



EMC[®] NetWorker

Release 8.1 Service Pack 1

Administration Guide

P/N 302-000-421

REV 04

EMC²

Copyright © 1990 - 2013 EMC Corporation. All rights reserved. Published in the USA.

Published May 26, 2014

EMC 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. EMC Corporation 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 EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to the technical documentation and advisories section on the EMC online support website.

CONTENTS

Preface

Revision History

Chapter 1

Overview

The NetWorker environment	32
NetWorker components.....	32
NetWorker features	33
Performance features	33
Ease of use	34
Scalability.....	35
Optional software additions	35
NetWorker user interfaces	35
NetWorker Management Console interface	35
NetWorker client interface	44
NetWorker character-based interface	47
NetWorker command-line interface	47
Common NetWorker tasks	47
How to add a new host.....	47
How to configure devices	48
How to label media	49
How to schedule a backup	50
How to view failed backups.....	50
How to perform a manual backup.....	50
NetWorker services	51
Services and programs on the NetWorker server.....	52
Services and programs on the NetWorker client.....	52
Services and programs on the NetWorker storage node.....	53
Services and programs on the NetWorker Management Console server ..	53
Stopping and starting the Console server	54
Stopping and starting a NetWorker server, client, or storage node	55

Chapter 2

Backing Up Data

Scheduled backups.....	58
Planning a scheduled backup	58
Setting up a scheduled backup	59
Save sets	66
Scheduling predefined save sets for backup	66
The ALL save set.....	68
Manual backups	70
Performing a manual backup on Windows.....	71
Performing a manual backup from the command prompt	73
Verifying backed-up data.....	76
Synthetic full backups.....	76
Differences between a synthetic full backup and a traditional backup ...	77
When to use synthetic full backups	78
How a synthetic full backup is created	78

Synthetic full requirements	79
Limitations of synthetic full backups	80
Recommended devices for synthetic full backups	81
Synthetic full scheduling considerations	81
Synthetic full and backup levels	82
Performing synthetic full backups	83
Checking the integrity of a synthetic full backup	86
Checkpoint restart considerations with synthetic full backups	87
Reporting and synthetic full backups	87
Running queries on synthetic full backups	87
Monitoring	87
Directives	87
Directing data from a synthetic full backup to a dedicated pool	88
Virtual synthetic full backups (for Data Domain systems)	88
VSF requirements	89
Enable/disable VSF in NMC	90
Performing VSF	90
Directives	91
Multiple storage node distribution	91
Concurrent operations	91
Avamar rehydrated save sets	91
Validating VSF	91
VSF limitations	93
Enable parallel save streams	93
Troubleshooting PSS	94
Probe-based backups	94
Client Direct backups	95
Checkpoint restart backups	95
Checkpoint restart usage	96
About partial non-NDMP save sets	97
About partial NDMP save sets	98
Configuring checkpoint enabled clients	98
Restarting a checkpoint-enabled backup	100
Monitoring checkpoint-enabled backups	101
Reporting checkpoint-enabled backups	102
Recovering checkpoint restart data	106
Cloning and scanning partial savesets	107
Cloud backup devices and partial savesets	107
Deduplication backups	107
Encrypting backup data	108
Microsoft Encrypting File System (EFS) and NetWorker AES encryption	108
Set the Datazone pass phrase for a NetWorker server	108
Apply AES data encryption to clients in the datazone	108
Compressing backup data	109
Applying compression to a scheduled backup	109
Applying compression to a manual backup	109
Special data handling for NetWorker clients on Windows	109
Backing up Console server management data	110
Scheduling backups for the Console server database	110
Performing a manual backup of the Console database	111
Managing the size of the Console database log file	111
Backing up Windows mount points	112
Including mount points in scheduled backups	112
Performing a manual backup of a mount point and its data	113
Performing a manual backup of nested mount points and their data ...	113

Backing up the Windows Content Index Server	114
Backing up CIS on Windows	114
Directing NetWorker software to skip CIS catalog backups.....	115
Backing up Windows DHCP and WINS databases.....	115
Backing up a DHCP database	115
Backing up a WINS database.....	115
Windows backup and recovery notes.....	116
Enabling short filename support	116
Enabling hard link support	116
Failed backup and recovery attempts	117
Granting full permissions for backup of Disk Quota database.....	117
Native VHD volume support.....	117
Recovery and case-sensitivity.....	118
Security settings for logging operations performed by backup operator.....	118
Customizing the backup command.....	119
Using the save command with a customized backup script	119
Using the savepnpc command with a customized backup program	124
Considerations for backing up raw partitions.....	128
Backing up raw partitions on UNIX.....	128
Backing up raw partitions on Windows.....	128
Backing up a mapped drive.....	128
Backing up access control lists.....	129
Backing up BOOT/BCD Data on Windows	129
Support for backing up renamed directories.....	129
Backing up only client file indexes and the bootstrap	130

Chapter 3

Storage Nodes and Libraries

Storage nodes.....	132
Requirements.....	132
Licensing	132
Configuring storage nodes	133
Configure a storage node	133
Modifying the timeout attribute for storage node operations	135
Configure timeouts for storage node remote devices	135
Configure the client's storage node affinity list.....	136
Storage node load balancing.....	137
Bootstrap backup on a storage node	139
Dedicated storage nodes	139
Troubleshooting storage nodes	140
Avamar deduplication nodes and replication nodes	140
Devices and libraries.....	140
SCSI libraries	141
NDMP libraries	141
Silo libraries.....	141
Autodetection of libraries and tape devices.....	141
Scanning for libraries and devices.....	142
Refreshing enterprise library views on request	143
Changing the polling interval for enterprise library views.....	143
Configuring libraries.....	144
Adding a library resource.....	144
Configuring a virtual tape library (VTL)	145
Queuing device resources for AlphaStor.....	145
Library reconfiguration	146

Reconfiguring a library	146
Using the jbedit command to configure a library.....	147
Specifying available library slots	147
Specify library slots.....	147
Miscellaneous library operations	148
Sharing libraries among NetWorker hosts.....	148
Sleeping periods for library tasks	149
Server Network Interface attribute	150
Tips for using libraries.....	150
Library notifications	150
Resetting a library	151
Using pools with libraries.....	152
Adding and removing media by using the library front panel	152
Library maintenance	152
Automatic tape device cleaning.....	153
Selecting a tape device manually for cleaning	153
Delaying tape device cleaning	153
Tape alert.....	154
Deleting libraries.....	155
Troubleshooting autoconfiguration failure	155
Silo libraries	156
NetWorker software interactions with a silo.....	156
Installing a silo.....	157
Naming conventions for silo devices	157
Configuring silo libraries	158
Releasing a silo device.....	159
Cleaning silo devices	160
Environment variables for StorageTek silos	160

