Unlike the client file indexes, there is only one media database per server.

The client file indexes and media database can grow to become prohibitively large over time. [Managing the size of the online indexes](#) on page 812 provides information about managing the size of these indexes.

Characteristics of the online indexes

The size of an index is proportional to the number of entries the index contains. The media database is usually smaller than the client file index, because the media database stores one entry for each volume, while the client file index stores one entry for each file that NetWorker saves on a volume. The NetWorker server selects which volume to mount to perform a recovery by mapping the saved files to their volumes.

Each entry in the client file index includes this information for a saved file:

- Filename
- Number of blocks
- Access privileges
- Number of links
- Owner
- Group
- Size
- Last modified time
- Backup time

The client file indexes grow with each backup, as entries are added for each newly saved file and save set. As long as an index entry for a file remains in the client file index, you can perform a browsable recovery of the file. Over time, the size of these indexes can grow very large.

 **NOTICE** If the file system that contains the indexes gets full, the NetWorker server cannot access the media database and cannot recover data. Unless you use browse and retention policies to control the size of the online indexes, the indexes continue to grow until they exceed the capacity of the file system.

NetWorker uses browse and retention policies to manage the lifecycle of the data, and to automatically control the size of the client file index. [Backup retention](#) on page 386 provides information about policies.

Automated index activities

The NetWorker server performs these online index activities:

- Inserts entries in the client file index for each file saved during a backup. For each new backup, the NetWorker server acquires more space from the file system for the new entries.
- Removes entries and returns disk space to the operating system. The browse and retention policies automatically determine when entries are removed from the index.

You can also remove index entries manually by clicking **Remove Oldest Cycle** in the **Index Save Sets** dialog box. [Removing the oldest save set cycles](#) on page 815 provides more information.

Checking online indexes

Each time the NetWorker server starts, the startup process uses `nsrck -ML1` to perform a level 1 consistency check on the client file indexes. In some circumstances, this consistency check will not detect corruption in the client file indexes. If you believe that an index may be corrupt, run a higher level check on the index, for example:

```
nsrck -L5
```

If the index is still corrupt, recover the index by using the procedure that is outlined in [Adding information about recyclable save sets to the client file index](#) on page 559.

It is recommended that you periodically run the `nsrck -F` and `nsrim -X` commands to check the integrity of the client and media indexes. The *NetWorker Command Reference Guide* or the UNIX man pages provide more information about these commands.

Viewing information about the indexes

The following table identifies the index information displayed for each client.

About this task

Table 137 Indexes window information

Column	Description
Client Name	Names of the NetWorker clients that have been backed up by the current server.
Size	Amount of disk space currently allocated to the client file index. As the index size increases, the allocated disk space automatically grows.

Procedure

1. From the **Administration** window, click **Media**.

2. In the left pane, click **Indexes**. The right pane displays index information for all clients of the server.

Index save sets

The **Index Save Sets** dialog box displays the save sets assigned to a particular client, along with detailed information about each save set. The dialog box also includes an option to remove old save set cycles.

Viewing client save set information

The following table identifies the information in the Save Sets dialog box for each save set.

About this task

Table 138 Index save sets dialog box information

Column	Description
Save Set Name	Name of the save set.
Size	Estimated amount of the index space used by the save set in the client file index.
Cycles	Number of backup cycles available for browsing. A cycle starts with a full backup and ends with the next full backup, and includes any incremental and level 1–9 backups that occur between full backups.
SSID	Unique identification number of the instance of the save set.
Files	Number of files backed up during that instance.
Size	Size of the backup.
Time	Date and time of the backup.
Level	Level of the backup (full, incr [incremental], or 1-9)

[Reduce the size of the client file index](#) on page 812 provides information about reducing the size of the client file indexes by using the Remove Oldest Cycle button.

Procedure

1. From the **Administration** window, click **Media**.
2. Click **Indexes**.
3. Right-click the client whose save sets you want to view, then click **Show Save Sets**. The **Index Save Sets** dialog box appears.
4. To view detailed information about a save set, click the save set name.

Querying the media database

You can query the media database for information about save sets. Queries apply to all complete, browsable save sets, not just those from the last 24 hours.

Procedure

1. From the **Administration** window, click **Media**.
2. Click **Save Sets**.
3. On the **Query Save Sets** tab, indicate the appropriate query parameters, then click the **Save Set List** tab to run the query and view the results.



NOTICE If the query is unsuccessful, an error dialog box appears, which indicates that NetWorker could not find save sets that matched the specified query. Click **OK** to close the dialog box.

Results

You can also use the `mminfo -av` command to query the media database. The *NetWorker Command Reference Guide* or the UNIX man pages provides detailed information about how to use the `mminfo` command.

Cross-checking client file indexes

Perform a cross-check to verify the consistency between the client file index and the media database. If the NetWorker server finds entries in the client file index that do not have corresponding entries in the media database, it removes the client file index entries. This feature is useful, for example, if you perform an index operation and the server fails before the NetWorker server has completely updated the indexes. Once the server is running again, cross-check to accurately update the online indexes.

Procedure

1. From the **Administration** window, click **Media**.
2. Click **Indexes**.
3. Right-click the client with the index to cross check, then select **Cross Check Index**.

The following prompt appears:

```
Cross-checking may take considerable time. Would you like to
cross-check
client_name?
```

4. Click **Yes** to continue. The NetWorker server displays a status box until the cross-checking is complete.

Refreshing index information

Occasionally refresh the information in the Indexes tab, particularly if you are connected to a server for a long period of time.

Procedure

1. From the **Administration** window, click **Media**.
2. Click **Indexes**.

3. From the **View** menu, select **Refresh**.

Client file index locations

During the initial client setup, the NetWorker software normally designates a default location for the client file index on the NetWorker server. This default location is:

- For UNIX: `/nsr/index/client_name`
- For Windows: `NetWorker_install_path\index\client_name`

However, you may need to designate a different index location when first configuring a Client resource, or you might need to move the file index of an existing client. These sections address these needs.

Designating the client file index location for a new client

Procedure

1. From the **Administration** window, click **Protection**.
2. Right-click **Clients**, then select **New**. The **Create Client** dialog box appears.
3. Click the **Globals (2 of 2)** tab.
4. In the **Index Path** attribute, type the full path of the directory where the client file index resides.
5. For the remaining tabs, type information as necessary to create the new client.
6. Click **Ok**.

Changing the client file index location for an existing client

To change the client file index location to a nondefault location for an existing client, you must first move the index to its new location.

Moving a client file index

You can move a client file index from its current location to a new location. For example, if the size of the client file index is too large, you can move it to a location with more space.

Procedure

1. Ensure that backups and recovers are not occurring on the NetWorker server.
2. Log in to the NetWorker server root on UNIX or as an administrator on Windows.
3. From the directory that contains the indexes, type:

```
uasm -s -i "client_index_directory_name" | (cd target_directory; uasm -r)
```

 **Note:** On Solaris and Linux platforms, `uasm` is installed in `/usr/lib/nsr`. On all other platforms, `uasm` is installed in the same location as the NetWorker binaries.

4. Use NMC to connect to the NetWorker server.
5. Click **Protection**, then click **Clients** in the left navigation pane.
6. Right-click the client that requires the client file index location update, and then select **Modify Client Properties**.

7. On the **Globals (2 of 2)** tab, in the **Index Path** attribute, specify the full path of the directory where the client file index now resides.
8. Click **OK**.

Updating the index location for a client in NetWorker

Procedure

1. From the **Administration** window, click **Protection**.
2. Click **Clients**.
3. Right-click the client with the client file index location to be changed, then select **Properties**. The **Properties** dialog box appears.
4. Click the **Globals (2 of 2)** tab.
5. In the **Index Path** attribute, type the full path of the directory where the client file index now resides.
6. Click **OK**.
7. (Optional) From a command prompt, run the `nsrck` or `nsrls` command and check the output for any errors.

For example, to run `nsrck` on client `jupiter`, type:

```
nsrck -L6 jupiter
```

Output similar to the following appears:

```
nsrck: checking index for 'jupiter'
nsrck: nsrindexesjupiter contains 54 records occupying 7 KB
nsrck: Completed checking 1 client(s)
```

NOTICE Depending on the size of the client file index, running either `nsrck` or `nsrls` can take a considerable amount of time. Running the `nsrck -L6` command, as shown in the example, also checks the index for corruption.

If no problems are found, then all future client file index information is saved to the new location.

Managing the size of the online indexes

Over time, the size of the online indexes on the NetWorker server can become prohibitively large. Reduce the size of these indexes by using the solutions suggested in these sections.

Reduce the size of the client file index

You can reduce the size of the client file indexes on the NetWorker server by using one or more of these methods:

- Remove save sets that comprise the oldest backup cycle from the client file index. [Removing the oldest save set cycles](#) on page 815 provides details.
- Delete volume-based entries from the client file index. [Deleting volume-based online index entries](#) on page 814 provides details.
- Adjust the Browse Policy and Retention Policy attributes of clients backing up to the NetWorker server to shorten the period of time that entries remain in the client file indexes. This solution works only for client backups that occur after you change these policy attributes.

- Modify the Retention policy associated with a particular save set by using the `nsrmm -e` command. Unless the associated save set contains a large number of files, this method may not be a practical method to reduce the index size. [Editing retention for a save set](#) on page 388 provides details.

If the size of the client file index for a client is still too large, consider moving the location of the index. [Moving a client file index](#) on page 811 provides details.

Reduce the size of the media database size

Use one or more of the following methods to reduce the size of media database on the NetWorker server.

- Remove volumes that contain recyclable save sets from the NetWorker inventory. [Removing volume-based entries from the online indexes](#) on page 813 provides details.

When you remove a volume from the media database, NetWorker removes the entries associated with that volume from the media database and the client file index for the client. If you select this option, you can use the `scanner` command to recover the data on the volume, if NetWorker has not relabeled the volume.

NOTICE You will gain very little disk space from removing a media database entry. Leaving index entries of a volume in the media database prevents the accidental labeling of another volume with the same name.

- Recycle volumes that contain recyclable save sets. [Changing the volume mode](#) on page 537 provides details.

When a volume mode changes to recyclable, the volume becomes eligible for reuse and NetWorker can perform the following operations:

- Relabel the volume
- Remove information about the save sets on the volume from the media database
- Reinitialize the volume

Once NetWorker relabels a volume, you cannot recover the contents.

To increase the number of currently recyclable save sets, modify the retention policy associated with the current media database by using the `nsrmm -e` command. [Editing retention for a save set](#) on page 388 provides details.

- Compress the media database. [Compressing the media database](#) on page 815 provides details.

Removing volume-based entries from the online indexes

The main purpose of removing volume-based entries from the online indexes is to eliminate damaged or unusable volumes from the NetWorker server. You can also use this feature to reduce the size of the online indexes by purging index entries associated with specific volumes.

Removing client file index entries

Use the `nsrmm` command to remove information about save set from the client file index. This changes the status of browsable save sets to recoverable.

Procedure

1. At the command prompt, type:

```
nsrmm -d -P -S ssid
```

where *ssid* is the save set ID for the save set.

2. Use `mminfo` to determine the save set ID. At the command prompt, type:

```
mminfo -v -c client_name
```

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrmm` and `mminfo` commands.

Results

When NetWorker marks a save set as recoverable, you cannot browse to recover these files. Use the save set recover procedure to recover data from a recoverable save set.

Removing client file index and media database entries

About this task

You can remove both the client file index and media database entries for a volume. This action removes all traces of the volume from the NetWorker server. Remove a volume from the media database only if the volume has been physically damaged and is unusable. However, if you remove the database entries for a volume, the volume is undamaged, and NetWorker has not relabeled the volume, you can use the `scanner` command to recover the data. [Adding information about recyclable save sets to the client file index](#) on page 559 provides details.

Typically, do not remove both the client file index and media database entries at the same time unless the volume is damaged or destroyed.

NOTICE The presence of a clone of a particular volume prevents the deletion of the volume entry in the media database. This is because the NetWorker server accesses the cloned volume rather than the original volume as needed. NetWorker does not purge the entry of the volume in the media database. Because of this functionality, removing volume entries from the media database is not a particularly effective way to reduce index size.

Deleting volume-based online index entries

You can use NMC or the `nsrmm` command to delete volumes from the media database and client file indexes. The NetWorker server first cross-checks the indexes before it clears a volume. As a result, the volume might still appear in the **Volumes** window in NMC for a brief period.

About this task

Procedure

1. From the **Administration** window, click **Media**.
2. Click **Volumes**.
3. Right-click the volume with the entry to delete from the online indexes, then select **Delete**.
4. Select one of these options to determine how volume entries are removed:
 - **File and Media Index Entries.** [Removing client file index entries](#) on page 813 provides details about this option.
 - **File Index Entries Only.** [Removing client file index and media database entries](#) on page 814 provides details about this option.
5. Click **OK**.

Deleting volumes from a command prompt

Use the `nsrmm` command to remove volume information from the media database and client file indexes.

About this task

To remove both client file index and media database entries for a volume, type the following command:

```
nsrmm -d -S ssid
```

To remove information about the volume from the client file index only, type the following command:

```
nsrmm -d -P volume_name
```

Deleting volumes in NMC

Use NMC to remove volumes from the client file index or from both the media database and client file index.

Procedure

1. From the **Administration** window, click **Media**.
2. Click **Volumes**.
3. Right-click the volume with the entry to delete from the online indexes, then select **Delete**.
4. Select one of these options to determine how volume entries will be removed:
 - **File and Media Index Entries** to remove the volume information from the media database and client file indexes.
 - **File Index Entries Only** to remove the volume information from the client file indexes only.

Compressing the media database

You can free up more space on the server by compressing the media database.

Procedure

1. Delete the appropriate file:
 - On Windows:


```
NetWorker_install_dir\mm\cmprssd
```
 - On UNIX:


```
/nsr/mm/.cmprssd
```
2. Type the following command at the command prompt:

```
nsrim
```

Removing the oldest save set cycles

Client file index entries for a full save set cycle include the last full backup and any dependent incremental or level saves. When you remove the oldest cycle, you free up disk space.

Procedure

1. From the **Administration** window, click **Media**.
2. Click **Indexes**.

3. Right-click the appropriate client, then select **Show Save Sets**.
4. Select the save set with the oldest cycle to remove, then click **Remove Oldest Cycle**.
5. When prompted, click **Yes** to confirm the removal.

Results

After the Remove Oldest Cycle operation has finished, NetWorker updates the statistics in the Index Save Sets dialog box to reflect the current state of the client file index.

Internationalization

The NetWorker software supports language packs, which you can install as part of the NetWorker installation, or you can install the language packs separately after you have installed the NetWorker software. The *NetWorker Installation Guide* provides more information.

Internationalization support in the NetWorker software depends on internationalization support of the underlying operating system. If you plan to use non-English data in the NetWorker software, ensure that you install and configure the appropriate support for that language on the operating system.

The following sections describe a number of issues and limitations that relate to the use of NetWorker software in a multi-language environment.

Log file viewer

To view NetWorker log files, use the `nsr_render_log` program.

Display issues

There are number of issues and limitations associated with displaying characters in various locales.

Character display at the command line

From the command line, characters supported by the current locale display correctly. Characters that the current locale of the user do not support will not appear correctly. For Microsoft Windows systems, if the user and system locales do not match, characters supported in the user locale but not the system locale may not appear correctly.

Character display in graphical user interfaces

How character display from within the different NetWorker GUIs vary and depend on the platform on which you run the GUI.

- On Microsoft Windows:
 - All Unicode encoded data will display correctly.
 - When you view UNIX path and filenames, path and filenames that you create with a character set that the current locale or UTF-8 supports, will display correctly. Paths that you create with another character set may not display correctly. Because Microsoft Windows does not have native support for many

of the character sets used on UNIX (for example, euc-jp, euc-cn and euc-tw), if a non-ASCII character is encoded by using these character sets, characters will not display correctly on Microsoft Windows.

- On Unix:
 - Characters that the current locale does not support may not display correctly.
- On OS-X:
 - Differences in Unicode support, non-ASCII paths, and filenames on OS-X machines can result in characters not displaying correctly when you browse the file system from a non-Mac platform.

Creating a Server Backup action

A Server Backup action performs a bootstrap backup of the NetWorker media and resource databases, and can also include the client file indexes. By default, the NetWorker server configuration contains a Server Protection policy that contains NMC server backup and Server db backup workflows. The Server db backup workflow contain a server backup action. This section describes how to create a new server db backup action, if required.

Before you begin

Create the policy and workflow that contain the action. The Server Backup action should be the first action in the workflow.

Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. From the **Action Type** list, select **Server Backup**.
3. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
4. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
5. Specify a weekly, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
 - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
6. Click the icon on each day to specify the type of backup to perform.

To perform the same type of backup on each day, select the backup type from the list and click **Make All**.

7. Click **Next**.

The **Server Backup Options** page appears.

8. From the **Destination Storage Node** list, select the storage node with the devices on which to store the backup data.
9. From the **Destination Pool** list, select the media pool in which to store the backup data.
10. From the **Retention** lists, specify the amount of time to retain the backup data.

After the retention period expires, the save set is marked as recyclable during an expiration server maintenance task.

11. Specify whether to include the client file indexes in the server backup by selecting or clearing the **Perform CFI** checkbox.

When you clear this option, the action will only backup the bootstrap.

12. Specify whether to include a bootstrap backup in the server backup by selecting or clearing the **Perform Bootstrap** checkbox.

When you clear this option, the action will only backup the client file indexes.

 **NOTICE** You must select either the **Perform CFI** checkbox, the **Perform Bootstrap** checkbox, or both checkboxes. Otherwise, the server backup action does not back up any data.

13. Click **Next**.

The **Specify the Advanced Options** page appears.

14. In the **Retries** field, specify the number of times that NetWorker should retry a failed probe or backup action, before NetWorker considers the action as failed. When the **Retries** value is 0, NetWorker does not retry a failed probe or backup action.

 **Note:** The **Retries** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option for other actions, NetWorker ignores the values.

15. In the **Retry Delay** field, specify a delay in seconds to wait before retrying a failed probe or backup action. When the **Retry Delay** value is 0, NetWorker retries the failed probe or backup action immediately.

 **Note:** The **Retry Delay** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. When you specify a value for this option in other actions, NetWorker ignores the values.

16. In the **Inactivity Timeout** field, specify the maximum number of minutes that a job run by an action can try to respond to the server.

If the job does not respond within the specified time, the server considers the job a failure and NetWorker retries the job immediately to ensure that no time is lost due to failures.

Increase the timeout value if a backup consistently stops due to inactivity. Inactivity might occur for backups of large save sets, backups of save sets with large sparse files, and incremental backups of many small static files.

Note: The **Inactivity Timeout** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option in other actions, NetWorker ignores the value.

17. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

18. From the **Failure Impact** list, specify what to do when a job fails:

- To continue the workflow when there are job failures, select **Continue**.
- To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.

Note: The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.

- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

Note: If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

19. From the **Send Notifications** list box, select whether to send notifications for the action:

- To use the notification configuration that is defined in the Policy resource to send the notification, select **Set at policy level**.
- To send a notification on completion of the action, select **On Completion**.
- To send a notification only if the action fails to complete, select **On Failure**.

20. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.

21. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.

22. (Optional) Configure overrides for the task that is scheduled on a specific day.

To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:

- Select the day in the calendar, which changes the action task for the specific day.

- Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.
 - To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

 Note:

- You can edit or add the rules in the **Override** field.
- To remove an override, delete the entry from the **Override** field.
- If a schedule is associated to an action, then override option is disabled.

23. Click **Next**.

The **Action Configuration Summary** page appears.

24. Review the settings that you specified for the action, and then click **Configure**.

After you finish

(Optional) Create a clone action to automatically clone the bootstrap backup when the backup completes or create an expire action.

 **Note:** NetWorker only supports one action after the server backup action.

Creating an expire action

The expire action removes all expired save sets from the client file index and marks the save sets as recyclable in the media database. Save sets expire when the retention period for the save set is exceeded. You can create an expiration action in an existing workgroup only after a server backup action.

Before you begin

Create the policy and workflow that contain the action. The expire action should be the first action in the workflow or you can create the expire action after a server backup action.

Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: \ * ? [] () \$! ^ ; ' " ` ~ > < & | { }

3. In the **Comment** field, type a description for the action.

4. To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.



Note: When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

5. From the **Action Type** list, select **Expire**.
6. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
7. Specify a weekly, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
 - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
8. Click the icon on each day to specify whether to perform expiration.

The following table provides details on the icons.

Table 139 Schedule icons for the expire action

Icon	Label	Description
	Execute	Perform expiration on this day.
	Skip	Do not perform expiration on this day.

To perform expiration every day, select **Execute** from the list, and click **Make All**.

9. Click **Next**.
The **Expiration Options** page appears.
10. Click **Next**.
The **Specify the Advanced Options** page appears.
11. From the **Failure Impact** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.



Note: The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.

- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

Note: If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

12. From the **Send Notifications** list box, select whether to send notifications for the action:
 - To use the notification configuration that is defined in the Policy resource to send the notification, select **Set at policy level**.
 - To send a notification on completion of the action, select **On Completion**.
 - To send a notification only if the action fails to complete, select **On Failure**.

13. (Optional) Configure overrides for the task that is scheduled on a specific day.
To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:

- Select the day in the calendar, which changes the action task for the specific day.
- Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.
 - To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

Note:

- You can edit or add the rules in the **Override** field.
- To remove an override, delete the entry from the **Override** field.
- If a schedule is associated to an action, then override option is disabled.

14. Click **Next**.

The **Action Configuration Summary** page appears.

15. Review the settings that you specified for the action, and then click **Configure**.

CHAPTER 15

NetWorker Host Management

This chapter contains the following topics:

- [Controlling access to a NetWorker client](#)..... 824
- [NetWorker host management](#)..... 824
- [Windows client interface](#)..... 826
- [Editing a client NSRLA database](#)..... 829

Controlling access to a NetWorker client

NetWorker uses the contents of the `/nsr/res/servers` (UNIX), or the `NetWorker_install_path\res\servers` (Windows) file on each NetWorker client to control who has client-tasking rights. Client-tasking rights provide a host with the right to request a program execution on another client. The following table provides a list of tasks that require an update to the servers file.

Table 140 When to modify the servers file

Operations	Update required on the NetWorker client's servers file
Archive request	Add the FQDN or shortname of the NetWorker server.
Scheduled backup	Add the FQDN or shortname of the NetWorker server. For a clustered NetWorker server, add the FQDN and shortname of the virtual NetWorker and all physical nodes.
Remote Directed Restore	Add the FQDN or shortname of the administering client to the server file on the destination client.
NDMP DSA backups	Add the FQDN or shortname of the NetWorker client that starts the backup.  Note: For NDMP, the servers file resides in the NetWorker Server.

 **Note:** Before adding the FQDN or shortname to the NetWorker server file, ensure that the host name resolution for FQDN or short name is working correctly.


The *NetWorker Security Configuration Guide* provides more information about client-tasking rights and how to modify the servers file.

NetWorker host management

The **Hosts** window on the **NetWorker Administration** window provides you with the ability to manage NetWorker hosts in the datazone.

The following figure provides an example of the **Hosts** window.

Figure 86 Hosts window



The **Hosts** window contains a taskbar and two window panels, the summary panel and a task monitoring panel.

The information that appears in summary panel changes based on the task option that you select on the taskbar. The following table provides an overview of the information that appears in the summary panel when you select a task on the taskbar.

Table 141 Summary pane

Selected Task	Summary panel description	Summary panel column description
Known hosts	Hosts pane— Displays a list of NetWorker hosts in the datazone that have an associated Client resource on the NetWorker server.	• Hostname—The name of the NetWorker host as it appears in the Name attribute of the NetWorker Client resource. • OS—The operating system of the client as it appears in the OS attribute of the NetWorker Client resource. The operating system attribute appears blank until you have performed one successful backup operation for the host or performed an inventory operation. • NetWorker version—The version of the NetWorker software on the host. This attribute appears blank until you have performed one successful backup operation for the host.
Software Inventory	Software pane— Displays information about the NetWorker software that is installed on known hosts in datazone. The information that appears in this	• Hostname—The name of the NetWorker host. • OS—The operating system of the host. • OS Platform—The operating system architecture of the host. • Package name—The names of the NetWorker packages that are installed on the host that you can use Package Manager to upgrade.

Table 141 Summary pane (continued)

Selected Task	Summary panel description	Summary panel column description
	view is based on information that is gathered during the last inventory operation. You can only run an inventory operation after you add software into the software repository.	- Version—The version of the detected NetWorker software. - Upgrade available—Displays Yes when the software repository contains a version of the NetWorker software that you can upgrade on the client.
Software Repository	Repository pane—Displays information about the NetWorker packages that are contained in the NetWorker software repository.	- Software—The name of the NetWorker software in the software repository. - Version—The version of the NetWorker software package. - Package Name—The name of the NetWorker package. - OS—The operating system for the package. - OS Platform—The OS architecture for the package. - Size—The size of the NetWorker package.

The *NetWorker Updating from a Previous Release Guide* describes how to use the Software Inventory and Software Repository panes to update the NetWorker software on known hosts.

The task monitoring panel is always visible for each task option. A splitter separates the task monitoring panel from the summary panel. You can click and move the splitter to resize the task monitoring panel.

The task monitoring panel contains three window panes:


- Software Operations—Displays status information about operations that are performed for each task option.
- Log—Displays the most recent notification logs for the NetWorker server.
- Alerts—Displays alerts that are generated by a NetWorker server.

[Monitoring NetWorker Server activities in the Administration window](#) on page 55 provides detailed information about the **Log** and **Alerts** window panes.

Windows client interface

The NetWorker User program provides the ability to manage clients in the NetWorker environment.

The following figure illustrates the Windows client interface.

Figure 87 NetWorker User program

Starting the NetWorker User program on Windows

There are two ways to start the NetWorker User program.

- Click the Windows **Start** button and select **Programs > EMC NetWorker > NetWorker User**.
- From the **Administration** window, click **Start** on the main menu, and select **NetWorker User...** If the NetWorker Module for Microsoft Applications (NMM) is installed on the client computer, this operation starts NMM instead. The NetWorker client package must be installed on the host where you start the NetWorker User program. Otherwise, you see an error message similar to the following:

The user program you are trying to run (winworkr) is either not installed on this computer, or is not in your path.