Chapter 4 Backup to Disk and Cloud

Types of disk storage devices	162
FTD.....	162
AFTD	162
DD Boost devices	162
Cloud	163
Example environment	163
Differences between FTDs, AFTDs and DD Boost devices.....	164
File type devices	166
Advanced file type devices.....	167
Memory requirements for AFTD backups.....	167
Create and configure an AFTD.....	167
Labeling and mounting an AFTD	179
Providing sufficient disk space for an AFTD.....	179
Verifying AFTD operations.....	182
Change the AFTD block size.....	182
Recover savesets by using AFTD concurrent operations	183
Limitations with concurrent AFTD recovery operations	183
Deactivating and erasing an AFTD.....	183
DD Boost devices	184
Cloud devices	185
Cloud backup devices compared to other device types.....	185
Cloud backup requirements and considerations.....	185
Cloud best practices	186
Creating and labeling a cloud storage device.....	186

Gathering report information on cloud backup	189
Staging with a cloud storage device	189
Cloning to a cloud storage device	189

Chapter 5

Backup to Tape and VTL

Overview of tape device storage	192
Stand-alone devices	192
Autodetecting and configuring a stand-alone tape drive	192
Adding a stand-alone device manually	193
SCSI data block size issues between UNIX and Windows	194
Determine the allowable block size	194
Set the block size for a device type	194
Device block size for read and write operations	195
Block-size mode (UNIX/Linux only)	195
Device parameter settings	195
Device settings in the NetWorker Administration interface	195
Device settings as environment variables	196
Common device interface	199
Device ordering	199
Persistent binding and naming	200
Detecting device ordering issues	201
Correcting drive ordering changes	201
Reordering tape drive numbers (Microsoft Windows only)	202
Device calibration	203
Dynamic drive sharing	203
Introduction to DDS	203
DDS block-size compatibility between UNIX and Windows	204
Preventing unintended access to DDS devices	205
Enabling DDS with NDMP	207
DDS attributes in the device properties	208
High availability and DDS	209
Improvements to deduplication rates for Data Domain VTL multiplexed backups	210
Nonrewinding tape device usage (UNIX/Linux only)	212
Support for LTO-4 hardware-based encryption	213
Recycling compared to adding more volumes	213
Display device operations messages	213
Service mode	214

Chapter 6

Media Management

Storage management operations	216
How the NetWorker server uses volume labels	216
How the NetWorker server selects a volume	216
Auto Media Management	218
Using Auto Media Management	219
Volume operations	220
Using the read-only mode	222
Changing a volume's mode	222
Recycling volumes	222
Labeling volumes	223
Using barcode labels	225
Mounting and unmounting volumes	227
Using libraries with a volume import and export capability	231

Inventorying library volumes	233
Working with volumes	234
Media handling errors	236
Re-enable a device	236
Media management in a silo	237
Numbering a silo slot	237
Mounting and unmounting silo volumes.....	237
Labeling a silo volume	238
Using a silo with volume import/export capability	238
Barcode IDs	239
Inventorying a silo.....	242
Volume save sets	242
Viewing save set details in the Volume Save Sets window	242
Viewing save set details from the Save Set detail table	244

Chapter 7

Backup Groups and Schedules

Overview of NetWorker scheduling	248
Backup groups	248
The NetWorker server and time-based backup groups	249
Preconfigured groups	250
Key Group attributes	251
Probe Group	252
Configuring a probe-enabled group	252
Aborted backup groups	253
How to create a group	253
How to edit a group	253
How to delete a group	254
How to copy a group.....	254
Setting the backup group time interval.....	255
Limiting full backups when the time interval is less than 24 hours	255
Running a backup group from the command line or a script.....	256
Managing backup groups	257
Previewing a backup group	257
Moving clients between groups	258
Estimating save set sizes of a backup group.....	258
Backing up status reports.....	258
Generating and printing bootstrap reports.....	259
Backing up open files.....	259
Opening files owned by the operating system	259
Opening files owned by a specific application	260
Files that change during backup.....	260
Backing up open files with VSS	260
Schedules.....	260
Schedules for Avamar deduplication clients	261
Preconfigured NetWorker schedules	261
Backup cycles	262
Scheduling and planning considerations.....	264
Scheduling large client file systems	265
Key components of a schedule	265
Working with schedules	266
Overriding a client's regular backup schedule	267
Backup levels.....	267
How NetWorker backup levels work.....	268
The NetWorker server and backup levels	270

	Backup levels for Windows SYSTEM and VSS SYSTEM save sets	272
Chapter 8	Browse and Retention Policies	
	About browse and retention policies	276
	Browse policies.....	276
	Retention policies	278
	Managing the data lifecycle.....	281
	Assigning multiple policies to a single client	282
	Preconfigured time policies	283
	Editing a time policy.....	283
	Delete a time policy.....	284
	Snapshot policies	284
	Working with snapshot policies.....	284
	Browse and retention policies for manual backups.....	286
	Modifying the browse and retention policy on a save set.....	286
	Reports on browse and retention policies for save sets	287
Chapter 9	Directives	
	Directives overview	290
	Types of local and global directives.....	290
	Global directives	290
	NetWorker User local directives (Windows only)	290
	Local directive files	290
	Creating a global directive resource.....	290
	Editing a global directive resource.....	291
	Deleting a global directive resource.....	291
	Copying a global directive resource	291
	Example directives	292
	Order of precedence of global and local directives	293
	Local directives within the NetWorker User program	293
	Set up a NetWorker User program local directive	293
	Preconfigured global directive resources	294
	Format of directive statements	297
	Directory specifications.....	297
	ASM specifications.....	298
	Save environment keywords.....	298
	Application Specific Modules (ASMs)	299
	Precautions when using rawasm to back up UNIX raw partitions.....	302
	File matching with multiple ASMs in a directive	302
Chapter 10	Sorting Backup Data	
	Media pools.....	304
	Using media pools	304
	NetWorker media pool types	304
	Sorting data with media pools.....	305
	Directing client file indexes and bootstrap to a separate media pool ...	306
	Directing consolidated backup data to a specific media pool	307
	Meeting the criteria for more than one media pool configuration	307
	When no customized media pool criteria is met	308
	Configuring media pools	308
	Using storage devices and media pool configuration to sort data	310
	Creating a media pool	311

Managing volumes in a media pool	313
Supporting WORM and DLTWORM tape drives	314
Working with media pools	317
Label templates	318
Using label templates	318
Using preconfigured label templates	319
Completing Label Template attributes	320
Naming label templates	321
Working with label templates	322

Chapter 11 Archiving

Overview of archiving	326
Archive requirements	326
How the NetWorker server archives data	326
Indexed and nonindexed archiving	327
Permissions for archiving	328
Enabling archive services for the client	328
Enabling or restricting archive access	328
Enabling public archive access	328
About archive pools	329
Preconfigured Indexed Archive pool and PC Archive pool	329
Preconfigured archive pool	329
Creating custom Archive pools	329
Archiving data procedures	330
Enabling archive services for a NetWorker client	330
Manually archiving data	330
Scheduling data archives	331
Retrieving archived data	334
Retrieval permissions	334
Retrieving archives from a client on UNIX	334
Retrieving nonindexed archives from a client on Windows	335
Recovering indexed archive data from a client on Windows	336
Archive request management	337
Starting a scheduled archive at any time	337
Stopping a scheduled archive while in progress	337
Disabling a scheduled archive	337
Viewing details of a scheduled archive	337

Chapter 12 Cloning

Overview of cloning	340
Cloning requirements	340
Save set cloning	341
Considerations for scheduled clone jobs	341
Setting up a schedule clone job	342
Starting a scheduled clone job manually	344
Monitoring scheduled clone jobs	345
Setting up automatic cloning from a backup group	345
Viewing the clone status of a save set	346
Cloning a save set manually	346
Additional manual clone operations	349
Specifying browse and retention policies for clone data	349
Specify a browse and retention policy in a scheduled clone job	349
Specify a browse and retention policy from the command prompt	350

Specify a retention policy for a Clone pool	350
Volume cloning	350
Creating a clone volume	350
Viewing clone volume details	350
Recovering cloned data	351
Recovering a clone save set from the command prompt	352
Recovering a save set when all cloned instances have expired	353
Cloning archived data	353
Directing clones to a special storage node.....	354
Storage node selection criteria for reading the clone data	354
Storage node selection criteria and settings for writing a clone.....	355
Storage node selection criteria for recovering cloned data.....	356
Using file type devices for clone operations.....	357
Differences in the cloning process.....	357
Manual cloning with advanced file type device.....	357
Backup-to-tape for Avamar deduplication clients	358
Cloning with Data Domain devices	358
Using the nsrclone command	358

Chapter 13 Staging Backups

Save set staging.....	362
Working with staging policies.....	362
Creating a staging policy	362
Editing a staging policy	364
Copying a staging resource	364
Deleting a staging policy	365
Consideration for staging a bootstrap backup	365
Staging and cloning from the command prompt	365

Chapter 14 Recovering Filesystem Data

NetWorker recovery overview	368
Local recoveries	368
Directed recoveries	369
Overview of NetWorker recovery methods.....	372
Browsable recovery	372
Save set recovery	372
Scanner recovery.....	373
VSS File Level Recovery (FLR)	373
Recovering the data	373
Using the Recovery Wizard	374
Using the recover command	379
Using the NetWorker User program.....	382
Save set recover by using NetWorker User	385
Using the scanner program	386
Using VSS file level recovery (FLR)	388
Recovering deduplication data	389
Recovering with BMR.....	389
Recovering ACL files	389
Recovering encrypted data	389
Recovering the Windows system configuration	390
Temporary disk space	391
Disable Antivirus For Windows System Drive Recovery.....	392
Recovering the Windows SYSTEM from the command prompt.....	392

Point-in-time recovery of the SYSTEM and VSS SYSTEM save sets	394
Point-in-time recovery of Microsoft SQL Server or Exchange Server	395
Preparing to recover the Windows SYSTEM STATE save set	395
Preparing to recover the SYSTEM DB save set	396
Recovering Windows volume mount points	396
Recovering mount points.....	396
Recovering a mount point and its data	397
Recovering nested mount points	397
Recovering special Windows databases	397
Recovering Windows DHCP and WINS databases.....	397
Restoring Windows Content Index Server on Windows.....	398
Recovering expired save sets	398
Recovering client files from an old NetWorker server	404
Recovering critical NetWorker server databases.....	405
Prerequisites to recover the NetWorker server databases	406
Consider your recovery options	408
Recovering the NetWorker server databases.....	408
Options for running the nsrdr command	414
Setting nsrdr tuning parameters	416
Recovering the NMC server database	417

Chapter 15 Enterprise reporting and events monitoring

Enterprise data reporting	420
Enabling/Disabling the gathering of report data	420
Data retention and expiration policies.....	421
Setting expiration policies for data retention	422
Report categories	422
Report types.....	424
Configuring reports	425
Viewing reports	428
Preconfigured reports.....	435
Customizing and saving reports	451
Sharing reports	452
Exporting reports.....	453
Command line reporting.....	454
Printing reports	455
Enterprise events monitoring	456
Events.....	456
Polling for System Events	456
Enabling or disabling the Capture Events option.....	457
Viewing events.....	457
Event priorities.....	458
Working with notes	458
Working with annotations	459

Chapter 16 NetWorker server events reporting and monitoring

Monitoring NetWorker server activities	462
Groups window	465
Clones window.....	467
Sessions window	468
Alerts window	469
Devices window	470
Operations window	470

Log window.....	472
Archive Requests window.....	473
Recover window	475
Notifications	479
Preconfigured notifications	480
Customizing notifications.....	484
Logging event notifications	490
Creating a custom notification.....	490
Editing a notification	491
Copying a notification	491
Deleting a custom notification.....	491
Owner notifications	492
Reporting group status and backup job status.....	493
Savegroup completion and failure notifications.....	493
Querying the job status	497
Reporting recover job status.....	500
Using nsrrecomp	501
Chapter 17	NMC Server Management
NMC server authentication.....	504
Native NMC-based authentication	505
An external authentication authority	506
Troubleshooting authentication errors.....	521
Troubleshooting login errors	526
Restricting a user's view of managed servers.....	526
Resetting the administrator password (native NMC authentication only)	528
Moving the NMC server	528
Setting system options	531
Setting a system option.....	531
Individual User Authentication	532
Setting environment variables.....	533
Setting environment variables on UNIX.....	534
Setting environment variables on Windows systems	535
Accessing the Console Configuration Wizard	535
NetWorker NMC server maintenance tasks	535
Changing the service port used by the NetWorker Console database ...	535
Changing database connection credentials	536
NMC server IP address/hostname updates.....	537
Displaying international fonts in non-US locale environments	538
NetWorker License Manager.....	538
Entering an enabler code.....	538
Deleting an enabler code	538
Entering an authorization code.....	538
Changing the License Manager server	539
Chapter 18	NetWorker Server Management
Enterprise	542
Enterprise components	542
Organizing NetWorker servers	542
Viewing the enterprise	543
Managing various servers in the enterprise	544
Managing folders in the enterprise	546