To start the NetWorker User program, you must belong to the appropriate Windows groups. The following table lists the groups that you must belong to in order to run the NetWorker User program.

The Backup Operators and Administrators groups are the local and remote Microsoft security groups.

Table 142 NetWorker User Groups requirements

Logged in	Workstation	Server	Server (domain controller only)
Locally	Backup Operators or Administrators	Backup Operators or Administrators	Not applicable
To the domain	Domain Administrators	Domain Administrators	Backup Operators or Administrators

Toolbar buttons

The NetWorker User program has a toolbar with buttons for common User program tasks. The following table describes the function of each button.

Table 143 NetWorker User toolbar functions

Button	Name	Function
	Backup	Starts a manual (unscheduled) backup of the client's data to a NetWorker server.
	Recover	Starts a recovery operation to retrieve copies of saved data back to the client computer.
	Archive	Starts an archive operation to save copies of data to a server for storage on an archive volume. Once the data is stored on the archive volume, you have the option of removing the data from the disk.
	Verify	Starts a verification operation to ensure that the data items that were just backed up are the same as the data items that are currently on the disk.

Browse window


If you select menu items or buttons, a **browse** window opens in the NetWorker User program .

A **browse** window opens if you select any of the following items:

- A toolbar button.
- A Backup, Recover, Archive, Verify, or Local Directive command from the NetWorker User File menu.

The **browse** window, which is shown in the following figure, displays the directory tree of the file system that is being browsed.

Figure 88 Example of the browse window



NOTICE When you mark a disk volume or directory for an operation, all its nested subdirectories and files are also marked.

A check mark beside an item name indicates that the item is selected for backup, recovery, archiving, or verification.

The **Attributes** column indicates any special handling option that was applied:

- P — The item is marked for password-protection.
- E — The item is marked for password-protection and encryption, using the PW2 ASM.
- C — The item is marked for compression.

Connecting to a NetWorker server

A typical user that runs the NetWorker User program needs to connect to the NetWorker server that performs scheduled backups. However, to perform a directed recovery or to back up files to another server, you might need to connect to a different NetWorker server.

About this task

Before the NetWorker User program can connect to a NetWorker server, the client computer must be set up as a Client resource on that NetWorker server:

Procedure

1. From the **Operation** menu, select **Change NetWorker Server**.
2. In the **Change Server** dialog box, select a server from the list of available NetWorker servers. If the server is not listed, do one of the following:
 - Click **Update List** to search the network for available NetWorker servers.
 - Type the server's hostname.
3. Click **OK**.

Editing a client NSRLA database

The NetWorker Client database, `nsrexec` contains a NetWorker resource, called the NSRLA resource. The NSRLA resource contains information about the client and

attributes that you can modify. Use the character-based `nsradmin` program to modify the NSRLA resource.

Procedure

1. Log in as root or as Windows Administrator on the NetWorker Client.
2. Type the following at the command prompt:

```
nsradmin -p nsrexec
```

The `nsradmin` prompt appears.

3. To determine the current settings for attributes in the NSRLA resource, perform the following two steps:
 - a. To determine the current settings for any hidden attributes (such as the Disable Directed Recover attribute), type the following at the `nsradmin` prompt:

```
option Hidden
```

- b. To display attributes, type the following:

```
print type:NSRLA
```

4. To change the value of attributes in the NSRLA resource, type the following line at the `nsradmin` prompt:

```
update attribute:value;
```

For example, to update the Disable Directed Recover attribute, type the following:

```
update disable directed recover:Yes
```

5. Type **Yes** when prompted to confirm the change.

Results

NOTICE When you modify an attribute with the `nsradmin` program, you must specify the attribute name and value correctly. If you do not specify the attribute name and value correctly, the `nsradmin` program does not update the attribute and `nsradmin` does not provide an error message.

The *NetWorker Security Configuration Guide* provides more information about the `nsrexec` database and how to modify attributes in the `nsrexec` database.

CHAPTER 16

Restricted Datazones

This chapter contains the following topics:

• Restricted Datazones overview	832
• Administrators and users of RDZ	832
• Setting up the RDZ	833
• Removing a resource association	839
• Backward compatibility	839

Restricted Datazones overview

Restricted Datazones (RDZ) allows NetWorker administrators to organize a NetWorker environment into a multi-tenancy configuration, providing the ability to add an extra layer of privilege control.

This additional layer of control allows you to isolate access to resources, and separate these restricted resources into specific groups. RDZs also provide the ability to set up communal, or shared, resources (resources that are not owned by a specific RDZ).

Restricted and shared resources

A restricted resource (a resource that is owned by an RDZ) can only be operated by users within the RDZ who have the appropriate privileges, and by the global administrator. Restricted resources can reference both other restricted resources within the same RDZ, and shared resources.

A shared resource can be operated on by any RDZ, but only modified by global administrators. Shared resources can only reference other shared resources (for example, a shared client can only reference other shared directives).

Resource type associations

You can associate the following resources to an RDZ:

- Clients
- Protection policies
- Protection groups
- Directives
- Labels
- Pools
- Jukeboxes
- Status of operations (for example, jukebox actions)
- Devices
- Storage nodes
- Scheduled recovery

Administrators and users of RDZ

The following section identifies the administrator and user roles in relation to the RDZ feature.

- **Global Administrator**—A full administrator that has access to all resources. This user is equivalent to a traditional NetWorker administrator. Global administrators oversee the setup and management of several RDZs and determine the access tenant administrators have. A Global Administrator is the only user who can set up the users and privileges of an RDZ. A full administrator may have access to all datazones.
- **Tenant/Restricted Administrator**—An administrator that exists only in the RDZ to which they are assigned, and therefore has a limited view and operation of NetWorker. A tenant administrator can only manage NetWorker resources within their assigned restricted datazone, although since this user has the Monitor NetWorker privilege they can also view shared resources. You cannot associate a tenant administrator with more than one RDZ, however you can associate with

more than one instance of the same RDZ. By using multiple instances of an RDZ, the global administrator can divide and assign specific tasks and privileges among the tenant administrators and users of that RDZ.

- **Tenant/Restricted User**—A user that exists only within the RDZ to which they are assigned, and who has no administrative privileges in that datazone. NetWorker does not support a tenant user in two RDZs simultaneously.

Administrator roles

Management and use of RDZs is divided among global administrators and tenant administrators. A global administrator creates and manages RDZs. The global administrator can perform all the RDZ tasks, or associate specific tasks and privileges within each RDZ to one or more RDZ users as tenant administrators.

Although there are many possibilities for the roles of administrators, most setups fall into the following two approaches:

- Global administrator sets up the initial configuration, and also configures everything, so there is no need for a tenant administrator. This approach may be preferred for a customer with a very large environment, where one individual controls the network and sets up RDZs for various divisions within their company.
- Global administrator sets up the initial configuration, and tenant Administrators can configure and operate clients and create, view, operate, manage, and modify the NetWorker resources that are associated with their own RDZ according to the privileges assigned by the global administrator. Controls can be put in place to limit a tenant administrator's impact on the server. The global administrator can restrict the NetWorker resources that each RDZ can use, such as the maximum number of clients, devices, jukeboxes, or storage nodes.

Using multiple instances of an RDZ

You can give different RDZ users different privileges or levels of access within the same RDZ. This is done by creating multiple instances of an RDZ.

For example, you can create one RDZ instance for a tenant administrator to configure the RDZ resources. You can then create a second instance of the same RDZ for a tenant user to only monitor the RDZ resources, but not to modify the RDZ resources. In this way, different NMC users are given different levels of access within the same RDZ.


The global administrator can create multiple instances of an RDZ. They can create an RDZ multiple times using the same RDZ name and the same restrictions (number of clients, number of devices, and so on). NetWorker propagates all the information in each instance (except for the tenant administrators, privileges, and comments) to all the instances of the RDZ that have the same name.

Setting up the RDZ

A NetWorker administrator or global administrator can set up the RDZ in the **Server** window of NMC.

An entry for **Restricted Data Zones** appears in the left navigation pane, as shown in the following figure.

Figure 89 Restricted Data Zones in NMC



Setting up RDZ Users

You can set up Users in the Restricted Datazone resource the same way as you would in the User Group resource, with the same set of privileges to choose from.

If you do not use the **External Roles** attribute, these are normal users. Privileges for the most part only apply to the resource they are associated to, excepting shared resources, which can be seen if the user has Monitor NetWorker privileges.

Note that privileges are additive. If you have a privilege in a User Group resource, that applies to everything, including users who are also simultaneously inside an RDZ. By default, users can see all resources in the User Group resource. You must ensure you modify the User Groups very carefully to make views more restrictive.


An NMC user account can only be assigned to one RDZ at a time. However, NetWorker might not prevent this from happening. Therefore, be cautious when setting up accounts, particularly when using **External Roles**, where an account might potentially overlap two RDZs. Having the same NMC user account in multiple RDZs results in unpredictable behavior and is not supported.

Adding a User

Add a user to the RDZ to allow them to do administrative tasks within the RDZ by right-clicking Restricted Datazones in the **Server** window and selecting **Properties**.

In the **User Configuration** section of the window, click the **+** next to **External roles** to add a group that contains a user, and check the privileges that this user has.

Figure 90 Restricted Datazone User Configuration



Note: Wildcard characters such as an asterisk (*) are not permitted.

Setting up an RDZ resource

The following procedure shows RDZ resource setup for a client, but you can use this procedure when setting up any type of RDZ resource.


Before you begin

Before you create an RDZ, you must first review the NMC User Groups.

- In the NMC **Administration** window, under **Server**, select **User Groups**.
- Review the properties of each existing **User Group** to check which **External Roles** and **Users** are specified, along with the respective privileges.
- By default, the **Users** user group contains all users ("*@"), and if left unmodified, any RDZ user that you create inherits those privileges for all shared resources.
- To prevent RDZ users from gaining unwanted access to shared resources, you must either restrict the **External Roles** or **Users** for the **Users Group** (or any others), or alternatively, uncheck some privileges.

Procedure

1. In the **Server** window, right-click **Restricted Data Zones** and select **New** .
The **Create Restricted Data Zone** window appears.

Figure 91 Create Restricted Data Zone in the NetWorker Administration Server window

2. Create the RDZ (for example, RDZ1) by naming the RDZ and specifying any restrictions.

Use the **Restrictions** subsection to set limits on the clients, devices, storage nodes, and jukeboxes that can be owned by the restricted datazone to prevent resource abuse and limit what the tenant administrator can create. Setting restrictions can provide more control for major events that may impact the server, licensing limitations, and so on. These restrictions are in place even if using the RDZ as a global administrator.

Note: Setting a resource restriction to a value of 0 indicates that the user cannot create this resource.

If you plan to have a tenant administrator or tenant user, do the following:

- a. Specify the user using the **External Roles** or **Users** parameters.
- b. Under **Privileges**, select the privileges that the user will have in the RDZ. Typically, an RDZ tenant administrator will have all privileges, while a tenant user will have limited privileges.

You can specify multiple users. However, they will each have the same privileges. To set up separate users with different privileges within the same RDZ, create a separate instance of the RDZ by repeating the above steps using the same RDZ name and restrictions, but with different External Roles, Users, and Privileges.


If the global administrator is going to administer the RDZ and if there are no RDZ users, then the **User Configuration** section of the **Create Restricted Data Zone** window can be ignored.

3. In the left navigation pane of the **Protection** window, right-click the desired resource (for example, Client) and select **Properties** to configure the resource with the RDZ. Note that in addition to using an existing resource, you can also create the resource for the RDZ.

Resources that you can associate to an RDZ displays a **Restricted Datazone** tab in NMC (or the **Restricted Datazone** attribute in `nsradmin`).


4. Select the **Restricted Datazone** tab. Resources automatically get associated to the Restricted Datazone a user belongs to when they create a resource.

Figure 92 Restricted Data Zone Client Properties



5. (Optional) In the left navigation pane of the **Devices** window, select a device if you want to give this RDZ client access to a specific device by right-clicking the device and selecting **Properties**. Give the RDZ client access to this device.


Figure 93 Restricted Data Zones in Device Properties window



Note: The RDZ can access the shared devices without any further device setup requirements for the shared resource if these devices are configured. Note, however, that multiple RDZs cannot simultaneously access the same device.


6. Create a policy. In the **Protection** window, right-click **Policy** in the left navigation pane and select **Create Policy**.
7. In the **General** tab, specify a name for the policy (in this example, RDZ1). In the **Restricted Data Zone** tab, select the RDZ from the drop-down, then click **OK**.

Figure 94 Restricted Data Zones in Create Policy window



8. Create a group. In the **Protection** window, right-click **Group**, and select **New**.
9. Name the group and select the desired client(s). In the **Restricted Datazone** tab, select the RDZ from the drop-down, then click **OK**.
10. In the **Protection** window, highlight the new policy and create a workflow. Associate this workflow with the new group.

Figure 95 New workflow associated with RDZ group



Note: You only must configure the policy and group resources for RDZ when using policies. The workflows and actions that are created as a result of it are kept within the policy feature and do not need any further RDZ configuration.

11. Return to the **Server** window and highlight Restricted Datazones. Right-click and select **Properties**.
12. Browse to the **Clients** tab. This tab now shows the clients that are associated with this RDZ. When a user belongs to an RDZ and creates a resource, this resource is automatically set to being owned by that RDZ.

Removing a resource association

You can remove a resource association in two ways:

- By deleting the resource itself. This involves deleting multiple instances of a resource if there are two or more resources with the same name.
- By unselecting the Restricted Datazone in the respective attribute.

Backward compatibility

RDZ is a feature of the server and storage node, so the client does not necessarily have to be upgraded to match the server version. RDZ is backward compatible with the NetWorker client if that client is supported with the NetWorker 19.1 server.

CHAPTER 17

Block Based Backup and Recovery

This chapter contains the following topics:

- [Overview](#) 842
- [Block based backups](#).....846
- [Block based recoveries](#).....851
- [Troubleshooting block based backup and recovery issues](#).....858

Overview

The NetWorker block based backups are high-performance backups which are supported on Windows and Linux.

During block based backups, the backup application scans a volume or a disk in a file system, and backs up all the blocks that are in use in the file system. Block based backups use the following technologies:

- The Volume Shadow Copy Service (VSS) snapshot capability on Windows and Logical Volume Manager (LVM) and Veritas Volume Manager (VxVM) on Linux to create consistent copies of the source volume for backups.
- The Virtual Hard Disk (VHDx), which is sparse, to back up data to the target device.

Block based backups support only the following Client Direct enabled devices as target devices:

- Advanced File Type Devices (AFTDs)
- Data Domain devices
- Cloud Boost devices

The block based incremental backups use the Change Block Tracking (CBT) driver to identify the changed blocks, and back up only the changed blocks.

Block based full and incremental backups are fast backups with reduced backup times because the backup process backs up only the occupied disk blocks and changed disk blocks respectively. Block based backups can coexist with traditional backups.

Block based backups provide instant access to the backups. The block based backups enable you to mount the backups by using the same file systems that you used to back up the data.

Block based backups provide the following capabilities:

- Mounting of a backup as a file system
- Mounting of an incremental backup
- Sparse backup support
- Backups to disk-like devices
- Backups of operating system-deduplicated file systems as source volumes on Windows
- Forever virtual full backups to Data Domain
- Data Domain retention lock
- 38 incremental backups to AFTD and Cloud Boost devices
- Synthetic full backups to AFTD and Cloud Boost devices
- Backups of volumes up to 63 TB each
- NetWorker-supported devices as secondary devices for backups
- Recoveries from Data Domain without using CIFS share
- Recovery of multiple save sets in a single operation
- Setting parallel save streams if the target or destination is Data Domain

The following table lists the backup scenarios and the recovery scenarios that block based backups support.

Table 144 Supported backup and recovery scenarios

Backup scenarios	Recovery scenarios
• Full backups • Virtual full backups • Synthetic full backups • Incremental backups • Full backups and incremental backups intermixed with built-in provisions to anchor the incremental backups with an appropriate backup type	• File level recovery by mounting the backup image on a target host • Image/destructive recovery at the block level • Image/destructive recovery from clones • Windows Bare Metal Recovery (BMR) by using a WinPE image

Supported operating systems and configurations

NetWorker supports block based backup and recovery of the following operating systems and configurations:

- Operating systems on x64:
 - Windows:
 - Windows client 8.1
 - Windows client 8
 - Windows Server 10
 - Windows Server 2016
 - Windows Server 2012 R2
 - Windows Server 2012
 - Windows Server 2016
 - Windows Server 2019
 - Linux:
 - Red Hat Enterprise Linux (RHEL) 6.0
 - RHEL 6.1
 - RHEL 6.2
 - RHEL 6.3
 - RHEL 6.4
 - RHEL 6.5
 - RHEL 6.6
 - RHEL 6.7
 - RHEL 6.8
 - RHEL 6.10
 - RHEL 7.0
 - RHEL 7.1

- RHEL 7.2
 - RHEL 7.3
 - RHEL 7.4
 - RHEL 7.5
 - RHEL 7.6
 - SuSE Linux Enterprise Server (SLES) 11 SP4
 - SLES 11 SP2
 - SLES 11 SP3
 - SLES 11 SP4
 - SLES 12
 - SLES 12 SP2
 - SLES 12 SP3
 - Community Enterprise Operating System (CentOS) 6.0
 - CentOS 6.1
 - CentOS 6.2
 - CentOS 6.3
 - CentOS 6.4
 - CentOS 6.5
 - CentOS 6.6
 - CentOS 6.7
 - CentOS 6.8
 - CentOS 6.10
 - CentOS 7.0
 - CentOS 7.1
 - CentOS 7.2
 - CentOS 7.3
 - CentOS 7.4
 - CentOS 7.5
 - Ubuntu 14.04 with kernel 3.13.0-24
 - Ubuntu 16.04 with kernel 4.4.0-21
- Operating systems on x86:
 - Windows client 8.1
 - Windows client 8
- File systems:
 - Windows:
 - New Technology File System (NTFS)
 - Resilient File System (ReFS)
 - Linux:
 - Third extended file system (ext3)

- Fourth extended file system (ext4)
 - Client Direct target devices
 - Concurrent backups of multiple volumes
 - Windows Server 2012 and Windows Server 2012 R2 deduplicated volumes without rehydrating the deduplicated data
 - Windows Server core installation role
 - Unified Extensible Firmware Interface (UEFI) based systems
 - GUID Partition Table (GPT) and Master Boot Record (MBR) volumes
 - Data Domain systems in a Fibre Channel environment
 - Full backup of Windows Server 2012 Cluster Shared Volumes on File Servers and Windows Clusters
 - LVM2 and VxVM managed volumes on Linux
-  **Note:** Each volume group on LVM2 and VxVM must have at least 10% free space for a block based backup to succeed.

Limitations

NetWorker's block based backups and recoveries do not support the following capabilities and configurations:

- FAT32 file system
In the case of the ALL save set backups, either unmount or remove the FAT32 volumes, and perform the backups.
- Live updates and service patches for Ubuntu 14.04 and 16.04
- Backup levels 1 through 9
- Backups of Microsoft 2012 clusters without Cluster Shared Volumes
- Incremental backups of Microsoft clusters
- Cloning of AFTD incremental backups
- Granular save sets at either the folder level or the file level, for example, D:\data
- Checkpoint restart
- Standard NetWorker directives
- The scanner command with the `-i` option for rebuilding indexes for block based backups
- Staging and the `nsrclone` command with the `-m` option for migrating block based backup save sets to other volumes
- Image recovery to a system volume
- Recoveries of ReFS volumes on Windows Server 2008 R2 SP1 and Windows 8 (x86 and x64)
- Recoveries of Windows deduplicated volumes to Windows Server 2008 R2 SP1 and Windows 8 (x86 and x64)
- Troubleshoot kernel on RHEL, and Trace and Xen kernels on SLES

Block based backups

This section provides information about block based backups (BBB).

Devices for block based backups

You must create a backup device and configure block based backups before you perform block based backups and recoveries.

You can create the following types of devices that depend on the backup requirements:

- AFTD
- Data Domain CIFS or NFS
- DD Boost
- Cloud Boost



Note:

For block based backups to succeed, ensure that you meet the following requirements:

- Create a separate pool.
- The pool must contain only one backup device.
- Perform all backups of a client to the same backup device.

If you want to make a local AFTD a Client Direct enabled device, specify either the CIFS path or the NFS path in the **Device access information** field of the **Create device properties** dialog box.

The "Backup Target" chapter describes how to configure devices.

Installing the lgtobbb package on Linux

About this task

You must install the `lgtobbb` package, which is packaged along with NetWorker, for block based incremental backups and recoveries to succeed on Linux. If you do not install the package, block based full backups succeed, but incremental backups and recoveries fail.

Procedure

1. Ensure that the NetWorker client is installed.
2. Install the `lgtobbb` package:
 - On RHEL:
 - a. Ensure that the `lsb` package from the operating system installation media is installed.
 - b. Run the following command:


```
rpm -ivh lgtobbb-18.0.0.0-1.x86_64.rpm
```
 - On SLES:
 - a. Ensure that the `lsb-release` package from the operating system installation media is installed.

b. Run the following command:

```
rpm -ivh lgtobbb-18.0.0.0-1.x86_64.rpm
```

- On Ubuntu:

a. Ensure that the following packages are installed:

- Shells: ksh and pdksh
- C++ library: libstdc++5
- gawk

b. Run the following command:

```
dpkg -i lgtotdclnt_99.0.99.8228_amd64.deb
```

Configuring block based backups

About this task

When you configure block based backups, consider the following notes:

- If you select the **Block based backup** option, and do not select the **Parallel save streams per save set** option, block based backups without parallel save streams are performed.
- If you select the **Parallel save streams per save set** option, and do not select the **Block based backup** option, file level backups with parallel save streams are performed.
- If you do not select the **Block based backup** and **Parallel save streams per save set** options, file level backups without parallel save streams are performed.

Procedure

1. Enable the block based backup feature when you use one of the following methods to configure the client:
 - NetWorker **Client Configuration** wizard
 - **Client Properties** window
 - The `nsradmin` program
2. Select the following fields to enable the block based backup feature:
 - **Client direct** (selected by default)
 - **Block based backup**
3. [Optional] To set parallel save streams if the target or destination is Data Domain:
 - a. In the NetWorker **Client Properties** dialog box, on the **Globals (1 of 2)** tab, select **Parallel save streams per save set**.
 - b. On the **Apps & Modules** tab, in the **Save operations** field, type one of the following values:
 - PSS:streams_per_ss=2
 - PSS:streams_per_ss=4
This is the default value.

 **Note:** Parallel save streams is applicable only for full backup.



Note: Consider the following notes about parallel save streams and save sets:

- When a backup contains more than four save sets, the parallel save streams value must be either greater than or equal to the number of save sets.
- The parallel save streams per save set value is the same for all the save sets of the client, that is, you cannot set the value of one save set to 2 and the value of another save set to 4 on the same client.

[Optional] Creating an AFTD CIFS share on Windows for block based recoveries

You must enable a CIFS share to access save sets on the device to recover data from an AFTD. The access credentials are the same as the administrator's credentials on the host.

Procedure

1. Right-click the folder that you want to share, and select **Share with > Specific people....**
2. In the **File Sharing** dialog box, select or add the people with whom you want to share the folder, and click **Share**.

Performing block based backups

The procedure for performing a block based backup is the same as the procedure for performing a NetWorker backup.

[Backing Up Data](#) on page 409 provides more information about how to back up data by using NetWorker.

You can perform a block based backup as any of the following types of backup:

- Scheduled backups
- Incremental backups
- Virtual full backups
- Synthetic full backups
- Manual backups or client-initiated backups
- Save set backups
- Exclude list backups
- Windows deduplication volume backups
- CSV backups
- Windows BMR backups

Scheduled backups

NetWorker supports block based backups for all scheduled backups.

The scheduled backup process is transparent to you and does not require any additional actions or considerations.

Incremental backups

You must perform an incremental backup of a volume only to the same device, to which a full backup of the volume was performed.

Note: Incremental backups can span across multiple storage units on the same Data Domain device.

On AFTDs, selecting any backup level apart from full or incremental results in performing an incremental backup.

An incremental backup shifts to a full backup when any of the following conditions occur:

- You restart the client host for any reason when the backup is either in progress or scheduled.
- The preceding incremental backup failed.
Note: This condition applies only to Windows. On Linux, an incremental backup continues even if its preceding incremental backup failed.
- You already performed 38 incremental backups to AFTD.
Note: After you perform a full backup, you can perform a maximum of 38 incremental backups.
- You add a volume for the backup of the ALL save set.
- You change the size of the volume.

The incremental backup process is transparent to you and does not require any additional actions or considerations.

Virtual full backups

Virtual full backups apply only to the Data Domain devices. When you perform an incremental backup to a Data Domain device, you perform the backup as a virtual full backup. However, the type of the backup that you have performed is displayed as full. A virtual full backup backs up only the changed blocks from its previous full backup while referencing the unchanged blocks to the corresponding blocks of the previous full backup.

Note: On Data Domain devices, selecting any backup level apart from full results in performing a virtual full backup.

Synthetic full backups

The synthetic full backups apply only to AFTDs. A synthetic full backup consolidates data from all the existing full and incremental backups.

Note: When you perform a synthetic full backup to a non-Windows remote storage node, you must create a client configuration for the storage node.

Manual backups or client-initiated backups

Use the `save` command with the `-z` option to perform a client-initiated block based backup from the command line.

Ensure that you meet the following requirements for a client-initiated backup:

- The device must be Client Direct enabled.
You can provide a pool of Client Direct enabled devices by using the `save` command with the `-b` option.

- The client-initiated block based backup supports the full level save sets that you define only at the volume level.
- Do not name a manual snapshot with the same name as the block based backup snapshot.
If a block based backup snapshot and a manual snapshot have the same name, performing the manual snapshot deletes the block based backup snapshot.