- Adding or deleting multiple servers by using a hostname file 548
- Configuring a NetWorker server 550
 - Set up the server..... 550
- Report home 551
 - Enabling the report home feature 551
 - Manually running a report home report 551
 - Disabling the report home feature 552
 - Specifying additional email recipients..... 552
 - Specifying the sender email address 553
- Parallelism and multiplexing 553
 - Parallelism 553
 - Multiplexing..... 556
- Managing server access 558
 - The administrator list 558
 - NetWorker User Groups 559
 - Restrict backup and recover access to the NetWorker server..... 568
- Working with the Multi-Tenancy Facility 569
 - Users within a restricted data zone..... 570
 - Configurations for the Multi-Tenancy feature 573
 - Configuring a restricted data zone 574
 - Restricted Data Zone resource associations 583
 - Viewing data within a Restricted Data Zone 584
- Server communication issues within Microsoft Windows 585
 - Name resolution..... 585
 - Backup Operators group 585
 - Dynamic Host Configuration Protocol 585
 - Backup and Recover Server service 585
- Indexes 585
 - Characteristics of the online indexes 586
 - Automated index activities 586
 - Checking online indexes 587
 - Viewing information about the indexes 587
 - Index save sets 587
 - Querying the media database 588
 - Cross-checking client file indexes 589
 - Refreshing index information 589
 - Client file index locations 589
 - Managing the size of the online indexes..... 591
- Managing Client Push 594
 - Changing the location of the software repository 594
 - Removing software package information from the software repository. 595
 - Transferring files and folders by using nsrpush..... 597
- Monitoring Changes to NetWorker Server Resources..... 599
 - How to disable/enable the Monitor RAP Attribute 599
- Log file size management..... 600
- Internationalization..... 602
 - Log file viewer 602
 - Interoperability with previous releases of NetWorker 602
 - Display issues..... 602
 - Maximum path and save set length..... 603
 - Locale-specific configuration issues on UNIX/Linux..... 603

Chapter 19

NetWorker Client Management

- NetWorker client overview..... 606

Client configuration.....	606
Creating a client	606
Editing a client	606
Copying a client	606
Changing a client name	607
Deleting a client	607
Recovering a deleted client	608
Editing a client NSRLA database	608
Creating a client probe	610
Associating a probe with a Client resource	611
Creating a lockbox to store and retrieve pass phrases securely.....	611
Set the Datazone pass phrase for a NetWorker server.....	611
Set the Remote access attribute for NetWorker client resource when a lockbox is created for a NetWorker server on a cluster	611
Error messages and error handling for the Datazone pass phrase	612
NetWorker authentication	612
Strong authentication (nsrauth)	612
Authentication for backwards compatibility (oldauth)	613
Access privileges for authentication configuration.....	613
Specifying the minimum authentication strength between hosts	614
Maintaining NetWorker local host authentication credentials	616
Maintaining local host Peer resources	619
Creating a custom certificate and private key for a host	621
Multiple clients from the same computer	622
Redefining a file system into multiple client and save set instances	622
Defining a client and save set combination	623
Scheduled backups of non-ASCII files or directories	623
Controlling access to a NetWorker client.....	624
Editing the servers file.....	625
Client priority	625
Dedicated client/server interface for backup and recover operations.....	626

Chapter 20

Block Based Backup and Recovery

Overview of NetWorker block based backup and recovery.....	628
Supported operating systems and configurations for block based backups and recoveries	629
Limitations of block based backups	630
Preparing for block based backups.....	630
Creating a backup device	630
Configuring block based backups	632
[Optional] Creating a CIFS share for block based recoveries.....	632
Performing block based backups.....	633
Scheduled backups	633
Incremental backups.....	634
Virtual full backups	634
Synthetic full and incremental synthetic full backups	634
Manual backups or client-initiated backups	635
Save set backups	635
Exclude list backups	635
Windows Server 2012 R2 and 2012 deduplication volume backups	635
CSV backups.....	635
Windows BMR backups	636
Verifying block based backups	636
Cloning block based backups.....	636

Preparing for block based recovery..... 636
 Performing block based recovery..... 637
 Performing Windows BMR recovery 639
 Performing clone recovery 640
 Troubleshooting block based backup and recovery issues 642

Chapter 21 NetWorker support for NDMP

Overview of NDMP 644
 Components in a NetWorker NDMP environment 644
 Configurations in a NetWorker NDMP environment 645
 NDMP local backup 645
 NDMP backups to non-NDMP devices (NDMP-DSA)..... 646
 Three-party backup with NDMP devices 649
 Pre-configuration requirements for NDMP data operations 650
 NDMP feature requirements 651
 Locale requirements with NDMP..... 652
 Memory and space requirements for NDMP FH updates..... 653
 Performance Considerations 653
 NDMP licensing requirements 654
 Configuring Devices for NDMP operations..... 654
 NDMP device limitations 654
 DinoStor-managed jukeboxes 655
 Configuring NDMP on Isilon filer 655
 Determining NDMP device path names..... 655
 Configuring NDMP devices 657
 Configuring NDMP-DSA devices..... 661
 Configuring the Clone Storage Node 662
 Creating resources to support NDMP clients 662
 Creating and configuring the NDMP client resource 663
 Using the Client Configuration wizard 663
 Configuring the NDMP client manually..... 671
 Performing NDMP backups..... 671
 Performing an NDMP backup from the command line 672
 Troubleshooting NDMP configuration and backup failures 674
 No PAX threads available 675
 Failed to store index entries 675
 IO_WritePage write failed - No space left on device (28): No space left on device 676
 Error reading the FH entries from save through stdin 676
 Cannot find file history info for filename...You may still be able to recover this file with a saveset recovery..... 676
 nsrmdmp_save: data connect: failed to establish connection 677
 nsrmdmp_save: get extension list: communication failure 678
 Cloning NDMP save sets..... 678
 Reporting NDMP Data..... 679
 Querying the NDMP volumes by backup type with the mminfo command ... 679
 Querying the NDMP save sets with the mminfo command 680
 Performing NDMP recoveries 680
 NDMP recovery requirements 681
 Performing an NDMP index-based file-by-file data recovery 683
 Performing a full or Directory Restore of NDMP data by using a save set recovery 687
 Performing destructive save set recoveries for vbb backups 690

	Troubleshooting NDMP recover	691
	RESTORE: could not create path pathname.....	691
	These files were not restored (Restore failed with error, or file/directory specified but not found in backup).....	691
Chapter 22	SNMP Module	
	SNMP traps.....	694
	Configuring NetWorker SNMP notifications.....	694
	Command line options for nsrtrap	695
	Modifying preconfigured NetWorker SNMP notification.....	695
	Creating NetWorker SNMP notifications.....	696
	Configuring SNMP management software.....	696
	NetWorker SMI Network Management Private Enterprise Code	697
	Receiving traps in the SNMP network management software	697
Chapter 23	DiskXtender Data Manager File System Support	
	Supported configurations	700
	Path information	700
	Permissions	700
	DiskXtender Data Manager file system overview.....	700
	File data in a DXDM file system.....	701
	Backup of DXDM file systems	702
	Aborted backups.....	703
	Recovery of DXDM file systems.....	705
	Initiating a recovery.....	705
	Restoring deleted files and previous file versions.....	705
	File system synchronization	706
	Automatic synchronization.....	706
	Manually synchronizing a file	707
Chapter 24	Recovery Support for Windows XP and 2003 Automated System Recovery	
	Microsoft Automated System Recovery.....	710
	Microsoft ASR documentation	710
	NetWorker support for ASR disaster recovery of Windows XP and 2003 clients .	710
	NetWorker ASR save set	710
	Network connection names	711
	ASR Limitations and special considerations	711
	FAT16 partitions are not supported	711
	OEM recovery CDs are not supported.....	711
	Vendor-specific drivers must be installed after Windows installation...	712
	Data and configuration changes since the last backup	712
	Creating an ASR disk	712
	Prerequisites.....	712
	Create an ASR disk locally	712
	Creating an ASR disk by using directed recovery.....	713
	Posting ASR disk creation task	714
	Using the ASR disk to recover a NetWorker client.....	715
	Requirements for an ASR recovery	715
	Performing an ASR recovery.....	716
	Performing a manual recovery on Windows 2003 x64 hosts	717

Components that require special handling after an ASR recovery..... 717
 Verifying the NetWorker client recovery 718

Chapter 25 Windows Bare Metal Recovery

Overview of Windows Bare Metal Recovery..... 720
 Changes from previous versions of NetWorker..... 720
 Supported Operating Systems..... 721
 BMR Support for Windows Features..... 721
 Offline recovery versus online recovery..... 723
 Components of the DISASTER_RECOVERY:\ save set..... 724
 Full versus incremental backups 727
 Synthetic full backups..... 728
 Online recovery of Active Directory, DFSR, or Cluster services 729
 Terminology 730
 Windows BMR Planning 731
 Road map for Windows BMR Planning 731
 Hardware Requirements for Windows BMR Backup and Restore 732
 Configuration requirements for Windows BMR backups..... 733
 Save set planning 734
 Best Practices for Windows BMR 736
 Windows BMR limitations and considerations 739
 Windows BMR Backup..... 747
 Include Windows BMR in scheduled backups..... 748
 Include Windows BMR in manual backups 749
 How to verify a valid Windows BMR backup..... 750
 Windows Bare Metal Recovery to Physical or Virtual Computers..... 751
 To perform a Bare Metal Recovery (BMR) to a Physical Computer..... 751
 To perform a BMR from a Physical Computer to a Virtual Machine (P2V) 763
 Troubleshooting Windows BMR..... 765
 Additional recovery options..... 769

Chapter 26 Volume Shadow Copy Service

Overview of VSS..... 772
 VSS and the backup process 772
 Provider support 774
 The importance of writers..... 774
 Controlling VSS from NetWorker software 775
 Controlling VSS from the Administration window..... 776
 Controlling VSS from the NetWorker client..... 776
 Control VSS from the command-prompt 777
 Globally disabling VSS..... 778
 VSS commands..... 778
 NetWorker Support for Microsoft Applications 781
 Authoritative restores of the Active Directory Application Mode (ADAM) writers..... 781

Chapter 27 Networking and connectivity

Name resolution and connectivity 784
 Troubleshooting name resolution and connectivity errors 784
 Verifying basic connectivity..... 785
 Verifying name resolution 788

Verifying the NetWorker configuration	791
Using multihomed systems	793
Multihomed system requirements	793
Configuring multihomed hosts in a datazone	793
NIC Teaming.....	799
Using DHCP clients.....	799

Chapter 28

Troubleshooting

Before contacting technical support	802
Determining the version of NetWorker software running on a client	802
Displaying diagnostic mode attributes	803
Viewing log files	803
Rendering log files in the current locale at runtime	804
How to view log files with the nsr_render_log program	805
Viewing log files from remote host machines.....	805
Log files from previous releases of NetWorker	806
Filtering log file information displayed by nsr_render_log.....	806
Locating savegroup job logs.....	806
NetWorker functionality issues	807
Backup and recovery.....	807
Backups fail to start when daylight savings time change occurs	809
Shut down NetWorker services prior to any significant changes to system date	809
Clone ID timestamp does not reflect the time the clone was created....	810
Backups fail to stop	810
Memory usage when browsing large save sets	810
Memory usage and nsrjobd	810
Media position errors encountered when auto media verify is enabled	810
PACKET RECEIVE BUFFER and NO ECB counters increase	811
The scanner program marks a volume read-only.....	811
The scanner program requests an entry for record size	811
Limitations for groups containing a bootstrap	811
Index recovery to a different location fails	812
Illegal characters in configurations.....	812
Error backing up large number of clients	812
Hostname aliases	813
Directory pathname restrictions	813
Backup of a new client defaults to level full.....	813
Non-full backup of Solaris files with modified extended attributes	814
Renamed clients cannot recover old backups.....	814
Client file index errors	814
Cannot use the Console interface to stop the savegrp command	815
Aborting a recovery	815
RPC error.....	815
Error message when relocating data	815
Desktop heap size limitation	816
The All save set and duplicate drive serial numbers.....	816
Disk label errors.....	816
Cannot print bootstrap information	817
Server index not forced	817
Copy violation	817
Converting sparse files to fully allocated files.....	817
Backing up large sparse files.....	818
The mminfo -N command is case-sensitive regarding save set names..	818