Save set backups

You can use a block based backup to back up the following save sets:

- Windows:
 - **ALL**—This save set includes VSS volumes, critical volumes, and non-critical volumes.
 - **DISASTER_RECOVERY**: —This save set includes VSS volumes and critical volumes.
 - **Volumes**—Specify any type of volume drive letters as save sets. For example:
D:\
 - **Volume mount points**—Specify volume mount points as save sets. For example:
D:\mount_point_name (for a single mount point)
D:\mount_point_name1\mount_point_name2\mount_point_name3
(for nested mount points)
- Linux:
 - **ALL**—All the mounted volumes that the `/etc/fstab` file lists.
Note: If the ALL save set contains block based backup supported and unsupported volumes, and you specify parallel save streams for the save set, the following types of backups will be performed:
 - Block based backup of the supported volumes, each with random parallel save streams, but not with the parallel save streams that you specified
 - NetWorker traditional backup of the unsupported volumes
 - **Volume mount points**—Specify volume mount points as save sets. For example:
/<mount_point_name> (for a single mount point)
/<mount_point_name1/mount_point_name2/mount_point_name3>
(for nested mount points)

Windows deduplication volume backups

The block based backups occur at the block level. The file system layout does not affect the backup. The backup virtual hard disk is deduplication in nature. The block based backups merge the blocks out of the deduplication volumes. In case the volume changes from deduplication to non-deduplication, the block based backup detects these events and forces the next backup to be a full backup.

CSV backups

You can simultaneously see Cluster Shared Volumes (CSVs) across all nodes. The block based backups support only full backups of CSVs, even in the case of a failover.

If you try to perform an incremental backup, the backup shifts to a full backup with a warning message.

Windows BMR backups

The procedure for performing a block based backup as a Windows BMR backup is the same as the procedure for performing a NetWorker Windows BMR backup. However, you must select the block based backup option when you configure the client using the NetWorker **Client Configuration** wizard, the **Client Properties** window, or the `nsradmin` program.

Verifying block based backups

Procedure

1. To list the block based backup save sets, run the following command:

```
mminfo -avot -q "ssattr=*BlockBasedBackup"
```

To list the block based virtual full backup save sets, run the following command:

```
mminfo -avot -q "ssattr=*BlockBased Virtual Full"
```

To list the block based synthetic full backup save sets, run the following command:

```
mminfo -avot -q "ssattr=*Synthetic full"
```

2. Verify whether all the selected save sets have been successfully backed up.

Cloning block based backups

The procedure for cloning a block based backup is the same as the procedure for cloning a NetWorker backup.

You can configure the NetWorker clone operations according to the environment and storage requirements. Block based backups support cloning of the full and incremental backups .

Block based recoveries

This section provides information about block based recoveries.

Preparing for block based recoveries

You must be familiar with the recovery operations, workflows, and interfaces that associate with the block based recovery. Use either NMC or the NetWorker command-line interface (CLI) to perform a block based recovery.

You typically complete the following tasks to perform a recovery by using NMC:

1. Selecting the save set.
2. Performing either file level recovery or image/destructive recovery.

If you want to perform a recovery by using the CLI, you must run the `recover.exe` command with the save set ID. Unlike a traditional backup, the block based backup does not maintain any indexes in the NetWorker client file index database.

The recovery process mounts all the save sets on a device that supports the Client Direct functionality.

If you want to recover data from either an AFTD or a Data Domain device by using the CIFS or NFS share, enable the CIFS or NFS share to access save sets on the device.

Performing block based recoveries

You can perform block based recoveries by using either NMC or the NetWorker CLI.

Using NMC to perform block based recoveries

Procedure

1. Open NMC.
2. Click **Recover**.
3. From the menu bar, select **Recover > New Recover**.
4. On the **Recovery Hosts** page:
 - a. Under **Source Host**, in the **Name** field, type the name of the host on which the backed-up data exists.
 - b. Under **Destination Host**, specify the host to which you want to recover the backed-up data.
 - c. Under **Available Recovery Types**, select **Block Based Backup**.
 - d. Click **Next**.
5. On the **Select the Data to Recover** page:
 - a. Select one of the following types of recovery that you want to perform:
 - File level recovery
 - Image level recovery
 - b. Select the timestamp of the backup that you want to recover.
 - c. Perform one of the following tasks that depend on the type of the recovery that you have selected:
 - For a file level recovery:
 - In the left panel, select the save sets that you want to recover.
 - In the right panel, select the relevant files that you want to recover.
 - For an image level recovery, in the left panel, select the save set that you want to recover.
 - d. Click **Next**.
6. On the **Select the Recovery Options** page, perform one of the following tasks that depend on the type of the recovery that you have selected:
 - For a file level recovery, select a file path for recovery and an appropriate option for duplication, and click **Next**.
 - For an image level recovery, select a file path for recovery, and click **Next**.
7. On the **Obtain the Volume Information** page, click **Next**.
8. On the **Perform the Recovery** page:
 - a. Under **Identity**, in the **Recover name** field, type a name for the recovery.
 - b. Select one of the following recovery start times:

- **Start recovery now**—Immediately starts the recovery.
- **Schedule recovery to start at**—Schedules the recovery according to the choice.

- If you want to stop the recovery at a certain time, in the **Specify a hard stop time** field, type the time.
- Select the **Recover Resource persistence** option according to the choice.
- Click **Run Recovery**.

The recovery log appears when the recovery progresses.

After the recovery succeeds, a successful completion message appears at the bottom of the recovery log.

To export the log file, click **Export Log File**.

- On the **Check the Recovery Results** page, click **Finish**.

Using the CLI to perform block based recoveries

Use the `recover.exe` command to perform a block based recovery. The command applies only to local clients. However, you cannot use the command to perform a remote or redirected recovery.

Performing file level recoveries

About this task

NOTICE For Windows hosts only, to ensure that you use the NetWorker `recover.exe` command and not the Windows OS `recover` command, perform one of the following tasks:

- Ensure that `NetWorker_install_path\bin` appears before `%SystemRoot%\System32` in the `$PATH` environment variable.
- When you start the `recover` command include the path to the binary. For example: `NetWorker_install_path\bin\recover.exe`.

Procedure

- On Windows:
 - Run the following command to mount the backup and start the command prompt at the mount point:

```
recover.exe -w -S <save_set_ID>
```

Use the Windows copy option and paste option to recover the backup.

After you perform the recovery, close the command prompt to exit the process.

- Run the following command to mount the backup and copy specific files from the input file to the destination:

```
recover.exe -w -S <save_set_ID> -I <input_file> -d  
<destination>
```

- On Linux:

Ensure that you meet the following prerequisites before you perform a file level recovery:

- a. You have disabled Security-Enhanced Linux (SELinux) by running one of the following relevant commands:
 - `setsebool -P nis_enabled 1`, if you use either RHEL 7.x or CentOS 7.x
 - `setsebool -P allow_ypbind 1`, if you use either RHEL 6.x or CentOS 6.x
- b. You have installed the **iscsiadm** utility by installing one of the following relevant packages on the Linux client:
 - `iscsi-initiator-utils<version_number>.rpm`, if you use either RHEL or CentOS
 - `open-iscsi<version_number>.rpm`, if you use SLES
- c. On SLES, if you want to start the **iscsiadm** utility for the first time, restart the iSCSI services by running the following command:
`service open-iscsi restart`

Perform a file level recovery:

- a. Run the following command to mount the backup:

```
recover.exe -w -S <save_set_ID>
```

Open a new terminal, and use Linux copy and paste commands to recover the data.

After you perform the recovery, type `quit` to exit the process.

- b. Run the following command to mount the backup and copy specific files from the input file to the destination:

```
recover.exe -w -S <save_set_ID> -I <input_file> -d  
<destination>
```

 **Note:** For Windows, the command is `recover.exe`. For Linux, the command is `recover`

Performing image and destructive recoveries

Ensure that you meet the following requirements to perform a recovery:

About this task

- The size of the target volume is either the same or more than the size of the source volume.
- The cluster size of the source volume is the same or more than the cluster size of the target volume.
- The target volume is not a system volume.

Run the following command to perform an image recovery:

```
recover.exe -S <save_set_ID> -r <target_volume>
```

Command-line options for recover.exe

The following table describes the key options that you can use with the `recover.exe` command to perform a block based recovery.

Table 145 Key options for the block based recover.exe command

Option	Description
<code>-r [volume GUID or mount point]</code> (On Windows) <code>-r [raw device name or mount point]</code> (On Linux)	Specifies the supported destinations for save set recovery on Windows: • Volume name • Raw pathname • Volume GUID • Existing mount point Specifies the supported destinations for save set recovery on Linux: • Mount point • Raw device name
<code>-S [save set ID or clone ID]</code>	Specifies the save set ID or the clone ID that you want to recover.
<code>-I [input file]</code>	Specifies a file that contains a list of files that you want to recover. This is useful to perform the disaster and remote recoveries.
<code>-w</code>	Specifies the file level recovery of Block Based backup.

Performing Windows BMR

The procedure to recover a block based backup through a Windows BMR is the same as the procedure to perform a NetWorker Windows BMR. However, you must select an appropriate block based backup on the **Select System Recovery** page of the wizard when you perform the block based recovery.

Performing block based clone recoveries

You can recover cloned data from the Client Direct enabled devices and the Client Direct disabled devices.

Recovering data from Client Direct enabled devices

Client Direct enabled devices include AFTD, DD Boost, and Data Domain CIFS devices.

About this task

Use one of the following methods to recover the data:

- NMC
Perform the steps that the [Using NMC to perform block based recoveries](#) on page 852 section describes.
- NetWorker CLI

Run one of the following commands:

- `recover.exe -w -S save_set_ID/clone_ID`
for file level recoveries
- `recover.exe -S save_set_ID/clone_ID -r target_volume`
for image recoveries

Recovering data from Client Direct disabled devices

Client Direct disabled devices typically include tape devices.

The file level recovery process requires a Client Direct enabled device. The recovery process first temporarily stages the data to a Client Direct enabled device that you have selected and then recovers the data from the device. The retention period of the staged data on the Client Direct enabled device is three days. You can delete the data before the retention period lapses.

The image recovery process by using the GUI is the same as the process to perform a file level recovery. However, you can perform image recoveries directly from the tape devices without mounting the backup images by using the CLI.

NOTICE To perform either a file level recovery or an image recovery of data from a CloudBoost device, first clone the data to a Client Direct enabled device and then recover the data from the Client Direct enabled device.

Use either NMC or the `recover.exe` command to perform recoveries.

Using NMC to perform clone recoveries

Procedure

1. Open NMC.
2. Click **Recover**.
3. From the menu bar, select **Recover > New Recover**.
4. On the **Select the Client to Recover** page:
 - a. Under **Source client**, in the **Name** field, type the name of the client on which the cloned data exists.
 - b. Under **Destination client**, specify the client to which you want to recover the cloned data.
 - c. For the type of backup that you want to recover, select **Block Based Backup (cloned to tape)**.
 - d. Click **Next**.
5. On the **Select a Block-Based Backup Clone** page:
 - a. Under **Found in**, specify the period during which you performed the clone and click **Query**.

The cloned save set groups appear in the **Block-Based backups** field.

- b. Select the save set group.
- c. Under **Select the Save Sets**, select either **All save sets**, or **Subset of save sets** and appropriate save sets that belong to the selected save set group.
- d. Under **Recovery Type**:

Select one of the following types of recovery that you want to perform:

- **File level recovery**
If you have selected this option, from the **Copy to Pool** list, select the pool that has the Client Direct enabled device to which you want to copy the cloned data.
 - **Image level recovery**
- e. Click **Next**.
- If you have selected **File level recovery** in step d, the **Copying the Backup to Disk** page appears.
After the cloning succeeds, click **Next**.
The **Select the Data to Recover** page appears.
 - If you have selected **Image level recovery** in step d, the **Select the Data to Recover** page appears.
6. On the **Select the Data to Recover** page:
- a. Perform one of the following tasks that depend on the type of the recovery that you have selected:
 - For a file level recovery, select the save set to recover from the left panes and select the files to recover from the right panes.
 - For an image level recovery, select the save set that you want to recover from the left pane.
 - b. Click **Next**.
7. On the **Select the Recovery Options** page, perform one of the following tasks that depend on the type of the recovery you have selected:
- For a file level recovery, select the **File path for Recovery** and **Duplicate File Options**, and click **Next**.
 - For an image level recovery, select the **File path for Recovery**, and click **Next**.
8. On the **Perform the Recovery** page:
- a. Under **Identity**, in the **Recover name** field, type a name for the recovery.
 - b. Select one of the following recovery start times:
 - **Start recovery now**—Immediately starts recovery.
 - **Schedule recovery to start at**—Schedules the recovery according to the choice.
 - c. If you want to stop the recovery at a certain time, in the **Specify a hard stop time** field, type the time.
 - d. Select the **Recover Resource persistence** option according to the choice.
 - e. Click **Run Recovery**.
The recovery log appears when the recovery progresses.
After the recovery succeeds, a successful completion message appears at the bottom of the recovery log.
To export the log file, click **Export Log File**.
9. On the **Check the Recovery Results** page, click **Finish**.

Using the CLI to perform the clone recovery

Run one of the following commands to recover the data from the Client Direct disabled devices:

About this task

- For file level recoveries:

```
recover.exe -w -S save_set_ID/clone_ID -b pool_name
```

 **Note:** The pool that you select must have a Client Direct enabled device. The pool must also be a backup clone type pool.

- For image recoveries:

```
recover -S save_set_ID/clone_ID -r target_volume
```

Troubleshooting block based backup and recovery issues

This section lists the common issues with the block based backups and recoveries and provides workarounds for these issues.

Table 146 Troubleshooting block based backup and recovery issues

Error message or Issue	Resolution
Block based backups are only supported with Client Direct.	In the Client Properties dialog box, select Client Direct .
VSS OTHER: ERROR: VSS failed to process snapshot: The shadow copy provider had an unexpected error while trying to process the specified operation. (VSS error 0x8004230f) 90108:save: Unable to save the SYSTEM STATE save sets: cannot create the snapshot.	Ensure that there is no recover session running on the client.
No save sets clone to clone device.	Block based backups clone only full backup save sets. Block based backups do not clone incremental backup save sets.
Unable to construct the recover list from input file.	Perform an image recovery if applicable. Otherwise, select all the files except the system files such as, System Volume Information and Recycle Bin to perform a file level recovery.
Failed to recover save set with error: To perform the recovery of a block based backup save set, the device must be enabled for Client Direct.	In the Client Properties dialog box, select Client Direct .
Though the size of a target volume is more than the size of a source volume, after performing an image recovery, the target volume file system can use its volume only up	To enable the target volume file system to use its volume to the actual size, extend the file system:

Table 146 Troubleshooting block based backup and recovery issues (continued)

Error message or Issue	Resolution
to the same size as the size of the source volume.	1. In the command prompt, type <code>diskpart</code> and press <Enter>. 2. In the DISKPART command prompt, select the target volume to extend the file system by running the following command: <code>select volume <target_volume></code> For example, select volume G: 3. Extend the file system by running the following command: <code>extend filesystem</code> 4. Exit from the DISKPART command prompt by running the following command: <code>exit</code>
Block based backup failed partially, when parallel save streams is enabled: 9904:nsrmmdbd: access denied to media database, `SYSTEM' on `hostname' must have 'Operate NetWorker' or 'Change Application Settings' privilege 90096:save: save of 'E:\' to 'hostname' failed: access denied to media database, `SYSTEM' on `hostname' must have 'Operate NetWorker' or 'Change Application Settings' privilege	To perform parallel save streams enabled backups for block based backup volumes: 1. In NMC, under Server, select User Groups. 2. Right click Application Administrators and select Properties. 3. In the User Group Properties Application Administrators window, under Users, add <code>user=system,host='hostname'.</code>

CHAPTER 18

Networking and Connectivity

This chapter contains the following topics:

- [Name resolution and connectivity](#)..... 862
- [Troubleshooting name resolution and connectivity errors](#)..... 863
- [Using multihomed systems](#)..... 871
- [NIC Teaming](#).....877
- [Using DHCP clients](#)..... 878
- [NetWorker TCP/IP keep-alive parameters](#) 878

Name resolution and connectivity

A NetWorker host must consistently and reliably connect to, and resolve, each destination NetWorker host by fully qualified domain name (FQDN), shortname, and IP address.

The NetWorker software requires consistent and predictable forward and reverse name resolution to work correctly. From NetWorker version 18.1 and later, you must enable nsrauth to work seamlessly in an environment where reverse DNS entries are unavailable. All features of NetWorker and NetWorker modules are supported in this environment. NetWorker performs name resolution checks during the following operations:

- NetWorker daemon startup.
- Client and Device resource configuration.
- Backup, recovery, and device operations.

NetWorker relies on the operating system to perform the following tasks:

- Handle name resolution requests.
- Resolve hostnames to IP addresses (forward name resolution lookups).
- Resolve IP addresses to hostnames (reverse name resolution lookups).

NOTICE On Windows Server 2008 R2 SP1, EDNS0 queries increase the size of the DNS UDP packet and some firewalls block UDP packets larger than 512 bytes. It is recommended that you disable EDNSprobes on hosts that operate in a firewall environment, as a DNS Server or Domain Controller. To disable EDNSprobes, run the following command:

```
dnscmd /config /EnableEDNSProbes 0
```

NetWorker supports the use of Internet Protocol version 6 (IPv6) in a dual stack or in a pure IPv6 environment. NetWorker does not support NetWorker resource configurations that use temporary or link-local IPv6 addresses.

When a NetWorker host uses IPv6 addressing, ensure that you add the IPv6 address for the host in DNS Server or the hosts file and to the alias field in the client resource. The *NetWorker Installation Guide* provides information about using NetWorker in an IPv6 environment.

Note:

1. In a configuration without reverse DNS entries, NetWorker supports FQDN only. For information on how to use IP address, see the troubleshooting section.
2. Forward DNS entry is mandatory for every host present in the NetWorker datazone.
3. NetWorker Client, NetWorker Server, NetWorker Storage Node and NetWorker Management Console (NMC) must be using NetWorker 18.1 or later.
4. NetWorker will not change the behavior of third party applications interacting with NetWorker that mandates the requirement of reverse DNS lookups.

5. If the IP address or the short name is used for any of the NetWorker operations, then the `/etc/hosts` should be updated or ALIASES should be defined.

Troubleshooting name resolution and connectivity errors

When NetWorker operations fail due to name resolution issues, the following types of error conditions can appear in the `daemon.raw` file or in the savegroup completion report:

- RPC errors
- Unknown host errors
- Failures in contacting the portmapper
- Connection failures or time outs
- Unexpected exits by programs
- Connection refused errors
- Failure of a remote command (`rcmd()` function) to an active client
- Failures in name-to-address translation
- Program not registered errors
- Failures of NetWorker services to start
- Failures of NetWorker services to remain active
- Invalid path errors

When NetWorker operations fail due to name resolution issues, the following error messages can appear in the `daemon.raw` file or in the policy report:

- Host name for IP address *IP_address* could not be determined through DNS
- IP address for host '*hostname*' could not be determined through DNS
- Warning, cannot resolve host *hostname* to *IP_address*, name server may be misconfigured.
- '*Client_name*': Couldn't look up the name of address: '*NetWorker_server_IP*':node name or service name not known.
- nsrexec: nsrexecd on (*client*) is unavailable. Using rsh instead
- nsrexec: host *hostname* cannot request command execution permission denied
- Cannot connect to nsrexecd on client *NetWorker_server*.rhost permissions do not allow rsh permission denied

Before you can troubleshoot name resolution and connectivity issues, you must determine between which hosts the connection problems occurred. The problems can occur between any two types of NetWorker hosts, for example, between the NetWorker server and a client or between a client and a storage node.

Complete the following steps to troubleshoot name resolution and connectivity errors:

1. Document the steps you take and the results, especially error messages, in case you need to contact Customer Service.
2. Use operating system tools to confirm that basic connectivity exists between the source and destination hosts. For example, `telnet`, `ping`, and `traceroute`. [Verifying basic connectivity](#) on page 864 provides more information.

3. Check that the source and destination hosts consistently and correctly resolves all names and IP addresses for each host. [Verifying name resolution](#) on page 866 provides more information.
4. Verify that the configuration of the source and destination host includes all relevant information for each host in the Aliases attribute and the servers file. [Verifying the NetWorker configuration](#) on page 870 provides more information.

Verifying basic connectivity

NetWorker requires reliable and consistent connectivity between the source and destination hosts. Confirm that you can remotely connect to the host. When the source and destination hosts reside on different networks, verify the network connectivity between the hosts.

Verifying remote host connectivity

Try to connect to the host. If a backup fails for a NetWorker client, then try to connect to the client by using other tools. For example, try to connect to the host by using Remote Desktop Connection on Windows or the `telnet` command on UNIX. If remote connections to the host fail, then investigate external host connectivity issues.

Verifying network connectivity

Use the `ping` command and the `tracert` command on UNIX and Linux, or the `pathping` command on Windows, to transmit packets between hosts and verify that network connectivity exists between the source and the destination hosts. Run each command from the source host and destination host and use each command with the shortname, FQDN, and the IP address of the destination host.

About this task

In the following example, the source host `mnd.corp.com` is a Linux host with the IP address `10.1.1.10`. The destination host `pwd.corp.com` is a Windows host with the IP address `10.1.1.20`.

Procedure

1. On the `pwd.emc.com` host, run the following `pathping` commands:

```
pathping pwd.corp.com
pathping pwd
pathping 10.1.1.20
pathping mnd.corp.com
pathping mnd
pathping 10.1.1.10
```

A successful `pathping` command displays the following information:

```
C:>pathping mnd.corp.com
Tracing route to mnd.emc.com [10.1.1.10]
over a maximum of 30 hops:
 0 pwd.corp.com [10.1.1.20]
 1 mnd.corp.com [10.1.1.10]
Computing statistics for 25 seconds...
Source to Here This Node/Link
Hop RTT Lost/Sent = Pct Lost/Sent = Pct Address
 0 pwd.corp.com [10.1.1.20]
 0/ 100 = 0% |
```



```
1 0ms 0/ 100 = 0% 0/ 100 = 0% mnd.corp.com [10.1.1.10]
Trace complete.
```

An unsuccessful pathping command displays the following information:

```
C:>pathping 10.1.1.10
Tracing route to 10.1.1.10 over a maximum of 30 hops
0 pwd.corp.com [10.10.10.20]
1 * * *
Computing statistics for 0 seconds...
Source to Here This Node/Link
Hop RTT Lost/Sent = Pct Lost/Sent = Pct Address
0 pwd.corp.com [10.10.10.20]
Trace complete.
```

2. Complete the following steps on the mnd.corp.com host:

a. Run the following ping commands:

```
ping pwd.corp.com
ping pwd
ping 10.1.1.20
ping mnd.corp.com
ping mnd
ping 10.1.1.10
```

b. Run the following traceroute commands:

```
traceroute pwd.corp.com
traceroute pwd
traceroute 10.1.1.20
traceroute mnd.corp.com
traceroute mnd
traceroute 10.1.1.10
```

Ensure that each ping and traceroute command succeeds. Lost packets can indicate a slow connection between hosts. If any try to transmit a packet fails with an error message, then verify the name resolution and ensure that all routers between the source host and destination hosts are operational.

Using nsrrpcinfo to report the status of registered RPC services

Use the nsrrpcinfo command to verify that you can establish sessions to the portmapper daemon on the source and destination host. The NetWorker Remote Exec service on Windows and the nsrexecd daemon on UNIX, starts the portmapper service that NetWorker uses.

Type the following commands on the source and destination host:

```
nsrrpcinfo -p hostname_of_NetWorker_server
nsrrpcinfo -p FQDN_of_NetWorker_server
nsrrpcinfo -p IP_address_of_NetWorker_server
nsrrpcinfo -p shortname_of_destination_host
nsrrpcinfo -p FQDN_of_destination_host
nsrrpcinfo -p IP_address of the destination host
```

 **Note:** On Windows, the `NetWorker_installation_dir\nsr\bin` contains the `nsrrpcinfo` program.

When the `nsrrpcinfo` command runs successfully, the output displays a list of port numbers and names. For example:

```
# nsrrpcinfo -p
program vers proto port
100000 2 tcp 7938 nsrportmapper
100000 2 udp 7938 nsrportmapper
390436 1 tcp 7943 nsrexecd
390435 1 tcp 9549 nsrexecd
390113 1 tcp 7937 nsrexecd
```

Ensure that the correct program number appears for each NetWorker process. If you do not see the correct program number or the appropriate NetWorker ports, and a personal or external firewall exists between the source and the destination hosts, then review the NetWorker configuration port requirements.

The *EMC NetWorker Security Configuration Guide* provides more information about how to configure NetWorker in a firewall environment and the correct program numbers for each NetWorker daemon.

Verifying name resolution

When NetWorker performs name resolution lookups, NetWorker uses the first entry in the name resolution resource that matches the request. Name resolution services include: the resolver cache, DNS, LDAP/AD, and the hosts file. Name resolution lookups check the resolver cache first. Entries that appear in the cache do not reflect changes made to the host tables and on the DNS server until a cache flush occurs.

A cache flush occurs for the following hosts:

- All hosts in the cache at intervals defined by the operating system, by system-specific commands, or by reinitialization of network components, including a reboot.
- A specific host in the cache each time that you use the operating system command `nslookup` to resolve the hostname.

Determining the IP name search order

NetWorker relies on the operating system to determine the order in which to check name resolution services. Before troubleshooting a possible name resolution error, determine the search order that is used by the operating system.

The name resolution search order differs for each operating system:

- Linux, Solaris, and HP-UX operating systems use the hosts database entry in the `/etc/nsswitch.conf` file to define the name resolution search order. For example, when the operating system checks the DNS Server and then the `hosts` file, the `nsswitch.conf` entry appears as follows:

```
hosts: dns files
```

- AIX operating systems use one of three methods to select the name resolution search order:
 - The `NSORDER` environment variable.

For example, when the operating system checks the `hosts` file first and then DNS, the `NSORDER` environment variable appears as follows:

```
NSORDER=local,bind4
```

- The `hosts` database entry in the `/etc/netsvc.conf` file.
For example, when the operating system performs name resolution checks by using the DNS Server and then the `hosts` file, the `hosts` entry in the `netsvc.conf` file appears as follows:

```
hosts=local,bind4
```

- The `/etc/irs.conf` file.
For example, when the operating system checks the `hosts` file first and then the DNS (IPv4 address), the `hosts` entries in `irs.conf` file appear as follows:

```
hosts local
hosts dns4
```

Note: The `NSORDER` environment variable setting overrides the settings in the `/etc/netsvc.conf` file and the `/etc/irs.conf` file. The `/etc/netsvc.conf` file setting overrides the `/etc/irs.conf` file setting.