Devices and Autochangers	820
Additional attributes in the Autochanger resource	820
Maintenance commands	820
Autodetected SCSI jukebox option causes server to stop responding ..	821
Autochanger inventory problems.....	821
Destination component full messages.....	821
Tapes do not fill to capacity.....	822
Tapes get stuck in drive when labelling on Linux Red Hat platform	823
Increasing the value of Save Mount Time-out for label operations	823
Server cannot access autochanger control port	823
Nonrewinding device requirement.....	825
Scanner command behaves differently with adv_file type device.....	825
Sleep times required for TZ89 drive types.....	825
Message displayed when CDI enabled on NDMP or disk FTD.....	826
Verifying firmware for switches and routers	826
Commands issued with nsrjb on a multi-NIC host fail	826
SCSI reserve/release with dynamic drive sharing	826
Device ordering issues	827
Recovery of save sets from a VTL	828
NetWorker locale and code set support	828
Resource database notes	828
Viewing resources	829
Repairing resource database corruption	829
Enabling service mode for NetWorker	829
Network and server communication errors	830
General issues	830
UNIX communication issues.....	832
Binding to server errors	832
Saving remote file systems.....	833
Microsoft Windows issues.....	833
NetWorker archiving and retrieval.....	834
Remote archive request from server fails	834
Multiple save sets appear as a single archive save set	834
Wrong archive pool is selected.....	834
Second archive request does not execute	834
The nsrchive program does not start immediately.....	834
Archive request succeeds but generates error when nsrexecd	
is not running.....	834
Empty annotations in retrieve list	835
Storage nodes.....	835
Storage node affinity errors	835
Storage node timeout errors.....	835
Console error messages and corrective actions	836
Console log files.....	838
The install log	839
The gstd log	839
Console troubleshooting notes and tips	839
Making sure the Console server is running	839
Enabling Java script.....	840
NMC user interface exits unexpectedly	843

Appendix A SYSTEM, VSS SYSTEM, and WINDOWS ROLES AND FEATURES Save Sets

SYSTEM save sets	846
Components of the SYSTEM STATE save set.....	846

Components of the SYSTEM FILES save set.....	847
Components of the SYSTEM DB save set.....	848
Components of the SHAREPOINT save set.....	849
VSS SYSTEM save sets	849
Components of the VSS SYSTEM BOOT save set	849
Components of the VSS SYSTEM FILESET save set	850
Components of the VSS SYSTEM SERVICES save set	850
Windows Server Cluster writers	850
VSS SYSTEM Recovery Considerations.....	850
WINDOWS ROLES AND FEATURES save sets	851
Considerations for the WINDOWS ROLES AND FEATURES save set.....	851

Appendix B Firewall Support

Overview.....	854
Service ports	854
Connection ports.....	854
Special considerations for firewall environments.....	855
Configuring TCP keep alives at the operating system level	855
Configuring TCP keep alives within the NetWorker software	856
Determining service port requirements.....	857
NetWorker client	857
NetWorker storage node.....	858
NetWorker server.....	859
NetWorker Management Console server	861
Configuring the port ranges.....	861
Determining the available port numbers.....	861
Configuring the port ranges in NetWorker	861
Configuring the port ranges on the firewall	864
Examples	867
Troubleshooting.....	872
Backups appear to stop responding or slow down dramatically.....	872
Fallback to RPC portmapper service on port 111	873
Cannot bind socket to connection port range on system hostname	874
Failed to bind socket for service_name service: Can't assign requested address.....	874
Service is using port port_number which is outside of configured ranges: range	875
Connection refused	875
Connection reset by peer.....	875
Unable to obtain a client connection to nsrmmgd (version #) on host hostname.....	875
nsrmdmp_save: data connect:failed to establish connection	875
Unable to execute savevfs job on host hostname: Remote system error - No route to host	876

Appendix C Backing Up and Restoring a Microsoft DFS

Overview of a Microsoft DFS.....	878
Save Set ALL-DFSR	878
DFS topology information.....	879
Configuring a scheduled DFS backup	879
Restoring a DFS.....	881
Authoritative restores of DFS Replication writers	882
Non-authoritative restores of DFS Replication writers	882
DFS backups and restores for Windows 2003	883

Appendix D	Additional Features of the Microsoft Windows Server	
	NetWorker Module for Microsoft	886
	Active Directory	886
	Backing up Active Directory	886
	Recovering Active Directory	886
	Encrypting file system	886
	Event logs	887
	Internet Information Server	888
	Windows registry	888
	Sparse files	888
	Windows Change Journal.....	888
	NetWorker support for Change Journal.....	889
	Configuring NetWorker software to use the Change Journal.....	890
	Advanced Configuration and Power Interface.....	891
	NetWorker support for ACPI	892
	Windows print queues	892
	Windows Optimized Deduplication	892
	Detecting Deduplication in a Backup.....	893
	Data Deduplication Backup and Restore.....	893
	Windows Data Deduplication Volume Best Practices	895
	Recommended Deduplication Workloads	896
Appendix E	UNIX and Linux Platform-Specific Notes	
	Solaris	898
	Support for Solaris zones	898
	NetWorker executables not found for Solaris client	898
	How to obtain support for devices not supported by Solaris	898
	Extended file attribute data included in Save Set File Size attribute	899
	The inquire command and Solaris 10	899
	Linux.....	899
	Backup considerations for Linux raw disk partitions	899
	Configure Linux operating system to detect SCSI devices	899
	The inquire command and the Scan for Devices operation do not detect more than 128 tape devices	899
	Configuration requirements for the inquire command.....	900
	Linux Journaled file system support.....	900
	HP-UX	900
	Autochanger installation on an HP-UX system	900
	How to test the device driver and device file installation	902
	Errors from unsupported media in HP tape drives	903
	Unloading tape drives on an HP-UX server or storage node	903
	SCSI pass-through driver required for HP-UX autochangers	904
	Symbolic link entries in the fstab file	904
	Customized backup scripts	904
	AIX.....	904
	STK-9840 drives attached to AIX.....	904
	LUS driver operation on AIX	904
	Recovering set-group-id or setuid binaries and files	904
Appendix F	MAC OS X Support	
	Support for Mac OS X.....	906
	Mac OS X metadata support	906
	Supported file systems.....	906

	Mac OS X backup considerations	906
	Scheduling a NetWorker client backup on Mac OS X.....	906
	Performing a manual backup on Mac OS X	908
	Recovering files and directories on Mac OS X using the command prompt .	908
	Task 1: Browse backed-up Mac OS X data	909
	Task 2: Recover individual files or directories	909
	Recovering files and directories on Mac OS X using NetWorker Recover	909
	Starting NetWorker Recover for the first time	910
	Navigating the NetWorker Recover window.....	911
	Configuring NetWorker Recover	914
Appendix G	Direct SCSI Backup and Recover	
	Introduction to direct SCSI backup and recover	920
	System requirements	920
	Unsupported features	920
	Performing direct SCSI backup	921
	Backing up data on a Symmetrix BCV device	921
	Backing up data on a raw device	922
	Backing up data from the command line	923
	Performing direct SCSI recover	924
	Recovering data to a Symmetrix BCV device	924
	Recovering data to a raw device	926
	Licensing	927
Appendix H	Security Configuration Settings	
	Access control settings.....	930
	User authentication.....	930
	User authorization.....	930
	Component access control	930
	Log settings	932
	Log files and their descriptions	932
	Log management and retrieval	933
	NetWorker Accountability	934
	Security audit logging overview	935
	Security audit logging configurations	935
	Security audit logging interoperability	939
	Modifying the security audit log resource	939
	Audit message format	942
	Communication security settings	943
	Port usage.....	943
	Encrypting backup data.....	944
	Encryption for cloud backup data	944
	Federal Information Processing Standard Compliance	944

Glossary

Index

PREFACE

As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your EMC representative if a product does not function properly or does not function as described in this document.