- Windows Server 2008 R2 SP1 operating systems use the following search order: WINS, network broadcast, `LMhosts` file, `hosts` file, then DNS. Windows Server 2008 R2 SP1 and earlier operating systems use a similar search order with the exception that the network broadcast occurs before the WINS lookup.

Verifying correct hosts file resolution

The operating system provides NetWorker with the first entry in the `hosts` file that matches the name resolution requirement. Additional instances of an IP address, FQDN, or shortname that appear in the `hosts` file for a host are ignored when NetWorker tries to resolve names.

When you create or modify the `hosts` file, ensure that you:

- Specify each hostname or IP address only once.
- Specify each FQDN and alias for a host on the same line as the IP address. For example:

```
IP address Canonical name FQDN alias alias...
```

- Specify the IPv6 loopback interface (`::1`) with the `localhost` on Linux and UNIX, when the operating system configures the IPv6 loopback interface. For example:

```
::1 localhost
127.0.0.1 localhost
```

Note: The IPv6 loopback entry must remain in the `hosts` file when the host exists in a pure IPv4, pure IPv6, or dual stack configuration.

Using the nslookup command

Use the `nslookup` command to verify that each DNS Server used by the source and destination hosts, correctly and consistently resolves both hosts by the short name, FQDN, and IP address.

About this task

Perform the following steps on the source host and destination host.

Procedure

1. Determine the Primary and Secondary DNS Servers that the host uses for name resolution:

- On UNIX, review the `/etc/resolv.conf` file.
- On Windows, type the following command from a command prompt:

```
ipconfig /all
```

2. Use the `nslookup` command in interactive mode to validate forward name resolution lookups with the Primary DNS Server:

- a. From a command prompt, type: `nslookup`
- b. At the `nslookup` command prompt, specify the following values:

```
Shortname_of_source_host
FQDN_of_source_host
IP_address_of_source_host
Shortname_of_destination_host
FQDN_of_destination_host
IP_address_of_destination_host
```

 **Note:** It is recommended that you resolve every name and IP address for each host three times to ensure that successive queries return correct and consistent values.

3. Complete the following steps when the host uses multiple DNS Servers for name resolution:

- a. Change the DNS Server that **nslookup** uses for name resolution.

In this example, the `ipconfig /all` command on a Windows host returns two DNS Servers, the Primary DNS Server 10.5.5.10, and secondary DNS Server 10.5.5.11.

To configure `nslookup` to use the IP address 10.5.5.11, type the following commands:

```
C:\>nslookup
Default Server: lad.corp.com
Address: 10.5.5.10
> server 10.5.5.11
Default Server: dmd.corp.com
Address: 10.5.5.11
```

- b. At the `nslookup` command prompt, specify the following values:

```
Shortname_of_source_host
Shortname_of_source_host
Shortname_of_source_host
FQDN_of_source_host
FQDN_of_source_host
FQDN_of_source_host
IP_address_of_source_host
IP_address_of_source_host
IP_address_of_source_host
Shortname_of_destination_host
```

```

Shortname_of_destination_host
Shortname_of_destination_host
FQDN_of_destination_host
FQDN_of_destination_host
FQDN_of_destination_host
IP_address_of_destination_host
IP_address_of_destination_host
IP_address_of_destination_host

```

Note: It is recommended that you resolve every name and IP address for each host three times to ensure that successive queries return correct and consistent values.

4. Use the `nslookup` command in interactive mode to validate reverse name resolution lookups in the reverse lookup zone with the Primary DNS Server:

- a. From a command prompt, type: `nslookup`.

- b. In the `nslookup` command prompt, type:

```
set q=ptr
```

- c. At the `nslookup` prompt, type:

```

IP_address_of_source_host
IP_address_of_destination_host

```

Clearing the resolver cache

Each operating system uses a local resolver cache. A local resolver cache removes the reliance on checking name resolution services for each name resolution request, which increases the hostname resolution speed. The operating system checks the cache first to resolve the host, and if the host record exists, the operating system does not check other name resolution services. The operating system adds an entry to the resolver cache after the first successful hostname resolution, and the entry remains in the cache for a predetermined time.

On Windows only, to display the contents of the resolver cache, type the following command:

```
ipconfig /displaydns
```

Use the appropriate command to even the contents of the resolver cache:

- On AIX and HP-UX:

- For bind 9, type:

```
rndc flush
```

- For bind 8, type:

```
refresh -s named
```

- On Solaris and Linux, restart the `nscd` daemon.

- On Windows, type:

```
ipconfig /flushdns
```

Verifying the NetWorker configuration

NetWorker contains two configurable options, the `servers` file that allows you to control access to a host and the `aliases` attribute in the Client resource, which allows you to define the names by which a host is known. When either option contains an incorrect host name, NetWorker operations can fail despite correct host name resolution and when an established connection exists between the source and destination hosts.

Ensure that the name that NetWorker uses primarily for a host appears consistently in all NetWorker resources. For example:

- Names of Client and Storage node resources. For example, if you specify the FQDN in the Name attribute when you create the Client resource for a storage node, ensure that you specify the FQDN in the Name attribute when you create the Storage Node resource.
- Names of the index database directory.
- Names specified in the Remote Access and Administrator attributes.
- Hostname references in resource attributes such as the Storage Node and Recover Storage Node attributes of a Client resource.
- Cached host certificates (NSR Peer information).

Verifying the validity of the `servers` file

The `servers` file defines a list of remote hosts that can ask the local `nsrexecd` process to start a program. For example, the NetWorker Server requests that the `nsrexecd` process on a client start the `save` process to begin a backup. The NetWorker Server installation process on certain operating systems prompts you to define remote hosts to add to the `servers` file. You can also manually modify the `servers` file at any time.

The `servers` file on a NetWorker Server host resides in the `res` subdirectory of the `nsr` directory. The location varies depending on the installation path.

When a host asks `nsrexecd` to start a process but the host does not appear in the `servers` file, a message similar to the following appears:

```
Cannot request command execution, permission denied
```

If you receive this message but the requesting host requires access, then manually edit the `servers` file on the destination host and add each short name and FQDN for the requesting host, on a separate line.

NOTICE After you make changes to the `servers` file, stop and then restart the NetWorker Server services on the host. The *NetWorker Security Configuration Guide* provides more information about how to modify the `servers` file.

Confirming the validity of Aliases attribute

Each Client resource contains an Aliases attribute that defines a list of known names that are associated with the client. The NetWorker server generates this list when you create the Client resource.

You can also manually edit the Aliases attribute value to add or remove hostname instances or IP addresses. Use the following guidelines when you modify the Aliases attribute value:

- Specify all short names and FQDNs for the host, including any retired hostnames.
- Specify each name on a separate line.

When the name returned by the operating system name lookup does not exist in any Aliases attribute for any client, a message similar to the following appears in the `daemon.raw` file:

```
hostname is not a registered client
```

Clearing the NetWorker name resolution cache

NetWorker processes maintain an internal name resolution cache of recent DNS lookups.

The amount of time that NetWorker maintains a cached entry depends on the success of the lookup:

- Successful lookup—30 minutes.
- Failed lookup—5 minutes.

When a NetWorker operation requires a name resolution lookup, NetWorker checks the internal cache first. If NetWorker finds the name in the internal cache, then NetWorker does not consult the operating system.

Use the `dbgcommand` command on the NetWorker server to send a list of cached names to the `daemon.raw` file:

```
dbgcommand -p nsrd_pid PrintDnsCache=1
```

where `nsrd_pid` is the process id of the `nsrd` process.

Use the `dbgcommand` command on the NetWorker server to even the internal name resolution cache:

```
dbgcommand -p nsrd_pid FlushDnsCache
```

where `nsrd_pid` is the process id of the `nsrd` process.

Using multihomed systems

When the NetWorker server, storage node, or client has more than one IP address, you can specify the exact TCP/IP network path that NetWorker uses during a backup.

A multihomed system is a system that has any of the following types of NICs:

- More than one NIC, each having separate IP address.
- A single NIC with multiple IP addresses.
- Multiple NICs in a single bond that has multiple IP addresses.

 **Note:** Dell EMC recommends that you set the FQDN as the hostname, instead of short names.

Multihomed system requirements

Before you configure NetWorker in a multihomed environment, review these requirements.

- Each IP address must always resolve to a unique primary hostname.

- Each IP address bound to a separate physical NIC must reside in a separate subnet.
- All the shortnames, FQDNs, and IP addresses for each NetWorker host must be correctly and consistently resolvable.
- Specify all of the hostnames that belong to a NetWorker server, storage node, or client in the Aliases attribute in the appropriate Client resource.
- Ensure that the `servers` file on each NetWorker client contains all the hostnames that resolve to the NetWorker server.

Configuring multihomed hosts in a datazone

The following table summarizes how to configure the NetWorker environment to use a multihomed NetWorker server, storage node, and client.

Table 147 Configuring multihomed hosts in NetWorker (continued)

Multihomed host	Required behavior	NetWorker configuration requirements
NetWorker server	The client sends metadata to the NetWorker server by using a specific NetWorker server NIC. The metadata includes the save set control session information and index database operations.	The <code>servers</code> file on each client must contain the shortname and FQDN for each NetWorker server NIC. The Server network interface attribute of each Client resource must contain the FQDN of the NetWorker server NIC. Each instance of the Client resource must have the same value for the Server NetWorker Interface attribute. The Alias field for the NetWorker server Client resource must contain an entry for the shortname and FQDN of each NIC.
	Each storage node device sends metadata to the NetWorker server by using a specific NetWorker server NIC. Metadata includes the device control session information and the media database operations that connect back to the <code>nsrmmdbd</code> process on the NetWorker server.	The Server network interface attribute of each Storage Node resource must contain the FQDN of the NetWorker server NIC. The Aliases attribute of the NetWorker server Client resource must contain an entry for the shortname and FQDN of each NIC.

Table 147 Configuring multihomed hosts in NetWorker (continued) (continued)


Multihomed host	Required behavior	NetWorker configuration requirements
	Each storage node library sends metadata to the NetWorker server by using a specific NIC on the NetWorker server. The metadata includes SCSI commands for the tape movements and the library inventory operations that connect back to nsrmmd process.	The Server network interface attribute of Library resource must contain the FQDN of the NetWorker server NIC. The Aliases attribute of the NetWorker server Client resource must contain an entry for the shortname and FQDN of each NIC.
Storage node	The client sends backup data to a NetWorker storage node over a specific NIC.	The Storage Nodes attribute of each Client resource must contain the FQDN of the storage node NIC. This also applies when the NetWorker server is the storage node. The Aliases attribute in the Client resource for the storage node must contain an entry for the shortname and FQDN of each NIC.
Client	The NetWorker server communicates with a client over a specific NIC.	When you create a Client instance for the client, specify a hostname for the client that is only reachable over the desired NIC.

Configuring NetWorker Server in a multihomed environment

This section provides an example of how to configure NetWorker in a multihomed environment when the NetWorker Server and the storage node have 2 NICs that communicate through different networks.

About this task


The following figure provides a graphical representation of the environment.

Figure 96 Multihomed environment


Complete the following steps to configure the multihomed environment:

Procedure


1. Update the **Aliases** attribute in the Client resource for the NetWorker Server to include the FQDN and the shortname for each NetWorker Server NIC. This figure shows the values in the Aliases attribute.

Figure 97 Configuring the Aliases attribute for NetWorker Server Client resource


2. Create a Client resource for the storage node. Update the **Aliases** attribute to include the FQDN and the shortname for each storage node NIC. This figure shows the values in the Aliases attribute.

Figure 98 Configuring the Aliases attribute for NetWorker Storage Node Client resource


3. Update the **Storage Nodes** attribute for each Client resource in VLAN1 to contain the hostname of the NIC for the storage node to which the client connects. For example, for NetWorker Client VLAN1_client, specify the storage node hostname sn1. This figure shows the values in the **Storage Nodes** attribute.

Figure 99 Storage Nodes attribute for clients in VLAN1


4. Update the **Aliases** attribute for each Client resource in VLAN1 to contain the FQDN and shortname of the client. The **Server network interface** attribute must contain the hostname of the NIC for the NetWorker Server to which the client connects. This figure shows the values in the **Aliases** and **Server network interface** attributes.

Figure 100 Aliases and Server network interface attributes for VLAN1 clients

5. Update the **Storage Nodes** attribute for each Client resource in VLAN2 to contain the hostname of the NIC interface for the storage node to which the client connects. For example, for NetWorker Client VLAN2_client, specify the storage node hostname sn2. This figure shows the values in the **Storage Nodes** attribute.

Figure 101 Storage Nodes attribute for clients in VLAN2

6. Update the **Aliases** attribute for each Client resource in VLAN2 to contain the FQDN and shortname of the client. The **Server network interface** must contain the hostname of the NIC interface for the NetWorker Server to which the client connects. This figure shows the values in the **Aliases** and **Server network interface** attributes.

Figure 102 Aliases and Server network interface attributes for VLAN2 clients

7. Create the Device resource on the remote storage node by specifying either one of the hostnames for the storage node.

NIC Teaming

NIC Teaming is a term that describes the use of multiple network interfaces in parallel. NIC teaming increases the link speed beyond the limits of any one cable or any one port and increases redundancy for higher availability.

Other terms for NIC Teaming include link aggregation, Ethernet trunk, port channel, port teaming, port trunking, link bundling, EtherChannel, Multi-Link Trunking (MLT), and NIC bonding.

NIC Teaming at the TCP level, regardless of the protocol or algorithm used, has no effect on a single TCP session. When you combine multiple links into a single link, the backup performance of a single session does not improve.

Depending on the algorithm used, starting parallel backup jobs with multiple NICs produces load balancing and can improve backup performance. To achieve load balancing, use a TCP session-based link aggregation algorithm and not a host-based algorithm. For example, use the IEEE 803.3ad/802.1ax Link Aggregation Control Protocol (LACP).

The use of trunked interfaces is transparent from a NetWorker point of view and the configuration of trunked interfaces inside NetWorker does not differ from the configuration of stand-alone interfaces. You can combine TCP trunking with multihoming, for example, by trunking some NICs on the system and leaving other NICs to work on separate subnets.

Using DHCP clients

NetWorker relies on forward and reverse hostname and IP address resolution for communication between NetWorker hosts. When DHCP allocation changes an IP address, NetWorker cannot correctly resolve the current client IP address back to a valid hostname.

To back up DHCP clients, choose one of the following solutions:

- Configure the clients and the DNS Server to allow Dynamic DNS Registration. In this configuration, each time a client receives a new IP address, the DHCP service registers the hostname and IP address with the central DNS Server.
- Configure the DHCP server to always issue the same IP address to a host. In this configuration, bind the MAC address of the host to an IP address. Register this IP address in DNS Server or add the IP address to the `servers` file on the client and the NetWorker server.

NOTICE It is recommended that you do not configure the NetWorker server as a DHCP client. If the NetWorker server is a DHCP client, then the NetWorker server must use a reserved address that the DHCP server synchronizes with the DNS server.

NetWorker TCP/IP keep-alive parameters

In NetWorker, the TCP/IP keep-alive parameters are enabled by default for connections in between “nsrjobd” and “nsrexecd”, and “nsrindexd” and save in case of client direct save. You can modify the TCP IP keep-alive parameters to prevent any potential connection failures.

The table lists the default TCP IP keep-alive parameters values in NetWorker. These values override the default operating system keep-alive parameters.

Table 148 TCP/IP parameters

TCP Parameter	Value in second
Keepalive time	300
Keepalive interval	30
Keepalive count	20

Modifying the NetWorker TCP/IP parameters in Linux platform

You can modify the NetWorker TCP/IP parameters by changing the environment variables in the “nsrrc” file for Linux based operating systems.

Procedure

1. Log in as root.
2. Create a file with a name as “nsrrc” in /nsr location
3. Enter the parameters and the values in seconds as per your network settings.

```
NW_TCP_KEEPIDLE_SECS=<value>
export NW_TCP_KEEPIDLE_SECS
NW_TCP_KEEPINTVL_SECS=<value>
export NW_TCP_KEEPINTVL_SECS
NW_TCP_KEEPCNT=<value>
export NW_TCP_KEEPCNT
```

A sample /nsr/nsrrc file.

```
NW_TCP_KEEPIDLE_SECS=200
export NW_TCP_KEEPIDLE_SECS
NW_TCP_KEEPINTVL_SECS=20
export NW_TCP_KEEPINTVL_SECS
NW_TCP_KEEPCNT=10
export NW_TCP_KEEPCNT
```

 **Note:** You must restart the NetWorker services for the changes to take effect.

Modifying the NetWorker TCP/IP parameters in Windows platform

You can modify the NetWorker TCP/IP parameters by changing the environment variables in the system properties.

Procedure

1. Open **System Properties** and click on the **Advanced** tab.
2. Click on **Environment variable**.
3. In the section **System Variables**, click **New** and a pop up is displayed.
4. Type the variable *name* and *value*.

 **Note:** You must restart the system for the changes to take effect.

CHAPTER 19

Cloud Supportability

This chapter contains the following topics:

- [CloudBoost appliance as the back up target](#).....882
- [Support for Azure Stack](#)..... 882
- [Cloud service provider support matrix for NetWorker](#)..... 883

CloudBoost appliance as the back up target

The CloudBoost appliance provides an integrated solution for existing supported backup environment by enabling the transfer of backups to public, hybrid, or private cloud storage. The CloudBoost appliance supports the following use cases: long-term retention to the cloud and backup to a private or public cloud.

CloudBoost decouples metadata from data. Encryption keys, metadata, and file system information are housed separately from the data, which removes a common bottleneck for cloud read/write operations. All advanced data services, such as chunking, encryption, inline deduplication, compression, and bulk data transfers are performed separately from metadata storage.

CloudBoost appliance can be configured on VMware ESXi and public clouds. For more information on installation and configuration, see the *NetWorker 19.1 with CloudBoost 19.1 Integration Guide*.

Support for Azure Stack

You can install NetWorker client agent on Azure stack and protect the virtual machines with guest level protection.

Figure 103 Azure stack backup and disaster recovery

Azure Stack Backup and Disaster Recovery


- Recoverability of Admin and Tenant Capabilities

- Tenant VM protection

- Primary Backup and Restore on-premises
- Cold Storage either on-premises or with a cloud service provider

- Admin Layer Protection

- Azure Stack service data and tenant/app metadata backed up via "tarball" to SMB3 share.



Dell - Internal Use - Confidential

DELL EMC

Cloud service provider support matrix for NetWorker

You can install NetWorker on both public and Hybrid cloud environment.

Table 149 Cloud service provider support matrix

Cloud Platform	Cloud type	AWS	AWS US Gov	Azure	Azure US Gov	Azure Stack	Google Cloud Platform
NetWorker Virtual Edition	Public	Supported in NetWorker 18.1 and later	Supported in NetWorker 18.1 and later	Supported in NetWorker 18.1 and later	Supported in NetWorker 18.1 and later	Not available	Not available
NetWorker Software Installation	Public	Supported in NetWorker 9.2.1 and later	Supported in NetWorker 9.2.1 and later	Supported in NetWorker 9.2.1 and later	Supported in NetWorker 9.2.1 and later	Not available	Supported in NetWorker 18.2 and later
NetWorker Software Installation	Hybrid	Not available	Not available	Not available	Not available	Guest based backup for Virtual Machines running inside stack to NetWorker server and DD running outside stack.	Not available

For more information on deploying NetWorker Virtual Edition, see *NetWorker Virtual Edition Installation Guide*.

CHAPTER 20

Troubleshooting

This chapter contains the following topics:

- [Before you contact technical support](#)..... 886
- [NetWorker log files](#)..... 887
- [NetWorker Authentication Service logs](#)..... 910
- [NetWorker functionality issues](#)..... 913
- [NetWorker locale and code set support](#)..... 926
- [Enabling service mode for NetWorker](#)..... 927
- [Resetting the NMC password for NetWorker 19.1 and later](#)..... 927
- [No privileges to view NetWorker server from NMC](#).....927
- [Network and server communication errors](#).....928

Before you contact technical support

If the solutions in this chapter do not solve the problem, go to the Online Support website at <https://support.emc.com> for technical assistance.

 **Note:** For more details on Support website, see the Preface section of the *NetWorker Administration Guide*.

Provide the following information.

- Software version of the NetWorker component.
- Operating system version.
For example:
 - For Solaris, at the command prompt type the `uname -a` command.
 - For AIX, at the command prompt type the `oslevel` command.
- Hardware configuration.
- Information about devices and other SCSI IDs.
To determine this information, use the following commands:
 - For AIX, Linux, and Solaris, type the `/usr/sbin/inquire` command.
 - For HP-UX, type the `/etc/ioscan` command.
- If you are using an autochanger, then the type of connection (SCSI or RS-232).
Also, provide the version of the autochanger driver you are using:
 - For Solaris, type the `pkginfo -x` command:
`# pkginfo LGTOdrvr`
 - For AIX, type the `lsllpp -l | grep EMC` command.
- Supply the following information:
 - How to reproduce the problem.
 - Exact error messages that you have encountered.
 - Number of times that you have seen the problem.
 - If the NetWorker operation was successful before you made any changes and, if so, the changes that you made.

Determining the version of NetWorker software running on a client

To determine the version of the NetWorker software running on a client, use either the client properties window in NMC, the NetWorker User program on Windows or the `nsradmin` command.

Determining the software version by using NMC

Use NMC to connect to the NetWorker server with a user that is a member of the Application Administrators user group on the NetWorker server.

1. On the **Protection** window, select **Clients** from the left navigation pane.
2. Right-click the client and select **Modify client properties**.
3. On the **Info & Licensing** tab, review the **NetWorker version** attribute.

NOTICE When you do not use the **Client Configuration** wizard to create the client, NMC updates the NetWorker version attribute after the first backup. When you update the NetWorker software on a client, the NetWorker version attribute does not reflect the new version until the first backup after the update.

Determining the software version by using NetWorker User

On Windows hosts, use the NetWorker User application to determine the NetWorker software version.

1. From the **Help** menu, select **About NetWorker User**. The NetWorker version number appears in the **About** dialog box.
2. Click **OK** to close the dialog box.

Determining the client software version by using nsradmin

Use the `nsradmin` program on the NetWorker server to determine the version of the NetWorker software that is installed on a host, from a command prompt.

1. At the command prompt, type:

```
nsradmin -p nsrexecd
```

2. At the **nsradmin** command prompt, type:

```
nsradmin> show NetWorker version
nsradmin> print type: NSRLA
```

The `nsradmin` output displays the version of NetWorker software running on each client.

Displaying diagnostic mode attributes

NetWorker resources such as clients and devices contain diagnostic attributes that are hidden by default from the Console server view.

1. Open the **Administration** window.
2. From the **View** menu, select **Diagnostic Mode**.
3. Right-click any resource and select **Properties** to see diagnostic attributes.

NetWorker log files

This section provides an overview of the log files that are available on NetWorker hosts and the NMC server.

NetWorker Server log files

This section provides a summary of the log files available on a NetWorker Server and log file management.

Table 150 NetWorker Server log files

Component	File name and default location	Description
NetWorker Server daemons	UNIX: <code>/nsr/logs/daemon.raw</code>	Main NetWorker log file.

Table 150 NetWorker Server log files (continued)

Component	File name and default location	Description
	Windows: C:\Program Files\EMC NetWorker\nsr\logs\daemon.raw	Use the <code>nsr_render_log</code> program to view the contents of the log file.
Client fix	UNIX: /nsr/logs/client_fix /nsr/logs/client_fix.raw Windows: - C:\Program Files\EMC NetWorker\nsr\logs\client_fix - C:\Program Files\EMC NetWorker\nsr\logs\client_fix.raw	Contains status information that is related to the use of the <code>nsr_client_fix</code> command.
NetWorker Server generated syslog messages and daemon.notice	UNIX: OS log file that is defined by system log configuration file. Windows: C:\Program Files\EMC NetWorker\nsr\logs\messages	Contains general NetWorker error messages.
NetWorker Server generated syslog messages local0.notice and local0.alert	Log file name and location that is defined by the system log configuration file.	UNIX only, OS log file. Note: NetWorker does not modify the <code>syslog.conf</code> file to configure <code>local0.notice</code> and <code>local0.alert</code> . Vendor specific documentation describes how to configure <code>local0.notice</code> and <code>local0.alert</code>
Disaster recovery command line wizard, nsrdr program	UNIX: /nsr/logs/nsrdr.log Windows: C:\Program Files\EMC NetWorker\nsr\logs\nsrdr.log	Contains detailed information about the internal operations that are performed by the <code>nsrdr</code> program. NetWorker overwrites this file each time you run the <code>nsrdr</code> program.
Index log	UNIX: /nsr/logs/index.log Windows: C:\Program Files\EMC NetWorker\nsr\logs\index.log	Contains warnings about the size of the client file index and low disk space on the file system that contains the index files. By default, the Index size notification on the NetWorker Server sends information to the log file.
Hypervisor	UNIX: /nsr/logs/Hypervisor/hyperv-flr-ui/hyperv-flr-ui.log	Contains status information about the Hyper-V FLR interface.