Note: This document was accurate at publication time. New versions of this document might be released on the EMC online support website. Check the EMC online support website to ensure that you are using the latest version of this document.

Purpose

This document describes how to configure and use EMC 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.

Related documentation

The following EMC information products provide additional information:

- ◆ *EMC NetWorker Installation Guide*
Explains how to install or update the NetWorker software for the clients, console, and server on all supported platforms.
- ◆ *EMC NetWorker Cluster Integration Guide*
Explains how to set up and configure the NetWorker software in supported cluster environments.
- ◆ *EMC NetWorker Release Notes*
Contain information on new features and changes, fixed problems, known limitations, environment and system requirements for the latest NetWorker software release.
- ◆ *EMC NetWorker and EMC Data Domain Deduplication Devices Integration Guide*
Provides planning and configuration information on the use of Data Domain devices for data deduplication backup and storage in a NetWorker environment.
- ◆ *EMC NetWorker and VMware Integration Guide*
Provides planning and configuration information on the use of VMware in a NetWorker environment.
- ◆ *EMC NetWorker Snapshot Management Integration Guide*
Provides the ability to catalog and manage snapshot copies of production data that are created by using mirror technologies on EMC storage arrays.

- ◆ *EMC NetWorker and EMC Avamar Integration Guide*
Provides planning and configuration information on the use of Avamar in a NetWorker environment.
 - ◆ *EMC NetWorker Cloning Integration Guide*
Contains planning, practices, and configuration information for using the NetWorker, NMM, and NMDA cloning feature.
 - ◆ *EMC NetWorker Error Message Guide*
Provides information on common NetWorker error messages.
 - ◆ *EMC NetWorker Performance Optimization Planning Guide*
Contains basic performance tuning information for NetWorker.
 - ◆ *EMC NetWorker Server Disaster Recovery and Availability Best Practices Guide*
Explains how to design and plan for a NetWorker disaster recovery. However, it does not provide detailed disaster recovery instructions. The Disaster Recovery section of the NetWorker Procedure Generator (NPG) provides step-by-step disaster recovery instructions.
 - ◆ *EMC NetWorker Command Reference Guide*
Provides reference information for NetWorker commands and options.
 - ◆ *EMC NetWorker Licensing Guide*
Provides information about licensing NetWorker products and features.
 - ◆ *EMC NetWorker License Manager 9th Edition Installation and Administration Guide*
Provides installation, setup, and configuration information for the NetWorker License Manager product.
 - ◆ *EMC NetWorker Software Compatibility Guide*
Lists supported client, server, and storage node operating systems for NetWorker, NetWorker Modules, and options.
 - ◆ EMC 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 Help, click Help in the main menu.
 - ◆ EMC NetWorker User Online Help
The NetWorker User program is the Windows client interface. The NetWorker User Online Help describes how to use the NetWorker User program, which is the Windows client interface connect to a NetWorker server to back up, recover, archive, and retrieve files over a network.
 - ◆ EMC NetWorker SolVe Desktop (also known as the NetWorker Procedure Generator (NPG))
The NetWorker Procedure Generator (NPG) is a stand-alone Windows application used to generate precise user driven steps for high demand tasks carried out by customers, support, and the field. With the NPG, each procedure is tailored and generated based on user-selectable prompts.
- To access the NetWorker Procedure Generator, log on to <https://support.emc.com/> and search for NetWorker Procedure Generator. You must have a service agreement to use this site.

- ◆ Technical Notes and White Papers
Provides an in-depth technical perspective of a product or products as applied to critical business issues or requirements. Technical Notes and White paper types include technology and business considerations, applied technologies, detailed reviews, and best practices planning.

Conventions used in this document

EMC uses the following conventions for special notices:

NOTICE

NOTICE is used to address practices not related to personal injury.

Note: A note presents information that is important, but not hazard-related.

Typographical conventions

EMC uses the following type style conventions in this document:

Normal	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> • Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus • Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities • URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications
Bold	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and man pages Used in procedures for: <ul style="list-style-type: none"> • Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus • What the user specifically selects, clicks, presses, or types
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> • Full titles of publications referenced in text • Emphasis, for example, a new term • Variables
Courier	Used for: <ul style="list-style-type: none"> • System output, such as an error message or script • URLs, complete paths, filenames, prompts, and syntax when shown outside of running text
Courier bold	Used for specific user input, such as commands
<i>Courier italic</i>	Used in procedures for: <ul style="list-style-type: none"> • Variables on the command line • User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections — the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

EMC support, product, and licensing information can be obtained as follows:

Product information. For documentation, release notes, software updates, or information about EMC products, licensing, and service, go to the EMC Online Support website (registration required) at:

<https://support.emc.com/>

Technical support — For technical support, go to EMC online support and select Support. On the Support page, you will see several options, including one to create a service request. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Online communities — Visit EMC Community Network <https://community.EMC.com/> for peer contacts, conversations, and content on product support and solutions. Interactively engage online with customers, partners, and certified professionals for all EMC products.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to:

DPAD.Doc.Feedback@emc.com

REVISION HISTORY

Email your clarifications or suggestions for this document to:

DPAD.Doc.Feedback@emc.com

The following table lists the revision history of this document.

Revision	Date	Description of added or changed sections
04	May 2014	Formatting changes, correction to DFSR in Appedic C.
03	December 2013	Update to data deduplication information in Appendix D. Update to document feedback email DPAD.Doc.Feedback@emc.com . Updated VMware details for P2V BMR in Windows Bare Metal Recovery chapter. Format updates.
02	November 2013	Update to data deduplication information. Added save set restart information to Troubleshooting chapter.
01	November 2013	First release of this document for the <i>EMC NetWorker 8.1 SP1 Release</i> .

CHAPTER 1

Overview

This chapter covers these topics:

- ◆ The NetWorker environment 32
- ◆ NetWorker features 33
- ◆ NetWorker user interfaces 35
- ◆ Common NetWorker tasks 47
- ◆ NetWorker services 51

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. 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 clients and servers from a web-enabled, graphical interface.