Table 150 NetWorker Server log files (continued)

Component	File name and default location	Description
	Windows: C:\Program Files\EMC NetWorker\nsr \logs\hyperv-flr-ui\hyperv-flr- ui.log	
VMware protection policies	UNIX: /nsr/logs/Policy/ <i>VMware_protection_policy_name</i> Windows: C:\Program Files\EMC NetWorker\nsr \logs\Policy \ <i>VMware_protection_policy_name</i>	Contains status information about VMware Protection Policy actions. NetWorker creates a separate log file for each action.
Policies	UNIX: /nsr/logs/policy.log Windows: C:\Program Files\EMC NetWorker\nsr \logs\policy.log	Contains completion information about VMware Protection Policies. By default, the VMware Protection Policy Failure notification on the NetWorker Server sends information to the log file.
Snapshot management	UNIX: /nsr/logs/nwsnap.raw Windows: C:\Program Files\EMC NetWorker\nsr \logs\nwsnap.raw /nsr/logs/ nwsnap.raw	Contains messages that are related to snapshot management operations. For example, snapshot creation, mounting, deletion, and rollover operations. Use the <i>nsr_render_log</i> program to view the contents of the log file.
Migration	UNIX: /nsr/logs/migration Windows: C:\Program Files\EMC NetWorker\nsr \logs\migration	Contains log files that provide detailed information about the migration of attributes in an 8.2.x and earlier resources during an update of the NetWorker Server. The <i>NetWorker Installation Guide</i> provides more information about all the migration log files.
Media management	UNIX: /nsr/logs/media.log Windows: C:\Program Files\EMC NetWorker\nsr \logs\media.log	Contains device related messages. By default, the device notifications on the NetWorker Server send device related messages to the <i>media.log</i> file on the NetWorker Server and each Storage Node.
Recovery Wizard	UNIX: /nsr/logs/recover/ <i>recover_config_name_YYYYMMDDHHMMSS</i>	Contains information that can assist you in troubleshooting recovery failures. NetWorker creates a log file on the NetWorker Server for each recover job.

Table 150 NetWorker Server log files (continued)

Component	File name and default location	Description
	Windows: C:\Program Files\EMC NetWorker\nsr \logs\recover \recover_config_name_YYYYMMDDHHMMSS	
Package Manager log	UNIX: /nsr/logs/nsrccd.raw Windows: C:\Program Files\EMC NetWorker\logs \nsrccd.raw	Contains information that is related to the Package Manager and the <code>nsrpush</code> command. Use the <code>nsr_render_log</code> program to view the contents of the log file.
Rap log	UNIX: /nsr/logs/rap.log Windows: C:\Program Files\EMC NetWorker\logs \rap.log	Records configuration changes that are made to the NetWorker Server resource database.
Security Audit log	UNIX: /nsr/logs/ NetWorker_server_sec_audit.raw Windows: C:\Program Files\EMC NetWorker\logs \NetWorker_server_sec_audit.raw	Contains security audit related messages.

NMC server log files

This section provides a summary of the log files available on an NMC server.

Table 151 NMC server log files

Component	File name and default location	Description
NMC server log files	Linux: /opt/lgtomc/management/ logs/gstd.raw Windows: C:\Program Files\EMC NetWorker\Management\logs \gstd.raw	Contains information that is related to NMC server operations and management. Use the <code>nsr_render_log</code> program to view the contents of the log file.
NMC server database conversion	Linux:	Contains the results of the NMC server database conversion that is

Table 151 NMC server log files (continued)

Component	File name and default location	Description
	/opt/lgtonmc/logs/ gstdbupgrade.log Windows: C:\Program Files\EMC NetWorker\Management\logs \gstdbupgrade.log	performed during an upgrade operation.
NMC web server	Linux: /opt/lgtonmc/management/ logs/web_output Windows: C:\Program Files\EMC NetWorker\Management\logs \web_output	Contains messages for the embedded Apache httpd web server on the NMC server.
NMC server database log files	Linux: /opt/lgtonmc/management/ nmcdb/pgdata/db_output Windows: C:\Program Files\EMC NetWorker\Management\nmcdb \pgdata\db_output	Contains messages for the embedded PostgreSQL database server on the NMC server.

NetWorker Client log files

This section provides a summary of the log files available on a NetWorker Client.

Table 152 Client log files

Component	File name and default location	Description
NetWorker Client daemons	UNIX: /nsr/logs/daemon.raw Windows: C:\Program Files\EMC NetWorker\nsr\logs \daemon.raw /nsr/logs/ daemon.raw	Main NetWorker log file. Use the <code>nsr_render_log</code> program to view the contents of the log file.
User log	C:\Program Files\EMC NetWorker\logs \networkr.raw	For Windows only, contains a record of every file that was part of an attempted manual backup or recovery operation

Table 152 Client log files (continued)

Component	File name and default location	Description
		that is started by the NetWorker User program. Subsequent manual backup or recover operations overwrite the file. Use the <code>nsr_render_log</code> program to view the contents of the log file.
Windows Bare Metal Recovery (BMR)	The following files in the <code>X:\Program Files\EMC NetWorker\nsr\logs\</code> directory: <code>ossr_director.raw</code>	Contains the recovery workflow of the <code>DISASTER_RECOVERY:\</code> and any errors that are related to recovering the save set files or Windows ASR writer errors. Use the <code>nsr_render_log</code> program to view the contents of the log file.
	<code>recover.log</code>	Contains the output that is generated by the NetWorker <code>recover.exe</code> program and error messages that are related to critical volume data recovery.
	<code>winPE_wizard.log</code>	Contains workflow information that is related to the NetWorker BMR wizard user interface.
	<code>winpe_nw_support.raw</code>	Contains output from the <code>winpe_nw_support.dll</code> library. The output provides information about communications between the NetWorker BMR wizard and the NetWorker Server. Use the <code>nsr_render_log</code> program to view the contents of the log file.
	<code>winpe_os_support.log</code>	Contains output information that is related to Microsoft native API calls.
CloudBoost - NetWorker Client	The following log files in the <code>/nsr/logs/cloudboost</code> directory:	These files appear on a client direct-enabled NetWorker Client and contain information about data stored on a CloudBoost device. The

Table 152 Client log files (continued)

Component	File name and default location	Description
	<code>MagFS.log.ERROR.date-timestamp.pid.txt</code> <code>MagFS.log.FATAL.date-timestamp.pid.txt</code> <code>MagFS.log.INFO.date-timestamp.pid.txt</code>	severity of the message determines which log file that error message is written to. The maximum size of the log files are 100 MB. Before a client direct backup, the <code>save</code> process checks the size of the file. When the maximum size is reached, <code>save</code> starts an automatic trimming mechanism, which renames and compresses the log file. The maximum number of versions for a file is 10. When the number of renamed log files reaches the maximum version value, NetWorker removes the oldest log when a new version of the log file is created. Note: The Troubleshooting manual backups section of the <i>NetWorker Administration Guide</i> describes how to use the <code>CB_LOG_DIR_LOCATION</code> environment variable to change the default log file location.
CloudBoost - CloudBoost Appliance	The following log files in the <code>/nsr/logs/cloudboost</code> directory: <code>MagFS.log.ERROR.date-timestamp.pid.txt</code> <code>MagFS.log.FATAL.date-timestamp.pid.txt</code> <code>MagFS.log.INFO.date-timestamp.pid.txt</code>	These files appear on the CloudBoost appliance and contain information about operations performed on a CloudBoost device. The severity of the message determines which log file that error message is written to. The maximum size of the log files are 100 MB. When the maximum size is reached, the <code>nsrmmd</code> process starts an automatic trimming mechanism, which renames and compresses the log file. The

Table 152 Client log files (continued)

Component	File name and default location	Description
		maximum number of versions for a file is 10. When the number of renamed log files reaches the maximum version value, NetWorker removes the oldest log when a new version of the log file is created.

View log files

NetWorker sends messages to two types of logs. Plain text log files that are saved with the `.log` extension and unrendered log files that are saved with the `.raw` extension.

The `.log` files and the messages that appear in NMC use the locale setting of the service that generates the log message. To view the contents of `.log` files, use any text editor. Before you can view `.raw` files in a text editor, render the `.raw` file into the locale of the local computer. You can use the `nsr_render_log` command manually render the raw log files or you can configure NetWorker to render the log files at runtime.

The `nsr_render_log` command renders internationalized NetWorker log files in to the current locale of the host that the user uses to run the program. All other log files, as well as messages displayed in NMC, use the locale of the service that is generating the log message. The `nsr_render_log` program is non-interactive. Use command line options to specify the log file that you want to view and the format of the output. The `nsr_render_log` program sends the results to `stdout`. You can redirect and save the output to a file.

Rendering a raw file manually

The `nsr_render_log` program is non-interactive. When you use the `nsr_render_log` program to render the contents of the `.raw` file to the locale of the host where you run the command, `nsr_render_log` prints the output to `stdout`. You can redirect this output to a file and view the output in a text editor.

Before you begin

The `bin` subdirectory in the NetWorker installation directory contains the `nsr_render_log` program. If the `bin` directory is not in the search path of the host where you run the command, include the full path when you use the `nsr_render_log` program. If you do not run the `nsr_render_log` command from the directory that contains the `.raw` file, include the path to the `.raw` file.

About this task

The `nsr_render_log` program supports a number of options that allow you to filter the contents of a `.raw` file and render the contents into an easy to read format.

Procedure

- To render a raw file into a format similar to a `.log` file and redirect the output to a text file, type: `nsr_render_log -c -empathy raw_filename 1>output_filename 2>&1`

where:

- *raw_filename* is the name of the unrendered file. For example, `daemon.raw`
- *output_filename* is the name of the file to direct the output to
- `-c` suppresses the category
- `-m` suppresses the message ID
- `-e` suppresses the error number
- `-a` suppresses the activity ID
- `-p` suppresses the process ID
- `-t` suppresses the thread ID
- `-h` suppresses the hostname
- `-y` suppresses the message severity

- To render a `.raw` file from a remote machine, type: `nsr_render_log -c -empathy -R hostname raw_filename 1>output_filename 2>&1`

where:

- *hostname* is the name of the host that contains the `.raw` file.
- *raw_filename* is the name of the unrendered file. For example, `daemon.raw`
- *output_filename* is the name of the file to direct the output to
- `-c` suppresses the category
- `-e` suppresses the error number
- `-m` suppresses the message ID
- `-p` suppresses the process ID
- `-a` suppresses the activity ID
- `-t` suppresses the thread ID
- `-h` suppresses the hostname
- `-y` suppresses the message severity

- To render a `.raw` file and only view log file messages for a specific device, type: `nsr_render_log -c -empathy -F devicename raw_filename 1>output_filename 2>&1`

where *devicename* is the name of the device.

- To render only the most recently logged messages, type: `nsr_render_log -c -empathy -B number raw_filename 1>output_filename 2>&1`

where *number* is the number of lines that you want to render.

The *NetWorker Command Reference Guide* provides detailed information about the `nsr_render_log` program and the available options.

- To render a `.raw` file and only view certain messages severities, type:
`nsr_render_log -c -empath -Y message_severity 1>output_filename
2>&1`

where *message_severity* is one of the severity types listed in the following table.

Table 153 Message types

Type	Description
Informational	Information that may be useful, but does <i>not</i> require any specific action.
Warning	A temporary problem that NetWorker software may resolve or prompt you to resolve.
Notification	An event has occurred that generated a message.
Error	Errors that you are required to resolve.
Critical	Errors that you are required to resolve, to ensure successful NetWorker operations.
Severe	Errors that cause NetWorker services to become disabled or dysfunctional.

The UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsr_render_log` program and the available options.

Rendering raw log files at runtime

You can instruct the NetWorker software to render the `daemon.raw` and `gstd.raw` files into the locale of the host at runtime, in addition to creating locale-independent log files. This allows you to view the log file in a text editor without using the `nsr_render_log` program to render the file first.

Before you begin

Log in to the NetWorker host with the root (UNIX) or Administrator (Windows) user account.

About this task

To instruct the NetWorker software to render logs in the locale of the computer hosting the file, set the **runtime rendered log file** attribute in the NSRLA database. For backward compatibility with previous releases of the NetWorker software, runtime rendered log files contain the following attributes:

- Message ID
- Date and time of message
- Rendered message

Procedure

1. To access the NSRLA database, from a command prompt, use the `nsradmin` program:

```
nsradmin -p nsrexec
```

2. Set the resource type to NSR log:

```
. type: NSR log
```


3. Display a list of all log file resources:

```
print
```

For example, on a Windows NMC server, output similar to the following appears:

```
nsradmin> print
type: NSR log;
administrator: Administrators,
"group=Administrators,host=bu-iddnwserver.iddlab.local";
owner: NMC Log File;
maximum size MB: 2;
maximum versions: 10;
runtime rendered log: ;
runtime rollover by size: Disabled;
runtime rollover by time: ;
name: gstd.raw;
log path: \
"C:\\Program Files\\EMC NetWorker\\Management\\GST\\logs\\
\\gstd.raw";

type: NSR log;
administrator: Administrators,
"group=Administrators,host=bu-iddnwserver.iddlab.local";
owner: NetWorker;
maximum size MB: 2;
maximum versions: 10;
runtime rendered log: ;
runtime rollover by size: Disabled;
runtime rollover by time: ;
name: daemon.raw;
log path: \
"C:\\Program Files\\EMC NetWorker\\nsr\\logs\\daemon.raw";
```

4. Define the log resource that you want to edit:

```
. type: NSR log; name: log_file_name
```

For example, to select the `daemon.raw` file, type the following:

```
. type: NSR log; name: daemon.raw
```

5. To define the path and file name for the rendered log file, use the **Runtime rendered log** attribute.

For example, to save rendered messages to the file `rendered.log` in the default NetWorker logs directory on a Windows host, type:

```
update runtime rendered log: "C:\\Program Files\\EMC NetWorker\\
\\nsr\\logs\\rendered.log"
```

6. When prompted to confirm the update, type: `y`

7. Verify that the attribute value update succeeds:

```
nsradmin> print

type: NSR log;
```

```

administrator: root, "user=administrator,host=bu-
iddnwserver.iddlab.local";
owner: NetWorker;
maximum size MB: 2;
maximum versions: 10;
runtime rendered log:C:\\Program Files\\EMC NetWorker\\nsr
\\logs\\daemon.log ;
runtime rollover by size: Disabled;
runtime rollover by time;;
name: daemon.raw;
log path: C:\\Program Files\\EMC NetWorker\\Management\\
\\GST\\logs\\daemon.raw;

```

8. Exit the `nsradmin` program.

Raw log file management

The NetWorker software manages the size and the rollover of the raw log files.

NetWorker automatically manages the `nwsnap.raw` and `nsrcpd.raw` files in the following ways:

- **`nwsnap.raw`:** Before a process writes messages to the `nwsnap.raw` file, the process checks the size of the `.raw` file. The process invokes the trimming mechanism when the size of the log file is 100 MB or larger. Snapshot management supports up to 10 `.raw` file versions.
- **`nsrcpd.raw`:** When the NetWorker daemons start on the machine, the startup process checks the size of the raw file. The startup process runs the trimming mechanism when the size of the log file is 2 MB or larger. Package Manager supports 10 raw file versions.

NetWorker enables you to customize the maximum file size, maximum number of file versions, and the runtime rollover of the `daemon.raw`, `gstd.raw`, `networkr.raw`, and `Networker_server_sec_audit.raw` files. Use the `nsradmin` program to access the NSRLA database, and modify the attributes that define how large the log file becomes before NetWorker trims or renames the log file.

The following table describes the resource attributes that manage the log file sizes.

Table 154 Raw log file attributes that manage log file size

Attribute	Information
Maximum size MB	Defines the maximum size of the log files. Default: 2 MB
Maximum versions	Defines the maximum number of the saved log files. When the number of copied log files reaches the maximum version value, NetWorker removes the oldest log when a new copy of the log file is created. Default: 10
Runtime rollover by size	When set, this attribute invokes an automatic hourly check of the log file size.

Table 154 Raw log file attributes that manage log file size (continued)

Attribute	Information
	When you configure the runtime rendered log attribute, NetWorker trims the runtime rendered log file and the associated <code>.raw</code> file simultaneously. Default: disabled
Runtime rollover by time	When set, this attribute runs an automatic trimming of the log file at the defined time, regardless of the size. The format of the variable is HH:MM (hour:minute). When you configure the runtime rendered log attribute, NetWorker trims the runtime rendered log file and the associated <code>.raw</code> file simultaneously. Default: undefined  Note: After setting this attribute, restart NetWorker services for the change to take effect.

How the trimming mechanism trims the log files differs depending on how you define the log file size management attributes. The following table summarizes the trimming behavior.

Table 155 Raw log file attributes that manage the log file trimming mechanism

Attribute configuration	Trimming behavior
When you configure runtime rollover by time or runtime rollover by size	- NetWorker copies the contents of the existing log file to a new file with the naming convention: <code>daemondate_time.raw</code> - NetWorker truncates the existing <code>daemon.raw</code> to 0 MB.  Note: When this mechanism starts on a NetWorker Server that is under a heavy load, this process may take some time to complete.
When you do not configure runtime rollover by time or runtime rollover by size	- NetWorker checks the log file size when the <code>nsrexecd</code> process starts on the computer. - When the log file size exceeds the size that is defined by the maximum size MB attribute, NetWorker renames the existing log file to <code>log_file_name_date_time.raw</code> then creates a new empty log file.

Table 155 Raw log file attributes that manage the log file trimming mechanism (continued)

Attribute configuration	Trimming behavior
	Note: When the <code>nsrd</code> daemon or NetWorker Backup and Recover Server service runs for a long time, the size of the log file can become much larger than the value defined by maximum size MB.

Managing raw log file size for the `daemon.raw`, `networkr.raw`, and `gstd.raw` files

To configure the NetWorker software to rollover the `.raw` file by time, perform the following steps.

Procedure

1. Log in to the NetWorker host with `root` on UNIX or `Administrator` on Windows.
2. To access the NSRLA database, use the `nsradmin` program:

```
nsradmin -p nsrexec
```

3. Set the resource type to NSR log:

```
. type: NSR log
```

4. Display a list of all log file resources:

```
print
```

For example, on a Windows NMC server, output similar to the following appears:

```
nsradmin> print
type: NSR log;
administrator: Administrators,
"group=Administrators,host=bu-iddnwserver.iddlab.local";
owner: NMC Log File;
maximum size MB: 2;
maximum versions: 10;
runtime rendered log: ;
runtime rollover by size: Disabled;
runtime rollover by time: ;
name: gstd.raw;
log path: \
"C:\\Program Files\\EMC NetWorker\\Management\\GST\\logs\\
\\gstd.raw";

type: NSR log;
administrator: Administrators,
"group=Administrators,host=bu-iddnwserver.iddlab.local";
owner: NetWorker;
maximum size MB: 2;
maximum versions: 10;
runtime rendered log: ;
runtime rollover by size: Disabled;
runtime rollover by time: ;
```

```
name: daemon.raw;
log path: \
"C:\\Program Files\\EMC NetWorker\\nsr\\logs\\daemon.raw";
```

5. Define the log resource that you want to edit:

```
. type: NSR log; name: log_file_name
```

For example, to select the `gstd.raw` file, type the following:

```
. type: NSR log; name: gstd.raw
```

6. Update the **runtime rollover by time** attribute with the time that you want to rollover the log file.

For example, to configure the `gstd.raw` file to rollover at 12:34 AM, type:

```
update runtime rollover by time: "00:34"
```

7. When prompted to confirm the update, type: `y`

8. Verify that the attribute value update succeeds:

```
nsradmin> print
```

```
type: NSR log;
administrator: root, "user=administrator,host=bu-
iddnwserver.iddlab.local";
owner: NMC Log File;
maximum size MB: 2;
maximum versions: 10;
runtime rendered log: ;
runtime rollover by size: Disabled;
runtime rollover by time: "00:34";
name: gstd.raw;
log path: C:\\Program Files\\EMC NetWorker\\Management\\
\\GST\\logs\\gstd.raw;
```

9. Exit the `nsradmin` program.

Configuring logging levels

This section describes how to modify the logging levels of the NetWorker and NMC processes to troubleshoot issues.

Setting the troubleshoot level for NetWorker daemons

How you configure the NetWorker daemons to run in troubleshoot mode depends on the daemon.

On a NetWorker server, you can configure the `nsrctld` and `nsrexecd` to start in troubleshoot mode. The `nsrctld` daemon starts other daemons, as required. To capture troubleshoot output for the daemons that the `nsrctld` daemon starts use the `dbgcommand`.

On an NMC server, you can start the `gstd` daemon in troubleshoot mode.

Starting `nsrctld` and `nsrexecd` daemons in troubleshoot mode on UNIX

The `nsrctld` daemon is the main process for the NetWorker server. To troubleshoot problems with the NetWorker server process, start the `nsrctld` process in

troubleshoot mode. The `nsrexecd` process is the main process for NetWorker client functions. To troubleshoot problems that are related to NetWorker client functions, start the `nsrexecd` process in troubleshoot mode.

Procedure

1. Log in to the NetWorker host with the root account, and then stop the NetWorker processes:

```
nsr_shutdown
```

2. From a command prompt, start the daemon, and then specify the troubleshoot level.

For example:

- To start the `nsrexecd` daemon in troubleshoot mode, type:

```
nsrexecd -D9 1>filename2>&1
```

- To start the `nsrctld` daemon in troubleshoot mode, type the following command:

```
source /opt/nsr/admin/networkerrc; source /opt/nsr/admin/  
nsr_serverrc; nsrctld -D 9 1>filename.log 2>&1
```

Where *filename* is the name of the text file that NetWorker uses to store the troubleshoot messages.

3. After you collect the necessary troubleshoot information, perform the following steps:
 - a. Stop the NetWorker processes by using the `nsr_shutdown` command.
 - b. Restart the processes by using the NetWorker startup script:

- On Solaris and Linux, type:

```
/etc/init.d/networker start
```

- On HP-UX, type:

```
/sbin/init.d/networker start
```

- On AIX, type:

```
/etc/rc.nsr
```

Starting the NetWorker daemons in troubleshoot mode on Windows

The NetWorker Backup and Recovery service starts the `nsrctld` process, which is the main process for a NetWorker server. To troubleshoot problems with the NetWorker server process, start the `nsrctld` process in troubleshoot mode. The NetWorker Remote Exec service starts the `nsrexecd` process which is the main process for NetWorker client functions. To troubleshoot problems that are related to NetWorker client functions, start the `nsrexecd` process in troubleshoot mode.

Procedure

1. Open the Services applet, `services.msc`.
2. Stop the NetWorker Remote Exec service.

On a NetWorker server, this also stops the **NetWorker Backup and Recover** service.

3. To put a `nsrexecd` process in troubleshoot mode:
 - a. Right-click the **NetWorker Remote Exec** service, and then select **Properties**.
 - b. In the **Startup Parameters** field, type `-D x`,
where `x` is a number between 1 and 99.
 - c. Click **Start**.
4. To put the `nsrdd` process in troubleshoot mode:
 - a. Right-click the NetWorker Backup and Recover service, and then select **Properties**.
 - b. In the **Startup Parameters** field, type `-D x`,
where `x` is a number between 1 and 99.
 - c. Click **Start**.

Results

NetWorker stores the troubleshoot information in the `daemon.raw` file.

After you finish

After you capture the troubleshoot information, stop the NetWorker services, remove the `-D` parameter, and then restart the services.

Starting the NMC server daemon in troubleshoot mode

When you can access the NMC GUI, use the Debug Level attribute in the System Options window to start the `gstd` daemon in troubleshoot mode.

When you cannot access the NMC GUI, use environment variables to start the `gstd` daemon in troubleshoot mode.

Starting the NMC server daemon in troubleshoot mode using NMC

The `gstd` daemon is the main NMC server process. To troubleshoot NMC GUI issues, start the `gstd` daemon in troubleshoot mode.

Before you begin

Log in to the NMC server with an administrator account.

Procedure

1. In the NMC Console, select **Setup**.
2. On the **Setup** menu, select **System Options**.
3. In the **Debug Level** field, select a number between 1 and 20.

Results

NMC stores the troubleshoot information in the `gstd.raw` file.

After you finish

After you capture the troubleshoot information, stop the NetWorker services, set the **Debug Level** to 0, and then restart the services.

Starting the NMC server daemon in troubleshoot mode using environment variables

Use environment variable to put the `gstd` daemon in troubleshoot mode when you cannot access the NMC GUI.

Setting the GST debug environment variable on Windows

To set the GST troubleshoot environment variable on Windows, use the Control Panel system applet on the NMC server.

Procedure

1. Browse to **Control Panel > System and Security > System > Advanced Settings**.
2. On the **General** tab, click **Environment Variables**.
3. In the **System variables** section, click **New**.
4. In the **Variable name** field, type: `GST_DEBUG`
5. In the **Variable value** field, type a number between 1 and 20.
6. Stop, and then start the **EMC gstd** service.

Results

NMC stores the troubleshoot information in the `gstd.raw` file.

After you finish

After you capture the troubleshoot information, stop the **EMC gstd** service, remove the environment variable from the startup file, and then restart the **EMC gstd** service.

Setting the GST troubleshoot environment variable on UNIX

Use a borne shell script to put the `gstd` daemon in troubleshoot mode.

Procedure

1. Modify the file permissions for the `gst` startup file. By default, the file is a read-only file.

The file location varies depending on the operating system:

- Solaris and Linux: `/etc/init.d/gst`
- AIX: `/etc/rc.gst`

2. Edit the file and specify the following at beginning of the file:

```
GST_DEBUG=x
```

```
export GST_DEBUG
```

where `x` is a number between 1 and 20.

3. Stop, and then restart the `gstd` daemon:

- Solaris and Linux: Type:

```
/etc/init.d/gst stop
```

then

```
/etc/init.d/gst start
```

- AIX: Type:

```
/etc/rc.gst start
```

then


```
/etc/rc.gst stop
```

Results

NMC stores the troubleshoot information in the `gstd.raw` file.

After you finish

After you capture the troubleshoot information, stop the `gstd` daemon, remove the environment variable from the startup file, and then restart the `gstd` daemon.

Using the `dbgcommand` program to put NetWorker process in troubleshoot mode

Use the `dbgcommand` program to generate troubleshoot messages for NetWorker daemons and processes without the stopping and starting the NetWorker daemons. You can also use the `dbgcommand` program to produce troubleshoot information for a process that another process starts. For example, use the `dbgcommand` to put the `nsrmmmd` process in troubleshoot mode.

Procedure

- From a command prompt on the NetWorker host, determine the process id (PID) of the daemon or process that you want to troubleshoot.
 - On Windows: To determine the PID, use the **Task Manager**.
 Note: If you do not see the PID for each process on the **Process** tab, browse to **View > Select Columns**, and then select **PID (Process Identifier)**.
 - On UNIX, use the `ps` command. For example, to get a list of all the NetWorker processes that start with `nsr`, type `ps -ef | grep nsr`.

- From a command prompt, type:

```
dbgcommand -p PID -Debug=x
```

where:

- PID* is the process id of the process.
- x* is a number between 0 and 9.

Note: 0 turns off troubleshoot.

Results

NetWorker logs the process troubleshoot information in the `daemon.raw` file.

After you finish

To turn off troubleshoot, type:

```
dbgcommand -p PID -Debug=0
```

Running individual clients in a group in troubleshoot mode

Modify the backup command attribute for a Client resource to send verbose backup information to the `daemon.raw` file, for individual clients in a group.

Before you begin

Use NMC to connect to the NetWorker server with a user that is a member of the Application Administrators or Database Administrators user group.

Procedure

1. From the **Administration** window, click **Protection**.
2. In the left navigation pane, click **Clients**.
3. Right-click the client, and then select **Modify Client Properties**.
4. On the **Apps & Modules** tab, in the **Backup command** attribute, type:

```
save -Dx
```

where *x* is a number between 1 and 99.

5. Click **OK**.

Results

At the scheduled time, NetWorker logs troubleshoot information for the client backup in the `daemon.raw`.

After you finish

When the group backup operations complete, edit the properties of the client and clear the **Backup Command** field.

Running client-initiated backups in troubleshoot mode from the command line

Use the `save` program to perform a client-initiated backup from the command line.

About this task

On the host you want to backup, type the following command:

```
save -Dx file_sytem_objects
1>filename 2>&1
```

where:

- *x* is a number between 1 and 99.
- *file_sytem_objects* is the name of the files or directory to backup.
- *filename* is the name of the file that stores the troubleshoot information.

 **Note:** The *NetWorker Command Reference Guide* provides detailed information about all the available backup options and how to use the `save` command.

Running Recoveries in troubleshoot mode

You can configure NetWorker to log verbose output for recoveries when you Recovery wizard, perform Windows disaster recoveries and by using the `recover` command.

Run Recovery wizard recover jobs in debug mode

You can run recover jobs that you created in the Recovery wizard by using the Recovery wizard or by using the `nsrtask` program from the command line.

Running a recovery job in troubleshoot mode

To send verbose recovery information to the recovery log file, set the troubleshoot level of a recovery job.

Before you begin

Use NMC to connect to the NetWorker server with a user that is a member of the Application Administrators or Database Administrators user group.

Procedure

1. On the **Administration** window, click **Recover**:
 - To modify a scheduled recover job, select the job in the **Configured Recovers** section, and then select **Properties**.
 - To configure a new recover job, select **New**.

 **Note:** You cannot modify an expired or failed to recover job.

2. To create or modify the recover job, use the Recovery wizard. On the **Select the Recovery Options** window, select **Advanced Options**.
3. In the **Debug level** attribute, select a troubleshooting level between 0 and 9.
4. Complete the remaining steps in the Recovery Wizard.

Results

NetWorker logs the troubleshoot recovery information to the recover log file.

Running a recovery job in troubleshoot mode by using nsrtask

Use the `nsrtask` command to run a recovery job that is created by the **Recovery** wizard, from a command prompt.

Procedure

1. On the NetWorker server, type: `nsradmin`.
2. From the `nsradmin` prompt:
 - a. Set the resource attribute to the **Recover** resource:


```
. type: nsr recover
```
 - b. Display the attributes for the **Recover** resource that you want to troubleshoot:


```
print name:recover_resource_name
```

where `recover_resource_name` is the name of the **Recover** resource.
 - c. Make note of the values in the **recover**, **recovery options**, and **recover stdin** attributes. For example:

```
recover command: recover;
recover options: -a -s nw_server.corp.com -c
mnd.corp.com -I - -i R;
recover stdin:
"<xml>
<browsetime>
May 30, 2013 4:49:57 PM GMT -0400
</browsetime>
<recoverpath>
C:
</recoverpath>
</xml>"
```

where:

- `nw_server.corp.com` is the name of the NetWorker server.
- `mnd.corp.com` is the name of the source NetWorker client.

3. Confirm that the `nsrd` process can schedule the recover job:

- a. Update the **Recover** resource to start the recover job:

```
update: name: recover_resource_name; start time: now
```

where *recover_resource_name* is the name of the **Recover** resource.

- b. Exit the `nsradmin` application

- c. Confirm that the `nsrtask` process starts.

If the `nsrtask` process does not start, then review the `daemon.raw` file on the NetWorker server for errors.

4. To confirm that the NetWorker server can run the `recover` command on the remote host, on the NetWorker server type the following command:

```
nsrtask -D3 -t 'NSR Recover' recover_resource_name
```

where *recover_resource_name* is the name of the **Recover** resource.

5. When the `nsrtask` command completes, review the `nsrtask` output for errors.

6. To confirm that the Recovery UI sends the correct recovery arguments to the `recover` process:

- a. On the destination client, open a command prompt.

- b. Run the `recover` command with the recover options that the **Recover** resource uses.

For example:

```
recover -a -s nw_server.corp.com -c mnd_corp.com -I - -i  
R
```

- c. At the **Recover** prompt, specify the value in the `recover stdin` attribute. Do not include the “,” or the “;” that appears with the `recover stdin` attribute.

If the `recover` command appears to stop responding, then review the `daemon.raw` file for errors.

- d. When the `recover` command completes, review the `recover` output for errors. If the `recover` command fails, then review the values that are specified in the **Recover** resource for errors.

7. To review the details of the Recover job, use the `jobquery` command. From a command prompt on the NetWorker server, type: `jobquery`

8. From the `jobquery` prompt, perform one of the following steps:

- Set the query to the **Recovery** resource and display the results of all recovery jobs for a **Recovery** resource:

```
print name: recover_resource_name
```

where *recover_resource_name* is the name of the Recover resource.

- Set the query to a particular jobid and display the results of the job.

```
print job id: jobid
```

Where *jobid* is the jobid of the Recover job that you want to review.

Note: Review the `daemon.raw` file on the NetWorker server to obtain the jobid for the recovery operation.


Running Windows BMR recoveries in troubleshoot mode

Use the WinPE registry to troubleshoot recoveries that are performed with the **BMR Recovery** wizard.

Procedure

1. From a command prompt, type: `regedit`
2. In the Registry Editor, browse to `HKEY_LOCAL_MACHINE\SOFTWARE\JavaSoft\Prefs\com\networker\win/P/E/Wizard`

Figure 104 WinPE registry key to troubleshoot recoveries



3. Change the **Data** value in the `debug_mode` attribute from 0 to 1.
4. Start the BMR Recovery wizard.

Results

The BMR Recovery Wizard logs the troubleshoot information that is related to the following in the `X:\Program Files\EMC NetWorker\nsr\logs\WinPE_Wizard.log` file.

After you collect the troubleshoot information, to turn off troubleshoot mode, modify the data value for the `debug_mode` attribute from 1 to 0.

Running client-initiated recoveries in troubleshoot mode from the command line

To perform a client started backup from the command line, use the `recover` program with the `-D` option.

About this task

For example, on the host you want to recover the data to, type the following command:

```
recover -Dx file_sytem_objects 1>filename 2>&1
```

where:

- *x* is a number between 1 and 99.
- *file_sytem_objects* is the name of the files or directory to recover.
- *filename* is the name of the file that stores the troubleshoot information.

 **Note:** The *NetWorker Command Reference Guide* provides detailed information about all the available recovery options and how to use the `recover` command.

NetWorker Authentication Service logs

This section provides an overview of the log files that are available for the NetWorker Authentication Service.

NetWorker Authentication Service log files

This section provides a summary of the log files available for the NetWorker Authentication Service.

Table 156 NetWorker Authentication Service log files

Component	File name and default location	Description
Installation log	Linux: /opt/nsr/authc-server/logs/install.log Windows: C:\Users\username\AppData\Local\Temp \NetWorker_date_seq_num_AuthC..log	Contains information about the installation of NetWorker Authentication Service.
authc_mgmt and authc_config	Linux: \$HOME/authc-cli.log Where <i>\$HOME</i> is the home folder for the currently logged in user. For example, when the root user runs the command, the file location is /root/authc-cli.log Windows: C:\Program Files\EMC NetWorker\nsr\authc-server\logs \authc-cli.log	Contains a list of error messages that appeared when a user ran the <code>authc_mgmt</code> and <code>authc_config</code> tools.
Authentication server log	Linux: /nsr/authc/logs/authc-server.log Windows: C:\Program Files\EMC NetWorker\nsr\authc\tomcat\logs \authc-server.log	Main authentication service log file.
Audit log	Linux: /nsr/authc/logs/authc-server-audit.log	Contains security audit messages for the NetWorker

Table 156 NetWorker Authentication Service log files (continued)

Component	File name and default location	Description
	Windows: C:\Program Files\EMC NetWorker\nsr\authc\tomcat\logs\authc-server-audit.log	Authentication Service.
Tomcat Access logger	Linux: /nsr/authc/logs/localhost_access_log.date.txt Windows: C:\Program Files\EMC NetWorker\nsr\authc-server\tomcat\logs\localhost_access_log.date.txt	Contains access information for the embedded Apache httpd web server.
Apache Catalina log	Linux: /nsr/authc/tomcat/logs/catalina.out Windows: C:\Program Files\EMC NetWorker\nsr\authc-server\tomcat\logs\catalina.date.log	Contain messages for the Apache Tomcat core component.

Refer to the Apache website for detailed information about the Apache Tomcat log files.

NetWorker Authentication Service server log file management

NetWorker Authentication Services uses the Apache log4j API to manage log files. To modify how NetWorker Authentication Services manage the `authc-server.log` log file, edit the `log4j.properties` file:

- **UNIX:** The `log4j.properties` file is located in `/nsr/authc/webapps/auth-server/WEB-INF/classes`.
- **Windows:** The file is located in `C:\Program Files\EMC\authc-server\tomcat\webapps\auth-server\WEB-INF\classes`.

This section describes how to modify the commonly used log attributes in the `log4j.properties` file. Apache documentation provides more detailed information about each attribute in the `log4j.properties` file.

Note: After you make changes to the `log4j.properties` file, you must stop and start the NetWorker Authentication Service daemon to reset the configuration settings.

Modifying the logging level

The `log4j.rootLogger=` attribute defines the level of logging that the NetWorker Authentication Service writes to the log files and where the messages appear. By default, the NetWorker Authentication Service sets the logging level to `warn` and messages appear in the log files, stdout, and in the Java application. There are five standard log levels: `debug`, `info`, `warn`, `error`, and `fatal`.

To change the logging level to `error`, modify the `log4j.rootLogger=` attribute to appear as follows: `log4j.rootLogger=error, stdout, app`

Modifying the maximum log file size

The `log4j.appender.app.MaxFileSize` attribute defines the maximum size of the `authc-server.log` file. When the log file reaches the maximum size, NetWorker Authentication Service renames the log file for archival purposes and

creates log file. By default, NetWorker Authentication Service sets the maximum size to 100 KB.

To increase the size of the log file to 2MB, modify the *log4j.appender.app.MaxFileSize* attribute to appear as follows:
`log4j.appender.app.MaxFileSize=2MB`

Modifying the number of rollover log files

The *log4j.appender.app.MaxBackupIndex* attribute defines the number of `authc-server.log` rollover log files that the NetWorker Authentication Service maintains. When the size of the `authc-server.log` reaches the maximum file size value, NetWorker Authentication Service copies the contents of the log file to a new log file with the naming convention `authc-serverdate.log`. By default, NetWorker Authentication Service maintains one rollover log file.

To increase the number of rollover log files to 4, modify the *log4j.appender.app.MaxBackupIndex* attribute to appear as follows:
`log4j.appender.app.MaxBackupIndex=4`

CLI log file management

NetWorker Authentication Services uses the Apache log4j API to manage log files. To modify how NetWorker Authentication Services manage the CLI log file, edit the `authc-cli-log4j.properties` file. On UNIX, the `authc-cli-log4j.properties` file is located in `/opt/nsr/authc-server/conf`. On Windows, the file is located in `C:\Program Files\EMC NetWorker\nsr\authc-server\conf`.

This section describes how to modify the commonly used log attributes in the `log4j.properties` file. Apache documentation provides more detailed information about each attribute in the `log4j.properties` file.

Note: After you make changes to the `authc-cli-log4j.properties` file, you must stop and start the NetWorker Authentication Service daemon to reset the configuration settings.

Modifying the logging level

The *log4j.rootLogger=* attribute defines the level of logging that the NetWorker Authentication Service writes to the log files and where the messages appear. By default, the NetWorker Authentication Service sets the logging level to *warn* and messages appear in the log files, stdout, and in the Java application. There are five standard log levels: *debug*, *info*, *warn*, *error*, and *fatal*.

To change the logging level to *error*, modify the *log4j.rootLogger=* attribute to appear as follows: `log4j.rootLogger=error, stdout, app`

Modifying the maximum log file size

The *log4j.appender.app.MaxFileSize* attribute defines the maximum size of the `authc-cli.log` file. When the log file reaches the maximum size, NetWorker Authentication Service renames the log file for archival purposes and creates a log file. By default, NetWorker Authentication Service sets the maximum size to 100 KB.

To increase the size of the log file to 2MB, modify the *log4j.appender.app.MaxFileSize* attribute to appear as follows:
`log4j.appender.app.MaxFileSize=2MB`

Modifying the number of rollover log files

The `log4j.appender.app.MaxBackupIndex` attribute defines the number of `authc-cli.log` rollover log files that the NetWorker Authentication Service maintains. When the size of the `authc-cli.log` reaches the maximum file size value, NetWorker Authentication Service copies the contents of the log file to a new log file with the naming convention `authc-clidate.log`. By default, NetWorker Authentication Service maintains one rollover log file.

To increase the number of rollover log files to 4, modify the `log4j.appender.app.MaxBackupIndex` attribute to appear as follows:
`log4j.appender.app.MaxBackupIndex=4`

NetWorker functionality issues

This section describes workarounds for NetWorker issues.

Backup and recovery

This section covers backup and recovery operations.

Checking the NetWorker services

If you have trouble starting NetWorker programs, the services might not be running correctly. On Windows systems, determine if these processes are running.

If they are not, start them:

- On Windows systems, go to **Control Panel > Administrative Tools > Services**.
- On UNIX systems, type one of the following commands:

```
ps -ef | grep nsr
ps -ax | grep nsr
```

You should receive an output similar to the following:

```
12217 ? S 0:09 /usr/sbin/nsr/nsrexecd -s jupiter
12221 ? S 2:23 /usr/sbin/nsr/nsrd
12230 ? S 0:00 /usr/sbin/nsr/nsrmmdbd
12231 ? S 0:01 /usr/sbin/nsr/nsrindexd
12232 ? S 0:00 /usr/sbin/nsr/nsrmmd -n 1
12234 ? S 0:00 /usr/sbin/nsr/nsrmmd -n 2
12410 pts/8 S 0:00 grep nsr
```

If the NetWorker daemons do not appear, start the NetWorker daemons.

Improper font size for the Client wizard with Netscape on Solaris

When you use the Netscape browser on Solaris, the font size of the **Client** wizard may appear too small.

To change the font type and size:

1. Open the `/usr/bin/nwviz` script file in a text editor.
2. Edit the following line to change the font size:

```
NSR_WIZARD_FONT_SIZE=size
```

3. Save and close the **nwwiz** file.

save: Unable to encrypt data

This message appears during a backup of a Windows host, when the host uses the encryption directive.

The `daemon.raw` file on the Windows host displays the following error message:

```
nsrexecd GSS critical An authentication request from
NetWorker_server was denied. The 'NSR peer information'
provided did not match the one stored by Windows_host. To
accept this request, delete the 'NSR peer information' resource
with the following attributes from Windows_host's NSRLA
database: name: NetWorker_server; NW instance ID: instance_id;
peer hostname: NetWorker_server
```

To resolve this issue, delete the NSR Peer Information resource for the NetWorker server on the Windows host.

Deleting the NSR Peer Information resource

When the local host credentials for a NetWorker host change, authentication attempts from the host to other hosts fail because the credential information stored in the target host does not match the local host credential information that is provided by the initiating host.

Use the `nsradmin` program or the **Local Host** window in NMC to delete the **NSR Peer Information** resource for the initiating host on the target host. The next time the initiating host attempts to connect to the target host, the `nsrauth` authentication process will use the current local host credentials to create a new **NSR Peer Information** resource for the initiating host.

Deleting the NSR Peer Information resource by using NMC

Use NMC to connect to the NetWorker server and delete the NSR Peer Information resource for a NetWorker host.

Before you begin

The account that you use to connect to the NetWorker server must have permission to access the NSRLA database on the target host.

 **Note:** You cannot use NMC to delete the **NSR Peer Information** resource for a NetWorker host that does not have an existing client resource that is configured on the NetWorker server.

Procedure

1. On the **Administration** window, select **Hosts**.

The **Hosts Management** window appears.

2. Right-click the NetWorker host with the **NSR Peer Information** resource that you want to delete, and then select **Host Details**.

 **Note:** The NetWorker host does not appear in the **Local Hosts** section when a client resource does not exist on the NetWorker server.

The **Certificate** window displays a list of NSR Peer Information resources stored in the `nsrexecd` database on the host.

3. In the **Certificate** pane, right-click the certificate that you want to delete, and then select **Delete**.

4. When prompted to confirm the delete operation, select **Yes**.

If you receive the error, `User username on machine hostname is not on administrator list`, you cannot modify the resource until you configure the NSRLA access privileges on the target host. The section "Configuring NSRLA access privileges" provides more information.

Results

The target host creates a new NSR Peer Information resource for the initiating host the next time that the initiating host attempts to establish a connection with the target host.

Deleting the NSR Peer Information resource by using `nsradmin`

To delete the NSR Peer Information resource for the initiating host, use the `nsradmin` command on the target host.

Before you begin

Connect to the target host with an account that has administrator access to the NSRLA database. You must configure access privileges to the NetWorker client database.

Procedure

1. Connect to the `nsrexec` database:

```
nsradmin -p nsrexec
```

2. Set the query type to the **NSR Peer Information** resource of the initiating host:

```
. type: nsr peer information;name:initiating_host_name
```

For example, if the hostname of the initiating host is `pwd.corp.com`, type:

```
. type: nsr peer information;name: pwd.corp.com
```

3. Display all attributes for the **NSR Peer Information** resource:

```
show
```

4. Print the attributes for the **NSR Peer Information** resource and confirm that the name and peer hostname attributes match the hostname of the initiating host:

```
print
```

5. Delete the `NSR Peer Information` resource:

```
delete
```

6. When prompted to confirm the delete operation, type `y`.

7. Exit the `nsradmin` program:

```
quit
```

Results

The target host creates a new NSR Peer Information resource for the initiating host the next time that the initiating host attempts to establish a connection with the target host.

Backups fail to start when the daylight savings time change occurs

When you schedule backup operations to occur during the hour in which the operating system moves the clock ahead or behind by one hour, NetWorker skips the backup operation. For example, the operating system is configured to move the clock forward one hour at precisely 2:00 A.M. and backups are scheduled to occur at 2:01 A.M. At 2:00 A.M., the operating system moves the clock forward to 3:00 A.M. NetWorker will skip all backup operations that are scheduled to start between 2:01 to 2:59 and NetWorker does not initiate the backup operation.

To avoid this situation, set the backup time to occur at least one minute before the time change occurs.

Note: When you use the `mminfo` command to get a weekly save set usage summary for the time period during the change to daylight savings time, `mminfo` does not display any information for the day of the change.

Shut down NetWorker services prior to any significant changes to system date

If you need to make a significant change to the system clock or date, for example, a change of more than a day, then ensure that you shut down the NetWorker services before you make the change. NetWorker services depend heavily on the system clock for many operations such as active sessions, volume mount and unmount operations, the expiration of save sets, and license enforcement.

Clone ID timestamp does not reflect the time the clone was created

To guarantee that the cloned save sets that NetWorker creates on different storage nodes do not have the same timestamp, the NetWorker software assigns a timestamp to cloned save sets that does not reflect the actual time that NetWorker creates the clone.

Memory usage when browsing large save sets

When you use the NetWorker User program to browse or perform a browsable recover from a large save set, such as a save set with one million or more files, the operation may consume all the memory on the host.

To avoid this issue, perform one of the following options:

- Perform a save set recovery.
- Use the `recover` command, which enables you to directly browse the client file index and select the files and directories that you want to recover. Use this option to browse large save sets or when memory is limited on the host systems.

Memory usage and `nsrjobd`

The `nsrjobd` daemon runs on the NetWorker server and is responsible for monitoring NetWorker activity during a backup or recovery operation. Depending on the size of your backup environment, `nsrjobd` can require large amounts of RAM.

Media position errors encountered when auto media verify is enabled

To verify media, the `nsrmmmd` process must reposition the volume to read previously written data.

The first try may not always succeed and the following warning messages appear in the message window of the NetWorker **Administration** window:

```
media warning: /dev/rmt2.1 moving: fsr 15: I/O error
media emergency: could not position jupiter.007 to file 44,
record 16
```

If the server can find the correct position, media verification succeeds and a successful completion message appears:

```
media info: verification of volume "jupiter.007" valid 30052
succeeded.
```

If the media verification fails, then perform the following tasks:

- Reset the device.
- Verify the device configuration.
- Verify that NetWorker can recognize the media.
- Verify that the device operations function correctly.

The scanner program marks a volume read-only

When you use the `scanner` program to rebuild the index of a backup volume, the `scanner` program marks the volume as read-only.

This is a safety feature that prevents NetWorker from overwriting the last save set on the backup volume.

Use the `nsrmm` command change the volume to write-enabled:

```
nsrmm -o notreadonly volume_name
```

The scanner program requests an entry for record size

If you use the `scanner` program with the `-s` option but without an `-i` or `-m` option, a message similar to the following may appear:

```
Please enter record size for this volume ('q' to quit)
```

If this message appears, specify a block size that is greater than or equal to 32.

Limitations for groups containing a bootstrap

NetWorker only writes bootstrap backups to a local device. When a group backup generates a bootstrap save set, ensure that device attached to the NetWorker server has an available volume for the bootstrap backup.

Index recovery to a different location fails

If you try to recover indexes to a directory that differs from the original location, an error message similar to the following appears:

```
WARNING: The on-line index for client_name was NOT fully
recovered. There may have been a media error. You can retry the
recover, or attempt to recover another version of the index.
```

To resolve this issue, ensure that you recover indexes to the original location then move the indexes to another directory. [Moving a client file index](#) describes how to move indexes to another directory.

Illegal characters in configurations

When you provide a name for label templates, directives, groups, policies, and schedules, do not use the following characters:

```
/ \ * [ ] ( ) $ ! ^ ' " ? ; ' ~ < > & | { }
```

Inaccessible object exception error when launching NMC with Java 9

When you are running NMC with Java 9, your system displays a number of `InaccessibleObjectException` errors for `JButton` and `BasicSplitPaneUI` objects in Console logs. You need to add JVM option `--illegal-access=permit` in **Java Control Panel** of Java 9 to avoid `InaccessibleObjectException` errors, when launching NMC.

Error backing up large number of clients

When you use a Windows NetWorker server to back up many clients, a `CMD.exe` application error window may appear with a message similar to the following:

```
The application failed to initialize properly (0xc0000142).
Click on OK to terminate the application.
```

If this problem occurs, edit the Windows registry on the NetWorker server to increase the desktop heap allocation.

1. In the **regedt32** application, browse to the following registry entry:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\
Session Manager\SubSystems\
```

2. Edit the **Windows** registry key.

3. Modify the third value of the **SharedSection** entry to increase the heap allocation size.

In the following example, the desktop heap allocation has been changed from a value of 512 KB to 1023 KB.

The original entry, with a desktop heap allocation of 512 KB appears as:

```
%SystemRoot%\system32\csrss.exe ObjectDirectory=\Windows
SharedSection=1024,3072,512 Windows=On SubSystemType=Windows
ServerDll=basesrv,1
ServerDll=winsrv:UserServerDllInitialization,3
ServerDll=winsrv:ConServerDllInitialization,2
ProfileControl=Off
MaxRequestThreads=16
```

The updated entry, with a desktop heap allocation of 1024 KB appears as:

```
%SystemRoot%\system32\csrss.exe ObjectDirectory=\Windows
SharedSection=1024,3072,1024 Windows=On SubSystemType=Windows
ServerDll=basesrv,1
ServerDll=winsrv:UserServerDllInitialization,3
ServerDll=winsrv:ConServerDllInitialization,2
ProfileControl=Off
MaxRequestThreads=16
```

4. Restart the NetWorker server.

The Microsoft Knowledge Base article 18480 on the Microsoft website provides more information.

Hostname aliases

When you incorrectly define an alias for a client, the backup fails. Under certain conditions, such as improperly configured DNS servers or hosts files, the NetWorker software does not create any aliases for a new client. If you use TCP/IP, ensure that you specify the hostname and the fully qualified domain name (FQDN) for a client in Aliases field of the client properties window.

When the alias field is incorrectly defined you can see the following behavior:

- Backup operations for the client fail with the following error message:
`No Client resource for client_name`
- NetWorker always performs backups for a client at a level full, regardless of the level of the scheduled backup.
- Automatic index management, as set up in the browse and retention policies, does not work.
- The `/nsr/index` directory, which contains the indexes for all the clients that are configured on the NetWorker server contains two directories for the same client, but each directory uses a different client name.

You must update the alias field for a client to include all hosts names for the client in the following situations:

- When a host have two or more network interfaces.
- When sites use a mixture of short and FQDNs for the same servers, for example, mars and mars.jupiter.com.
- When the datazone uses both (Network Information Services (NIS) and DNS.

 **NOTICE** Do not include aliases that are shared by other hosts in the datazone.

Directory pathname restrictions

A file manager (but not Windows Explorer) restriction causes errors when a pathname contains too many characters.

To avoid these errors, use a pathname that has fewer than 128 characters.

Backup of a new client defaults to level full

The first time you back up a new client, a message similar to the following appears:

```
client: save: There are no save sets in the media database;
performing a full backup
```

This message indicates that NetWorker has not previously performed a back up of the specified save set. Before you can perform an incremental or level backup on a save set, perform a full backup of the save set.

If a level full backup exists for this save set, this error message can appear in the following situations:

- The clocks on the client and server are not synchronized.
- The `savegrp` session begins before midnight and ends after midnight.
- Multiple client ids exist for the client.

Non-full backup of Solaris files with modified extended attributes

When you change the extended attributes for a Solaris file, but you do not modify the file, the action does not update the change time (ctime) for the file. As a result, the NetWorker software does not know that the extended attributes for the file have changed since the last incremental backup, and any non-full scheduled backup of the file system will not back up the file.

To ensure the file is backed up, use the `touch` command or otherwise modify the file so that the ctime is updated. Alternatively, perform a manual backup of the file. [Manual backups](#) on page 488 provides more information.

Client file index errors

This section provides information about issues are related to client file indexes.

Renamed clients cannot recover old backups

The NetWorker server maintains a client file index for every client that has been backed up. When you change the name of the client, NetWorker uses the new hostname to create a new client file index, as a result you cannot recover files that were backed up under the old client name.

To recover data that was backed up by using the old client name, perform a directed recovery and specify the old client name as the source host and the new client name as the destination host. [Directed recoveries](#) on page 546 provides information about how to perform directed recoveries.

Missing client file indexes

Before you use the `scanner` program with the `-i` option, ensure that a client file index exists for the client that is associated with each save set. If you try to recover a client file index with the `scanner -i` command when the client file index does not exist, a message similar to the following appears:

```
scanner: File index error, file index is missing.
Please contact your system administrator to recover or recreate
the index.
(severity 5, number 8)
scanner: write failed, Broken pipe
scanner: ssid 25312: scan complete
scanner: ssid 25312: 91 KB, 13 file(s)
scanner: done with file disk default.001
```

To resolve this issue, use the `nsrck -l2 clientname` to create a client file index for the client, then try the `scanner` command again.

Check failure of client file indexes

Each time the NetWorker server starts, the startup process uses a `nsrck -ML1` command to perform a level 1 consistency check on the client file indexes. In some circumstances, this consistency check does not detect corruption in the client file indexes. If you believe that the NetWorker server may have a corrupted client file index, run a higher level check on the index.

For example:

```
nsrck -L5
```

If the command does not resolve the index corruption, refer to [Adding information about recyclable save sets to the client file index](#) on page 559 for more information.

No notification of client file index size growth

The NetWorker Server does not notify you when a client file index is getting too large. Monitor the system regularly to check the size of client file indexes. [Reduce the size of the client file index](#) on page 812 provides information about how to manage the NetWorker client file indexes.

The *NetWorker Command Reference Guide* or the UNIX man pages provide more information about how to use the `nsrls`, `nsrck`, and `nsrim` commands to monitor and manage client file indexes.

Aborting a recovery

When you stop a recovery operation on a client, the following could occur:

- The recovery might stop immediately.
- The recover program will display a list of the files that were not recovered.
- Messages similar to the following appears, which indicates that the recovery operation did not stop cleanly:

```
Recover: ***Canceled***
Recover: Unable to read checksum from save stream
Recover: error recovering C:\WINDOWS\CURSORS\APPSTART.ANI
Didn't recover requested file C:\WINDOWS\CURSORS
\APPSTART.ANI
```

xdr of win32 attributes failed for *directory*

This error appears when the backup operation cannot back up the directory path. The rest of the save set completes successfully.

To resolve this problem, perform another backup of the directory.

Cannot create directory *directory*

This error message appears when you attempt to relocate data to a directory that does not exist on the target host. You can ignore this message. The recovery process creates the new directory and completes successfully.

The All save set and duplicate drive serial numbers

The All save set, which backs up all locally mounted drives and the VSS SYSTEM save sets, uses the serial numbers assigned to drives as part of the backup logic that

determines when the backup operation should include a drive. If more than one local drive uses the same serial number, the All save set will only back up one of the drives.

To resolve this issue, perform one of the following solutions:

- Use the `DiskProbe` utility to set the serial numbers to unique numbers. The `DiskProbe` utility is part of the Windows Support Tools and is available for all versions of Windows supported by NetWorker software.
- Avoid using the All save set. Instead, specify each drive letter or the VSS SYSTEM save set separately. [The DISASTER_RECOVERY:\ save set](#) on page 433 provides more information about the All save set.

No disk label errors

This error message appears when you configure a non-optical device as an optical device.

To resolve this issue, verify that the Media Type attribute in the Device resource matches the expected media for the device, and correct if necessary.

Resolving copy violation errors

If you install the NetWorker server software on multiple hosts and more than one server uses the same NetWorker enabler code, a message similar to the following appears in the policy reports and notifications:

```
--- Unsuccessful Save Sets ---
* mars:/var save: error, copy violation - servers 'jupiter' and
'pluto' have the same software enabler code, 'a1b2c3d4f5g6h7j8'
(13)
* mars:/var save: cannot start a backup for /var with NSR
server 'jupiter'
* mars:index save: cannot start a backup for /usr/nsr/index/
mars with NSR server 'jupiter'
* mars:index save: cannot start a backup for bootstrap with NSR
server 'jupiter'
* mars:index save: bootstrap save of server's index and volume
databases failed
```

To resolve this issue, perform one of the following actions:

- Remove the NetWorker server software from all hosts but one.
- Contact Dell EMC Licensing and request new licenses for each additional NetWorker server.

 **Note:** After you perform one of the resolutions, stop and then restart the NetWorker services on the NetWorker server that performs the backups.

Converting sparse files to fully allocated files

The NetWorker server determines which files are sparse by comparing the allocated blocks with the byte size. If the allocated blocks do not account for the size of the file, NetWorker considers the file as sparse and the save operation replaces long strings of zeros with “holes” in the recovered file.

A recovery operation may recover some files as sparse when the files were not sparse at the time of the backup operation. Oracle databases are susceptible to this problem because they are zero-filled, fully allocated files, but are not sparse.

To workaroud this issue, use the `cp` command to copy the file after recovery:

```
cp recovered_filename zero_filled_filename
```

This command converts a sparse file to a fully allocated file.

 **NOTICE** Ensure that you have enough free disk space to accommodate a duplicate of each copied sparse file.

Backing up large sparse files

To conserve backup media, NetWorker compresses sparse files before the save operation writes the file to the backup media. While NetWorker compresses the file, the backup job may stop and the following message appears:

```
savegrp: Aborting inactive job (633).
```

This can occur when the backup operation does not write any data to the backup media during the compression operation and time the backup is idle reaches the time that is specified by the group Inactivity Timeout attribute. To resolve this issue, increase the Inactivity Timeout attribute for the backup group.

To help determine an adequate timeout limit:

1. Set the **Inactivity Timeout** value to zero. A value of zero results in no timeout limit.
2. Determine the time that the backup requires to complete a full save of the file system, and specify this time as the inactivity timeout limit.

Queries using the `mminfo -N` command are case-sensitive

When you use the `mminfo` command to query the media database, the `-N name` option is case-sensitive. The save set name the `-N` option references must match the case of the save set name that you specify in the save set attribute of the client resource.

However, when NetWorker performs a back up of drive partitions on Windows (for example, `C:\`), the NetWorker server stores the save set name in uppercase in the media database.

For example, if the save set name that represents the drive partition was typed in the client resource is lowercase `c`, you must use uppercase `C` to query the media database:

```
mminfo -N C:\
```

Renamed directories and incremental backups

By default, if the name of a directory changes after a full backup, but no files or subfolders in the directory change, NetWorker will not include the renamed directory in subsequent incremental backups.

To include renamed directories in an incremental backup, select the Backup renamed directories attribute in the Client resource.

 **Note:** NetWorker will only backup renamed directories with unchanged files and subfolders only when you explicitly list directory names in the save set attribute of the Client resource.

For example, if the save set field contains `E:\` and you rename the `E:\test` directory to `E:\test1`, NetWorker does not back up the `E:\test1` directory when you enable **Backup renamed directories**. When the save set field contains

`E:\test` and you rename the `E:\test` directory to `E:\test1`, NetWorker performs a backup of the `E:\test1` directory when you enable **Backup renamed directories**.

Resolving names for multiple network interface cards

If any NetWorker host (client, storage node, server) has multiple network interface cards (NICs) with unique IPs and host names, you must configure all NICs and ensure that the host names are resolvable, even if the host does not use one or more of the NICs. Failure to have all NICs resolvable may cause problems with host connectivity to the NetWorker server.

Follow these steps to ensure that NetWorker uses the appropriate hostname for an IP address, and to ensure that you properly configure the hosts file and routing table on the host:

- Set up DNS to associate each IP address with a separate name.
- Configure the hosts file and routing table on each host that has multiple NICs with the appropriate IP address.
- Configure NetWorker to use the names that you configured in the DNS and hosts file.

Configuring multiple NICs

In the following example, a dual-interface client connects to the NetWorker Server and Storage Node over interface1 which has an IP address of 1.1.1.1 and has a dedicated connection to the Storage Node over interface2 with an IP address of 2.2.2.1. The user wants to send all data to the Storage Node over interface2 instead of the default interface1.

1. Configure DNS with unique host names for IP addresses **1.1.1.1** and **2.2.2.1**. For example, **client-1** maps to **1.1.1.1** and **client-2** maps to **2.2.2.1**. DNS should also be configured with unique host names for the IPs on the Storage Node. For example, **node-1** maps to **1.1.1.2** and **node-2** maps to **2.2.2.2**.
2. Configure the routing table on the client to route the traffic through the correct interface, and then add the two IP addresses to the local hosts file.
3. On the NetWorker server, enter **node-2** in the Storage Node Affinity List of the client. [Configuring the client's storage node affinity list](#) provides more information.

The section on [Configuring NetWorker in a multihomed environment](#) provides more details.

Libraries entering ready state

When you start the NetWorker service or after you configure a tape library, the library does not immediately enter the Ready state within NetWorker. This is normal behavior.

Successful save sets listed as failed in the Group Backup Details window

Certain backup operations, such as on some NetWorker modules, create multiple sessions to perform a single backup job. If one of these sessions fails, the Console reports that the entire backup job has failed.

To determine the status of each session, click the **Show Messages** button in the **Failed** table of the **Savegroup Completion** dialog box. This information also appears in the **Logs** tab, under monitoring, and in the savegroup completion report.

The NetWorker Server window does not appear on HP-UX

On HP-UX, the following error message appears after you log in to the NMC server and attempt to connect to a NetWorker server:

Unable to connect to server: Failed to contact using UDP ping

To resolve this issue:

1. In the NetWorker Console, select **Setup**.
2. Select **Setup > System Options**.
3. Unselect the **RPC ping via UDP when connecting to NetWorker** checkbox.

Backup fails with Win32 error 0x2

The scheduled backup of a physical node fails, if the backup start time of the virtual server and the physical node is the same. The following error is displayed:

```
WINDOWS ROLES AND FEATURES: ERROR: Failed to save Backup Comp
Doc, C:\Program Files\EMC NetWorker\nsr\tmp\backcomp.xml
90110:save: save of 'WINDOWS ROLES AND FEATURES:\' to server
'<server_name>' failed: The system cannot find the file
specified.(Win32 error 0x2)
```

To resolve the issue, ensure that the backup start time of the virtual server and the physical node is different.

Error displaying workflow details

From the NMC GUI, when you right-click a workflow and select **Show Details**, the following error message appears:

```
Unable to obtain root job list: Unable to accept message bus block,
while expecting basic-consume method, received another method.:
Message bus encountered library exception: (unknown error)
```

In addition, the `gstd.raw` file had messages indicating that there is an error with the RabbitMQ server.

```
gstd NSR warning nsm: Initial jobs info RPC failed for server:
<NetWorkerServer> - Check if RabbitMQ server is running!
```

This can occur if RPC communication with the RabbitMQ server fails because of missing exchanges such as `amq.direct`, `amq.fanout`, `amq.headers`, `amq.match`, and so on.

To resolve the issue, do the following:

1. Shutdown NetWorker.
2. From the command prompt, type the following command to go to the directory that contains the RabbitMQ files:
 - Windows: `cd C:\Windows\system32\config\systemprofile\AppData\Roaming`
 - Linux: `/opt/nsr/rabbitmq-server-3.2.4/var`
3. Type `ren RabbitMQ RabbitMQ.v01` to save the existing RabbitMQ directory.

4. Start NetWorker.

After NetWorker starts, new RabbitMQ files are created.

- **Windows:** C:\Windows\system32\config\systemprofile\AppData\Roaming\RabbitMQ
- **Linux:** /opt/nsr/rabbitmq-server-3.2.4/var

To verify if the missing exchanges are present, run the following commands:

Windows:

```
% set ERLANG_HOME=C:\PROGRA~1\EMCNET~1\nsr\rabbitmq-server-3.2.4
% set HOMEPATH=\Windows
% rabbitmqctl report
```

Linux:

```
# HOME=/nsr/rabbitmq /opt/nsr/rabbitmq-server-3.2.4/sbin/
rabbitmqctl report
```

Back up of All Save Sets takes a long time to complete

Back up of **All Save Sets** might be very slow and can take a significant amount of time to complete.

This can be because of your antivirus software configurations. Check your antivirus configuration settings.

GSS-API authentication error

If there is a change in the IP address of the server, storage node, or the client, you might encounter the following authentication error:

```
Authentication error; why = GSS-API credential problem
```

This error is resolved automatically after about an hour. There is no action required from the user.

NetWorker locale and code set support

The NetWorker software does not support locales that are defined by the operating system or code sets that remap characters, which have a special meaning for file systems. Depending on the file system, the special characters may include the slash (/), the backslash(\), the colon (:), or the period(.). `De_DE.646` is an example of one unsupported locale.

The NetWorker software might not function normally after you change the locale to an unsupported locale. Client file indexes that were created in a supported locale can become inaccessible.

Enabling service mode for NetWorker

To enable and disable access to the NetWorker Server, use the **Accept new sessions** and **Accept new recover sessions** attributes in NMC. When you unselect these attributes, the server does not accept new backup and recovery sessions.

The *NetWorker Security Configuration Guide* provides more information about these attributes.

When you restrict NetWorker Server access, NetWorker takes all storage nodes offline, effectively putting NetWorker into a service mode operational state. In this state, you can stop any external client backup and recovery requests and prevent the start of scheduled group backups. Service mode provides you with a maintenance period where you can diagnose and troubleshoot issues before you return the server to normal operation.

You can also enable/disable specific storage nodes or devices to prevent use and allow for service operations. [Storage node configuration](#) on page 99 describes how to enable/disable specific storage nodes. [Re-enabling a device](#) on page 173 describes how to enable/disable a specific device.

Resetting the NMC password for NetWorker 19.1 and later

To reset the administrator password, create a JSON file on the NetWorker server that contains the new password in a Base64 encoded format.

Procedure

1. Use Base64 encoding utilities to determine the Base64 password value for the new password.
2. Use a text editor to open the `authc-local-config.json.template` file, which is located in the `C:\Program Files\EMC NetWorker\nsr\authc-server\scripts` folder on Windows and the `/opt/nsr/authc-server/scripts` directory on Linux.
3. In the template file, perform the following steps:
 - a. Replace the `username` variable with the name of the administrator account for which you want to reset the password.
 - b. Replace the `encoded_password` variable with the base64 encoded password value.
4. Rename the `authc-local-config.json.template` file to `authc-localconfig.json`.
5. Copy the `authc-local-config.json` file to the Tomcat conf folder.
6. Stop and then start the services on the NetWorker server (NetWorker service and GST service).

No privileges to view NetWorker server from NMC

The NetWorker administrator will not have the privilege to view the NetWorker resources if there is a mismatch in the external roles in the NetWorker user group and the external role displayed in NMC. The user should ensure that there is no mismatch in the external roles attribute which is displayed in NMC and the NetWorker user group attribute. When the NetWorker server resolves to more than one hostname, (when it has more than one alias), there are chances that the NetWorker server has created an external role with one of the hostname. Ensure that the NetWorker server

is resolved to only one hostname and hosts entry for the NetWorker server is same in NMC and NetWorker server hosts.

Network and server communication errors

This section provides general, UNIX and Windows network and communication issues that you may encounter in a NetWorker environment.

The errors that are encountered during registering a new client are as follows:

- The NetWorker sever and client running on the same box- During the registration of the client, register for the client hostname and not for the client IP address. If the Client IP is registered, then during recovery, based on client IP address, recovery doesn't display the backed up save set list in the browse tab.
- Networker server and client running on different box- Client IP address can be registered and recovery also proceeds without any issues.

To help ensure successful communication between NetWorker clients and servers, each NetWorker host configured must not have any invalid or inactive IP addresses stored in the hostname resolution service (DNS, NIS, Active Directory, hosts file, and so on). Each IP address that maps to a host must have a configured network interface (NIC).

Unapproved server error

If an unapproved server tries to contact a client to start a backup, a message similar to the following appears: `client_name: server_name cannot request command execution`.

To provide additional servers access to the NetWorker client, perform the following steps:

1. Modify the `servers` file on the client and ensure that the file contains both the short name and the long name of the server. For example, the `servers` file on a NetWorker client should contain these names for a NetWorker server that is named `mars` in the `jupiter.com` domain:

```
mars
mars.jupiter.com
```

2. In the **Alias** attribute of the Client resource, specify both the short name and the long name, and any other applicable aliases for the client.

Unapproved server error during client setup

If you add a Windows client to a UNIX NetWorker server, and the `servers` file on a Windows client does not include the UNIX server hostname, the message similar to the following may appear:

```
client_name: saveset_name Host server_name cannot request
command execution
client_name: saveset_name 10/13/00 11:48:26 nsrexec: Host
server_name cannot request command execution
client_name: saveset_name Permission denied
```

Ignore the message, and continue to add the client to the UNIX server. To avoid the message, add the UNIX server hostname to the `servers` file on the client after you add the client to the UNIX server.

Server copy violation

When the **Alias** attribute of the Client resource for the NetWorker server does not contain all of the host names or aliases for the NetWorker server, the NetWorker server may become disabled and an error message similar to the following appears:

```
nsrd: registration info event: server is disabled copy
violation
```

To resolve this issue, add all of the server aliases that are related to any additional network interfaces to the alias list of Client resource for the NetWorker server.

Remote recover access rights

You can control client recover access with attributes in the Client resource. The **Remote Access** attribute displays a list of the users that can recover save sets for a client. Add or remove user names depending on the level of security the files require.

Note: If you type a hostname or `host=hostname` in the **Remote Access** attribute, you allow any user on that host to recover files for the client. To enter a username without specifying the host, type `user=name`.

The following users have permission to recover any files on any client, regardless of the users who are listed in the **Remote Access** attribute:

- 'Root' user on a UNIX host
- Member of the 'Administrators' local group on a Windows host
- Members of a 'Application Administrator' User group on the NetWorker Server
- Members of a NetWorker Server User group that has the 'Change Security Settings' privilege

Other users can only recover files for which they have read permission, which is based on file permissions at the time of backup. Files recovered by a user other than root, operator, or the operator group are owned by that user.

NetWorker server takes a long time to restart

The consistency check of the media database, which occurs when the NetWorker server services start, can take a significant amount of time to complete when the media database is very large. While the NetWorker server performs the consistency check, client connections with the NetWorker server are delayed.

To reduce the size of the media management database, run the `nsrim -C` command when the NetWorker server is idle. Be aware that this command may take a long time to run and that the NetWorker server will be unavailable during this time. Run the command when the NetWorker server is not busy.

Note: The `nsrim -C` command can take a long time to complete and you cannot perform NetWorker server operations until the command completes.

[Reduce the size of the media database size](#) on page 813 provides more information about reducing the size of the media database.

Changing the NetWorker server address

When the IP address changes on the NetWorker server, the NetWorker hostid also changes. The authorization code assigned to each NetWorker license depends on the hostid. When the hostid of the NetWorker server changes, you must contact Dell EMC

Licensing to generate new authorization codes based on the new hostid, then update each NetWorker license with the new authorization code.

If you do not re-register the software with the new authorization codes within 14 days of the hostid change, the NetWorker becomes disabled and you cannot perform any operations with the exception of recovery operations.

 **Note:** If you are using DHCP, use a static IP address for the NetWorker server.

Binding to server errors

NetWorker architecture follows the client/server model, where the NetWorker servers use RPC to provide services to the client. These services reside in daemon processes.

When the daemons start, they register with the registration service provided by the portmapper.

If the NetWorker services are not running and an operation requests a NetWorker service, a message similar to the following may appear in the savegroup completion email:

```
Server not available
RPC error, no remote program registered
```

These messages indicate that one or more NetWorker services are not running on the NetWorker server. The following table summarizes the startup commands that you can use to startup the services on a UNIX NetWorker server.

Table 157 NetWorker Startup commands

Operating system	Startup command
Solaris, Linux	/etc/init.d/networker start
HP-UX	/sbin/init.d/networker start
AIX	/etc/rc.nsr

Identifying connection lost to device

When a device loose connection with Networker server, the Device Connectivity Check (DCC) framework discovers device unavailability. Upon discovery of the device unavailability, NetWorker server notifies the user through NMC console, Networker logs, or by email.

The various scenarios where the devices becomes unavailable are as follows:

- Data Domain hosting the devices becomes unavailable.
- Data Domain administrator accidentally deletes the MTree having the devices.
- The mount point hosting the AFTD devices becomes offline.
- The AFTD device folder is deleted accidentally from the OS file system.

 **Note:** Starting with NetWorker 19.2, DCC framework will be enhanced to verify the existence and availability of DD, DDCT, and AFTD devices.

The Networker DCC framework provides the following capabilities:

- The DCC framework periodically checks the device connectivity.
- The DCC feature is enabled by default. It is visible in the diagnostic mode under NMC console.

- The NetWorker storage node reports DCC result to NetWorker server.
- NetWorker server parses the report and moves the device either to normal or suspected state and informs the broker about the status of the device.
- If the suspect status of the device has changed, NetWorker server reports this event to customer, through NMC console, NetWorker logs or through email.

The supported NetWorker device types are:

- DD and DDCT device
- AFTD device.

The limitations are as follows:

- DCC does not check the existence of savesets residing in a volume. However, it checks the existence of `VolHdr` file created inside the volume to determine the device availability and accessibility.
- In case of DDCT devices, NetWorker metadata is still stored in the cloud tier volume and the actual saveset file might have moved to the cloud storage. DCC checks the availability and accessibility of only the DDCT device and not the cloud unit configured which may or may not be available.

Reasons for Device unreachability

There are three reasons for device unreachability and in each case we update the device state accordingly.

- DD is unavailable: If the DD is unavailable, DCC marks all the devices under that DD to Suspected state automatically. If the DD is available again, DCC marks all the devices under that DD to Normal state automatically.
 ⓘ **Note:** DCC sends notifications to NMC or logs (server and storage node) for the above two scenarios. DCC does not check the existence of M-Tree or devices if the DD is unavailable.
- M-Tree is unavailable: If M-tree is deleted, all the devices under the M-Tree is marked to Suspected state. If M-tree is undeleted, all the devices under the M-Tree is moved to Normal state from Suspected state.
 ⓘ **Note:** DCC sends notifications to NMC or logs (server and storage node) for the above two scenarios. DCC does not check the existence of the device if M-Tree is deleted.
- DD and AFTD device unavailability: By default DCC does not check for the device availability if the device is in Disabled or Service Mode state. DCC checks for the existence of `volhdr` in device directory. If `volhdr` exists device is considered as Available, otherwise it is considered as Unavailable. If the devices are unavailable DCC sends notification to NMC or logs (server and storage node)
 ⓘ **Note:** DCC feature works only when both NetWorker server and storage nodes are on 19.1 version. Starting with NetWorker 19.1, the default value of DCC interval is 3 minutes.

Device Connectivity Check Configurations

Device Connectivity Check is enabled on both server and storage node by default. The default time interval is 3 minutes and it can be extended up to 60 minutes. DCC

settings can be edited either through nsradmin CLI (hidden) or through NMC in diagnostic mode.

Settings on Server CLI and Storage Node CLI

Figure 105 Settings on Server CLI

```
Policy migration attempted: Yes;
Policy migration status: Not applicable;
Policy migration logs: ;
    VMWS enable: Yes [No];
    VMWS user name: VMUser;
    VMWS user password: *****;
    VMWS port: 8080;
Authentication config date: "Wed Apr 10 12:20:34 2019";
Authentication Service port: 9090;
Authentication Service database: /nsr/authc/data;
Authentication Proxy port: 7999;
    Disable RPS Clone: [Yes] No ;
device connectivity check enabled: [Yes] No ;
device connectivity check interval: 3;
vmware policy allowed: [Yes] No ;
```

Figure 106 Settings on Storage Node CLI


```
dedicated storage node: Yes [No];
snode device connectivity check enabled: [Yes] No ;
snode device connectivity check interval: 3;
shared device creation: Yes [No];
    version: 99.0.99.9985;
    daemon versions: "nsrsnmd:99.0.99.9985.Build.9985",
                    "nsrexecd:99.0.99.9985.Build.9985"
type of storage node: [scsi] ndmp silo ;
storage node is configured: [Yes] No ;
date of registration: "Wed Apr 10 12:20:53 2019";
enabled: [Yes] No ;
```

Example 14

NMC Server and Storage Node Settings

To set the server in NMC, select **Server properties** in the **Configuration** tab (Diagnostic mode).

Figure 107 NMC Server settings



The screenshot shows the 'Operations' section of the NMC interface. It contains two settings: 'Device connectivity check enabled' with a green checkmark icon, and 'Device connectivity check interval' with a dropdown menu set to '3'.

To set storage node in NMC, select **Storage node properties** in the **Configuration** tab (Diagnostic mode).

Figure 108 NMC Storage Node Settings

The screenshot shows a dialog box titled 'Operations'. Inside, there are two settings: 'Snode device connectivity check enabled' which has a green checkmark in a box, and 'Snode device connectivity check interval' which is a text box containing the number '3' with up and down arrow buttons. At the bottom of the dialog are three buttons: 'OK', 'Reset', and 'Cancel'.

DCC can be enabled or disabled based on the settings of server and storage Node.

Table 158 Enabling and Disabling DCC based on Server and Storage node settings

DCC check on server	DCC check on storage node	Resultant DCC
On	Off	Off
On	On	On
Off	On/Off	Off
Interval 1	Interval 2	Whichever is greater

When DCC identifies any unreachable device due to various reasons, it moves the device to suspected state and it sets two attributes Suspected device and Suspected mark time.

- Suspected device and suspect time attributes are set when the device state is changed to suspected state from normal state and these attributes are cleared when the device is available.
- If there is any device that is in suspected state and is moved to service mode manually, suspected flag is cleared off and there won't be any DCC check on that going forward. DCC starts checking this device only if it again moves to enabled state.

New.Net and NetWorker software are incompatible

Software from New.Net, Inc. loads a dynamic link library (DLL) named newdotnet.dll, which modifies the Windows TCP/IP stack in ways that are incompatible with NetWorker software.

This causes many NetWorker programs, including save.exe, to fail on exit. This is a New.Net problem that the NetWorker software cannot work around. The Go!Zilla, BearShare, Mp3.com, iMesh, Babylon, Cydoor, Webshots, and gDivx products include the New.Net software. If you suspect that the New.Net DLL is the cause of problems, uninstall the New.Net software.

NOTICE If you manually delete the newdotnet.dll file, the system will become unusable.

GLOSSARY

This glossary provides definitions for terms used in this guide.

A

access control list (ACL)	List that specifies the permissions assigned to a specific file or directory. See administrator .
active group	NetWorker backup group that has its Autostart attribute enabled.
administrator	Person who normally installs, configures, and maintains software on network computers, and who adds users and defines user privileges.
Administrators group	Microsoft Windows user group whose members have the rights and privileges of users in other groups, plus the ability to create and manage the users and groups in the domain.
advanced file type device (AFTD)	Disk storage device that uses a volume manager to enable multiple concurrent backup and recovery operations and dynamically extend available disk space.
agent	Term used by Sun Microsystems to denote a cluster server. Also known as a package (HP-UX), and a virtual server (Microsoft).
annotation	1. Comment associated with an archive save set. 2. Comment associated with an event.
application specific module (ASM)	Program that is used in a directive to specify how a set of files or directories is to be backed up or recovered. For example, compressasm is a NetWorker directive used to compress files.
archive	Process that backs up directories or files to an archive volume to free up disk space for regular backups. Archived data is not recyclable. See groom .
archive request	NetWorker resource used to schedule and manage archiving.
archive volume	Volume used to store archive data. Archive data cannot be stored on a backup volume or a clone volume.
attribute	Name or value property of a resource.
authentication	Process by which a user or software process is determined to be trusted or not trusted.
authorization	Privileges assigned to users.
authorization code	Unique code that in combination with an associated enabler code unlocks the software for permanent use on a specific host computer. See license key .
autochanger	See library .

auto media management Feature that enables the storage device controlled by the NetWorker server to automatically label, mount, and overwrite a volume it considers unlabeled.

B

backup

1. Duplicate of database or application data, or an entire computer system, stored separately from the original, which can be used to recover the original if it is lost or damaged.
2. Operation that saves data to a volume for use as a backup.

backup cycle Full or level 0 backup and all the subsequent incremental backups that are dependent on that backup.

Backup Operators group Microsoft Windows user group whose members have the capability to log in to a domain from a workstation or a server, whose data they may back up and restore. Backup Operators can also shut down servers or workstations.

backup volume A volume used to store backup data. NetWorker backup data cannot be stored on an archive volume or a clone volume.

bootstrap Save set that is essential for disaster recovery procedures. The bootstrap consists of three components that reside on the NetWorker server: the media database, the resource database, and a server index.

browse policy NetWorker policy that specifies the period of time during which backup entries are retained in the client file index. Backups listed in the index are browsable and readily accessible for recovery.

C

canned report Preconfigured report that can be tailored by the user.

carousel See [library](#).

client Host on a network, such as a computer, workstation, or application server whose data can be backed up and restored with the backup server software.

client file index Database maintained by the NetWorker server that tracks every database object, file, or file system backed up. The NetWorker server maintains a single index file for each client computer. The tracking information is purged from the index after the browse time of each backup expires.

client-initiated backup See [manual backup](#).

Client resource NetWorker server resource that identifies the save sets to be backed up on a client. The Client resource also specifies information about the backup, such as the schedule, browse policy, and retention policy for the save sets.

clone

1. Duplicate copy of backed-up data, which is indexed and tracked by the NetWorker server. Single save sets or entire volumes can be cloned.
2. Type of mirror that is specific to a storage array.

clone volume	Exact duplicate of a backup or archive volume. NetWorker software can index and track four types of volumes (backup, archive, backup clone, and archive clone). Save sets of these different types may not be intermixed on one volume. Clone volumes may be used in exactly the same way as the original backup or archive volume.
cluster	Group of linked virtual or physical hosts, each of which is identified as a node, with shared storage that work together and represent themselves as a single host.
common internet file system (CIFS)	Formerly known as Server Message Block (SMB). Message format used by Microsoft DOS and Windows to share files, directories, and devices.
connection port	Port used to perform functions through a firewall.
Console application administrator	Console server user role whose members can configure features, except security features, in the Console sever application.
Console security administrator	Console server user role whose members can add Console users and assign them to Console roles.
Console server	See NetWorker Management Console (NMC) .
consolidate	To create a full backup by merging a new level 1 backup with the last full level backup.
continued save set	Save set data that is continued from a previous volume.
control zone	Group of datazones managed by the NetWorker software.
conventional storage	Storage library attached to the NetWorker server or storage node, used to store backups or snapshot backups. Also known as secondary storage. See primary storage .

D

daemon	Process on UNIX systems that runs in the background and performs a specified operation at predefined times or in response to certain events.
database	1. Collection of data arranged for ease and speed of update, search, and retrieval by computer software. 2. Instance of a database management system (DBMS), which in a simple case might be a single file containing many records, each of which contains the same set of fields.
data management application (DMA)	Application that manages a backup or recovery session through an NDMP connection.
data mover (DM)	Client system or application, such as NetWorker software, that moves data during a backup, recovery, snapshot, or migration operation. See proxy host .
data server agent (DSA)	Functionality that enables the NetWorker server to communicate with a non-NetWorker NDMP host and package images of save streams. For example, an NDMP host that generates proprietary save data may send that data to a NetWorker storage device to have a save set associated with it.

data service provider (DSP)	Feature that controls access to disk storage during an NDMP back up.
datazone	Group of clients, storage devices, and storage nodes that are administered by a NetWorker server.
deduplication backup	Type of backup in which redundant data blocks are identified and only unique blocks of data are stored. When the deduplicated data is restored, the data is returned to its original native format.
destination client	Computer to which database files are restored in a directed recovery.
device	1. Storage folder or storage unit that can contain a backup volume. A device can be a tape device, optical drive, autochanger, or disk connected to the server or storage node. 2. General term that refers to storage hardware. 3. Access path to the physical drive, when dynamic drive sharing (DDS) is enabled.
Device Central	Interface from which one can manage all NetWorker libraries.
DFS component	1. A namespace for files and DFS links, called a DFS root. 2. A connection to a shared file or folder, called a DFS child node. See distributed File System (DFS).
direct access restore (DAR)	NDMP operation that can recover data in the middle of a tape set without having to parse the tape set sequentially, thereby reducing the recovery time of large backups.
directed recovery	Method that recovers data that originated on one client host and re-creates it on a different client host, known as the destination client.
directive	Instruction that directs NetWorker software to take special actions on a given set of files for a specified client during a backup or recovery operation. Directives are ignored in manual (unscheduled) backups.
disaster recovery	Restore and recovery of data and business operations in the event of hardware failure or software corruption.
distributed File System (DFS)	Microsoft Windows add-on that creates a logical directory of shared directories that span multiple hosts across a network.
document mode	Display mode that presents static reports such as charts or tables in a format that resembles the Print Preview mode in a PDF viewer.
drill-down	Organization of report information by granularity. For example, within a group summary report, a client report may be viewed, and then a report for a selected save set for that client.
drive	Hardware device through which media can be read or written to. See device .
DSA save set	Save sets of an NDMP client that are backed up to non-NDMP tape device. See data server agent (DSA) .

dynamic drive sharing (DDS) Feature that allows NetWorker software to recognize and use shared drives and when they are available.

E

enabler code Unique code that activates the software:

- Evaluation enablers or temporary enablers expire after a fixed period of time.
- Base enablers unlock the basic features for software.
- Add-on enablers unlock additional features or products, for example, library support.

See [license key](#).

enterprise Computers and folders organized into a tree-based visual representation.

event Notification generated by an application that could require user action, such as the impending expiration of a software enabler key that appears in the daemon log of the Console server.

event-based backup See [probe-based backup](#).

exit code Indicator that specifies whether a backup or recovery session succeeded. An exit code of zero (0) indicates the session completed successfully. A nonzero exit code indicates that the session did not complete successfully.

expiration date Date when a volume changes from read/write to read-only.

expired save set Save set that has exceeded its browse time and has been removed from the NetWorker client file index. Expired save sets can no longer be browsed.

F

file index See [client file index](#).

file system

1. Software interface used to save, retrieve, and manage files on storage media by providing directory structures, data transfer methods, and file association.
2. Entire set of all files.
3. Method of storing files.

firewall Security software designed to prevent unauthorized access to or from a private network.

folder An icon on a computer screen that can be used to access a directory.

full backup Type of backup that backs up all data objects or files, including the transaction logs contained in databases, regardless of when they last changed. See [level](#).

G

generic services toolkit (GST)	Software framework that underlies the Console server.
groom	Process that removes the original files from a local disk after a successful archive operation.
group	One or more client computers that are configured to perform a backup together, according to a single designated schedule or set of conditions.

H

hash	Number generated from a string of text that is used to encrypt a user password. See salted hash .
heterogeneous network	Network with systems of different platforms and operating systems that interact across the network.
high-availability system	System of multiple computers configured as cluster nodes on a network that ensures that the application services continue despite a hardware or software failure. Each cluster node has its own IP address with private resources or disks that are available only to that computer.
high-water mark	Percentage of disk space that, when filled, automatically starts the staging process.
host	Computer on a network.
host authentication	Encryption and verification services between NetWorker hosts. See user authentication .
host ID	Eight-character alphanumeric number that uniquely identifies a computer.
hostname	Name or address of a physical or virtual host computer that is connected to a network.

I

inactivity timeout	Time in minutes to wait before a client is considered to be unavailable for backup.
incremental backup	See level .
individual user authentication	Process by which Console administrators restrict or grant user access to NetWorker servers, based on Console usernames.
insertion time	Time that the save set record was most recently introduced into the save set database.
Interactive mode	Console mode that displays reports (as charts or tables) that users can interact with. For example, one can sort, rearrange, and resize columns in a table-format report that was run in this mode.
Internationalization (I18N)	Process of adapting software to accept input and output of data in various languages and locales.

J

JAR (Java Archive)	A file that contains compressed components needed for a Java applet or application.
Java	Type of high-level programming language that enables the same, unmodified Java program to run on most computer operating systems. See Java Virtual Machine (JVM) .
Java plug-in	JVM that can be used by a web browser to run Java applets.
Java Virtual Machine (JVM)	Execution environment for interpreting the Java programming language. Each operating system runs a unique JVM to interpret Java code.
jukebox	See library .

L

label	Electronic header on a volume used for identification by a backup application.
legacy method	Use of special-case Microsoft APIs to back up and recover operating system components, services, and applications.
level	Backup configuration option that specifies how much data is saved during a scheduled or manual backup: • A full backup backs up all data objects or files, regardless of when they last changed. • An incremental backup backs up only data objects or files that have changed since the previous backup.
library	Hardware device that contains one or more removable media drives, as well as slots for pieces of media, media access ports, and a robotic mechanism for moving pieces of media between these components. Libraries automate media loading and mounting functions during backup and recovery. The term library is synonymous with autochanger, autoloader, carousel, datawheel, jukebox, and near-line storage.
library sharing	Shared access of servers and storage nodes to the individual tape drives within a library. The drives are statically assigned to hosts.
license key	Combination of an enabler code and authorization code for a specific product release to permanently enable its use. Also called an activation key.
License Manager (LLM)	Application that provides centralized management of product licenses.
Lightweight Directory Access Protocol (LDAP)	Set of protocols for accessing information directories.
live backup	See rollover-only backup .
local cluster client	NetWorker client that is not bound to a physical machine, but is instead managed by a cluster manager. It is also referred to as a logical or virtual client.
localization (L10N)	Translation and adaptation of software for the user language, time formats, and other conventions of a specific locale.

logical cluster client See [virtual cluster client](#).

logical device Virtual device used in the integration of NetWorker software with SmartMedia. Many logical devices can be assigned to a single physical device.

low-water mark Percentage of disk space filled that, when reached, automatically stops the migration process.

LUS Driver used by EMC software products as a proprietary device driver that sends arbitrary SCSI commands to an autochanger. Also known as the EMC User SCSI.

M

managed application Program that can be monitored or administered, or both from the Console server.

managed node Storage management application under the control of Console. For example, a system running NetWorker on a backup server or storage node is considered to be a managed node.

man pages Online technical reference manual, normally provided on UNIX servers, for the syntax and function of program commands that may be issued from the command line.

manual backup Backup that a user performs from the client, also known as an unscheduled, on-demand, or ad hoc backup.

media Physical storage, such as a disk file system or magnetic tape, to which backup data is written. See [volume](#).

media index Database that contains indexed entries of storage volume location and the life cycle status of all data and volumes managed by the NetWorker server. Also known as media database.

member Physical host that occupies a node in a cluster environment. Each member has its own IP address.

mount To make a volume physically available for use, such as the placement of a removable disk volume or tape into a drive for reading or writing.

mount host Host in a network that is used to mount storage array snapshot volumes to perform snapshot restore and rollover operations.

mount point See [volume mount point](#).

multiple session See [parallelism](#).

multiplex To simultaneously write data from more than one save set to the same storage device.

N

NDMP server Instance of one or more NDMP services, such as a data, tape, or SCSI server, that is managed by a single control connection.

NDMP service	Virtual machine that is controlled by a data management application (DMA) such as NetWorker software. Example services include: • Server with a directly attached storage appliance • Storage device system with one or more tape drives • Software process that reads two datastreams and multiplexes them into one stream
NDMP storage node	Host or open system with NDMP services. For example, Netapp Filer and EMC Filer.
near-line storage	See library .
network attached storage (NAS)	Disk array or storage device (NAS filer) that connects directly to the messaging network or LAN interfaces and uses the common communication protocols of TCP/IP or NDMP.
Network Data Management Protocol (NDMP)	Software component that uses TCP/IP standards to specify how heterogeneous network components communicate for the purposes of backup, recovery, and transfer of data between storage systems.
NetWorker administrator	NetWorker server user who may add, change, or delete NetWorker server users.
NetWorker application administrator	NetWorker server user who may operate NetWorker software, configure the NetWorker server, and create and modify NetWorker resources.
NetWorker Management Console (NMC)	Software program that is used to manage NetWorker servers and clients. The NMC server also provides reporting and monitoring capabilities for all NetWorker processes.
NetWorker security administrator	NetWorker server user who may add, change, or delete NetWorker server user groups.
NetWorker server	Computer on a network that runs the NetWorker server software, contains the online indexes, and provides backup and restore services to the clients and storage nodes on the same network.
NetWorker Snapshot Management (NSM)	Technology that provides point-in-time snapshot copies of data. NetWorker software backs up data from the snapshot. This allows applications to continue to write data during the backup operation, and ensures that open files are not omitted.
network file system (NFS)	Communications protocol that enables users to access shared files on different types of computers over a network.
NFS server	Host that contains exported file systems that NFS clients can access. See network file system (NFS) .
node	See cluster .
non-critical volume	A volume that contains files that are not part of the system state or an installed service.
notification	Message sent to the NetWorker administrator about important NetWorker events.
nsrd	Master NetWorker server process.
nsrhost	Logical hostname of the NetWorker server.

O

offline backup	Backup of database objects performed while the corresponding database or instance is shut down and unavailable to users. Also known as a cold backup.
offline restore	Automated restore that does not require the manual installation of an operating system. A bare metal recovery (BMR) is an offline restore.
online backup	Backup of database objects performed while the corresponding database or instance is running and available to users. Also known as a hot backup.
online indexes	Databases located on the NetWorker server that contain all the information pertaining to the client backups (client file index) and backup volumes (media index).
online restore	Restore operation that is performed from a NetWorker recover program. An online restore requires that the computer has been booted from an installed operating system. See also offline restore.
operator	Person who performs day-to-day data storage tasks such as loading backup volumes into storage devices, monitoring volume locations and server status, verifying backups, and labeling volumes.
override	Different backup level that is used in place of the regularly scheduled backup.

P

package	A term used by HP-UX to denote a cluster server. Also known as an agent (Sun) or virtual server (Microsoft).
parallelism	Feature that enables a maximum number of concurrent streams of data during backup or restore operations. For example, parallelism values can be set for the NetWorker server, clients, pools, and groups.
pathname	Set of instructions to the operating system for accessing a file: • An absolute pathname indicates how to find a file by starting from the root directory and working down the directory tree. • A relative pathname indicates how to find a file by starting from the current location.
peer	NetWorker host that is involved in an authentication process with another NetWorker host.
permanent enabler	Enabler code that has been made permanent by the application of an authorization code. See enabler code .
physical cluster client	Backup client that is bound to a physical host in the cluster and can have its own resources (private or local).
physical host	Node or host that forms part of a cluster.
point-in-time copy (PIT copy)	Fully usable copy of a defined collection of data, such as a consistent file system, database, or volume that contains an image of the data as it appeared at a specific point in time. A PIT copy is also called a snapshot or shadow copy.

policy	Set of defined rules for client backups that can be applied to multiple groups. Groups have dataset, schedule, browse, and retention policies.
pool	1. NetWorker sorting feature that assigns specific backup data to be stored on specified media volumes. 2. Collection of NetWorker backup volumes to which specific data has been backed up.
primary storage	Server storage subsystem, such as a disk array, that contains application data and any persistent snapshots of data.
probe-based backup	Type of scheduled backup, also known as an event-based backup, where the NetWorker server initiates the backup only when specified conditions are met, as determined by one or more probe settings.
proxy host	Surrogate host computer that performs backup or clone operations in place the production host by using a snapshot copy of the production data. See mount host .
purge	Operation that deletes file entries from the client file index.
Q	
quiesce	State in which all writes to a disk are stopped and the file system cache is flushed. Quiescing the database prior to creating the snapshot provides a transactionally consistent image that can be remounted.
R	
recover	To restore data files from backup storage to a client and apply transaction (redo) logs to the data to make it consistent with a given point-in-time.
recyclable save set	Save set whose browse and retention policies have expired. Recyclable save sets are removed from the media database.
recyclable volume	Storage volume whose data has exceeded both its browse and retention policies and is now available to be relabeled and reused.
Registry	Microsoft Windows database that centralizes all Windows settings and provides security and control of system, security, and user account settings.
remote device	1. Storage device that is attached to a storage node that is separate from the NetWorker server. 2. Storage device at an offsite location that stores a copy of data from a primary storage device for disaster recovery.
remote procedure call (RPC)	Protocol used by the backup server to perform client requests over a network.
repository	Console database that contains configuration and reporting information.
requester	A VSS-aware application that creates and destroys a shadow copy. NetWorker software is a requester. See shadow copy .

resource	Software component whose configurable attributes define the operational properties of the NetWorker server or its clients. Clients, devices, schedules, groups, and policies are all NetWorker resources.
resource database	NetWorker database of information about each configured resource.
resource owner	Logical cluster host that owns the resource. If a Cluster resource, such as a shared disk, is not owned by a virtual host, it is assumed to be owned by the physical node that hosts the resource.
restore	To retrieve individual data files from backup media and copy the files to a client without applying transaction logs.
retention policy	NetWorker setting that determines the minimum period of time that backup data is retained on a storage volume and available for recovery. After this time is exceeded, the data is eligible to be overwritten.
retrieve	To locate and recover archived files and directories.
retry mechanism	Action that NetWorker software performs when client operations fail. This situation might occur because the rate of transmission is either low or undetectable.
role	Grant of user privileges to the Console. There are three roles: Console Application Administrator, Console Security administrator, and the Console User. See user groups.
roll forward	To apply transactional logs to a recovered database to restore it to a state that is consistent with a given point-in-time.
rollover	Backup of a snapshot to conventional storage media, such as disk or tape. Previously known as a live backup.
rollover-only backup	Rollover whereupon the snapshot copy is deleted. Previously known as a serverless backup, live backup, or nonpersistent backup.
root	1. (UNIX only) UNIX superuser account. 2. (Microsoft Windows and UNIX) Highest level of the system directory structure.
S	
salted hash	Added string of random data that provides a unique identifier to a user's password. See hash.
save	NetWorker command that backs up client files to backup media volumes and makes data entries in the online index.
save set	1. Group of files or a file system copied to storage media by a backup or snapshot rollover operation. 2. NetWorker media database record for a specific backup or rollover.
save set consolidation	Process that performs a level 1 backup and merges it with the last full backup of a save set to create a new full backup.
save set ID (ssid)	Internal identification number assigned to a save set.

save set recover	To recover data by specifying save sets rather than by browsing and selecting files or directories.
save set status	NetWorker attribute that indicates whether a save set is browsable, recoverable, or recyclable. The save set status also indicates whether the save set was successfully backed up.
save stream	Data and save set information that is written to a storage volume during a backup. A save stream originates from a single save set.
scanner	NetWorker command used to read a backup volume when the online indexes are not available.
scheduled backup	Type of backup that is configured to start automatically at a specified time for a group of one or more NetWorker clients. A scheduled backup generates a bootstrap save set.
secondary storage	Storage media managed by a NetWorker server or storage node that stores conventional or snapshot data. Configure a storage device on a NetWorker server or storage node for each secondary storage.
security event	Operation related to authorization, authentication, or configuration.
service port	Port used to listen for backup and recover requests from clients through a firewall.
shadow copy	Temporary, point-in-time copy of a volume created using VSS technology. See VSS (Volume Shadow Copy Service) .
shared disk	Storage disk that is connected to multiple nodes in a cluster.
shell prompt	Cursor in a shell window where commands are typed.
silo	Repository for holding hundreds or thousands of volumes. Silo volumes are identified by bar codes, not by slot numbers.
simple network management protocol (SNMP)	Protocol used to send messages to the administrator about NetWorker events.
skip	Backup level in which designated files are not backed up. See level .
Smart Media	EMC software application that manages media resources within a distributed environment.
snapset	See snapshot save set .
snapshot	Point-in-time, read-only copy of specific data files, volumes, or file systems on an application host. Operations on the application host are momentarily suspended while the snapshot is created on a proxy host. Also called a PiT copy, image, or shadow copy.
snapshot policy	Sets of rules that control the life cycle of snapshots. These rule specify the frequency of snapshot creation, how long snapshots are retained, and which snapshots will be backed up to conventional storage media.
snapshot save set	Group of files or other data included in a single snapshot. Previously called a snapset.

stage	To move data from one storage medium to a less costly medium, and later removing the data from its original location.
stand-alone	In a cluster environment, a NetWorker server that starts in noncluster (stand-alone) mode.
stand-alone device	Storage device that contains a single drive for backing up data. Stand-alone devices cannot automatically load backup volumes.
STL	Silo Tape Library.
storage node	Computer that manages physically attached storage devices or libraries, whose backup operations are administered from the controlling NetWorker server. Typically a “remote” storage node that resides on a host other than the NetWorker server.
synthetic full backup	Backup that combines a full backup and its subsequent incremental backups to form a new full backup. Synthetic full backups are treated the same as ordinary full backups.
T	
tape service	NDMP DSP service that controls access to tape storage. A system can simultaneously host multiple tape services corresponding to multiple backup streams.
target client	NetWorker client on which data is to be restored This may be the same as the original source client from which the data was backed up, or it may be a different client.
target database	Database that the NetWorker server backs up as a safeguard against data loss.
target sessions	The number of simultaneous backup data streams accepted by a backup device.
temporary enabler	Code that enables operation of the software for an additional period of time beyond the evaluation period. See enabler code .
transaction log	Record of named database transactions or list of changed files in a database, stored in a log file to execute quick restore and rollback transactions.
transmission control protocol / internet protocol (TCP/IP)	Standard set of communication protocols that connects hosts on the Internet.
trap	Setting in an SNMP event management system to report errors or status messages.
U	
update enabler	Code that updates software from a previous release. It expires after a fixed period of time.
user	1. A NetWorker user who can back up and recover files from a computer. 2. A Console user who has standard access privileges to the Console server.
user alias	Username seen by the NetWorker server when a Console user connects to the NetWorker server.

user authentication Feature that validates user sign-on attempts. NetWorker can validate sign-on attempts against either a central authority, such as an LDAP database, or a local Console database. **See** [host authentication](#).

user data Data that is generated by users, typically for the purposes of a business function. A Microsoft Word document or an Excel spreadsheet is an example of user data.

user groups Feature that assigns user privileges. **See** [role](#).

V

versions Date-stamped collection of available backups for any single file.

virtual cluster client NetWorker client that is not permanently bound to one physical host but is managed by a cluster manager. It is also referred to as a logical cluster client or a virtual client.

virtual server

1. Server, usually a web server, that shares resources with other virtual servers on the same computer to provide low-cost hosting services.
2. In a cluster configuration, a set of two nodes, which are physical computers, and virtual servers. Each node and virtual server has its own IP address and network name. Each virtual server also owns a subset of shared cluster disks and is responsible for starting cluster applications that can fail over from one cluster node to another.

virtual tape library (VTL) Software emulation of a physical tape library storage system.

volume

1. Unit of physical storage medium, such as a disk or magnetic tape, to which backup data is written.
2. Identifiable unit of data storage that may reside on one or more computer disks.

volume ID (valid) Internal identification that NetWorker software assigns to a backup volume.

volume mount point Disk volume that is added into the namespace of a host disk volume. This allows multiple disk volumes to be linked into a single directory tree, and a single disk or partition to be linked to more than one directory tree.

volume name Name that you assign to a backup volume when it is labeled.

VSS (Volume Shadow Copy Service) Microsoft technology that creates a point-in-time snapshot of a disk volume. NetWorker software backs up data from the snapshot. This allows applications to continue to write data during the backup operation, and ensures that open files are not omitted.

VSS component A subordinate unit of a writer. **See** [writer](#).

W

Windows disaster recovery Bare metal recovery of a host. NetWorker provides an automated bare metal recovery solution for Windows.

writer Database, system service, or application code that works with VSS to provide metadata about what to back up and how to handle VSS components and applications during backup and restore. **See** [VSS \(Volume Shadow Copy Service\)](#).