NetWorker components

[Figure 1 on page 33](#) illustrates the main components in a NetWorker environment.

Console server

The NetWorkerConsole server manages all NetWorker servers and clients. The Console server also provides reporting and monitoring capabilities for all NetWorker servers and clients.

Console user interface

The Console server uses a graphical interface run from any computer that has a supported web browser and Java Runtime Environment (JRE). The NetWorker Installation Guide provides information on supported web browsers and supported versions of the JRE. Multiple users can access the Console server concurrently from different browser sessions. A computer that hosts the web browser can also be a NetWorker client.

NetWorker server

NetWorker servers provide services to back up and recover data on NetWorker client computers in a datazone.

Datazone

A datazone is a single NetWorker server and its client computers. Datazones can be added as backup requirements increase.

NetWorker storage node

A NetWorker storage node can be used to improve performance by offloading from the NetWorker server much of the data movement involved in a backup or recovery operation.

NetWorker client

A NetWorker client computer is any computer whose data must be backed up. The NetWorker Console server, NetWorker servers, and NetWorker storage nodes are also NetWorker clients.

Deduplication storage systems

NetWorker supports backup data deduplication on EMC Avamar® storage nodes and on Data Domain® storage systems.

The *NetWorker Avamar Integration Guide*, provides detailed information about setting up Avamar deduplication to work with NetWorker.

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

Virtual environments

NetWorker clients can be created for virtual machines for either traditional backup or VMware Consolidated Backup (VCB). Additionally, the NetWorker software can automatically discover virtual environments and changes to those environments on either a scheduled or on-demand basis and provides a graphical view of those environments.

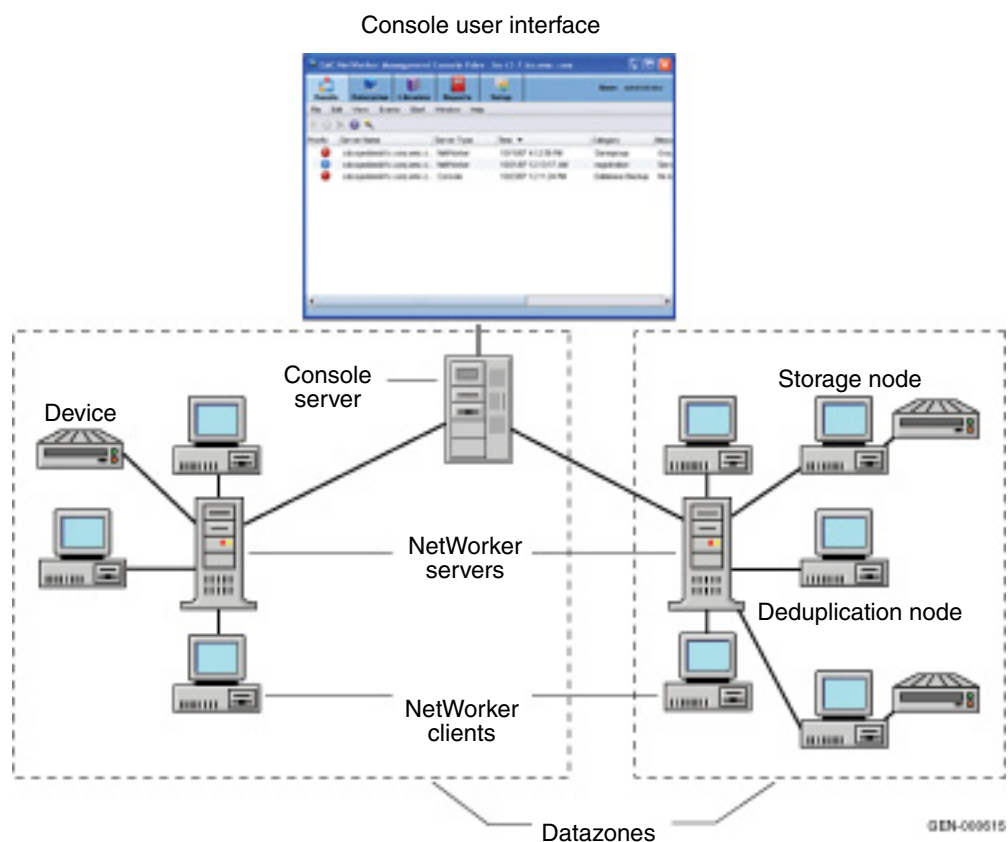


Figure 1 NetWorker components

NetWorker features

This section describes the major features that distinguish the NetWorker software. Optional additions to the NetWorker software are also listed.

Performance features

Standard NetWorker performance features include:

- ◆ Server parallelism, which enables several *save streams* to flow to the server or storage node at the same time.
- ◆ Multiplexing, which enables more than one save stream to write to the same device at the same time.

- ◆ Client parallelism, which enables the client to send more than one save stream at a time.
- ◆ Client Direct, which enables client backups to bypass the storage node and send deduplicated backup data directly to DD Boost storage devices, or to send non-deduplicated backup data directly to AFTD storage.
- ◆ Session management, which enables one to control the number of save streams per device to maximize the performance of each device.
- ◆ Backup to file-based devices and optional subsequent staging and cloning to near-line or offline volumes.
- ◆ Backup to a cloud storage configuration.

Ease of use

NetWorker software provides tools to make protection of critical data easy to manage. With these tools, you can:

- ◆ Use either the graphical interfaces or command-line programs to manage NetWorker tasks and functions.
- ◆ Use wizards to set up the following NetWorker items:
 - Client resources
 - AFTD and Data Domain devices
 - Common Console configuration tasks
 - LDAP user authentication
- ◆ Administer and configure NetWorker functions from any network computer with a web browser.
- ◆ Grant permission to provide directed recovery operations. Directed recovery is the capability for recovery of one client's data to another client computer.
- ◆ Obtain immediate answers to questions by accessing online help and UNIX man pages. Microsoft Windows users can also access the NetWorker command reference guide, which provides information similar to the UNIX man pages.
- ◆ Take advantage of the automatic media management feature to enable the NetWorker server or storage node to label and mount volumes as needed for backups.
- ◆ Drag-and-drop functionality allows for an easy transfer of single or multiple objects.
- ◆ Use the integrated knowledge base and technical bulletins at the EMC online support[®] website to find answers to common questions.
- ◆ Automatically discover and view a graphical map of virtual environments.
- ◆ Set up NetWorker Console server authentication to an external LDAP v3 compliant server.
- ◆ Support for automated Windows Bare Metal recovery.

Scalability

NetWorker software can be scaled as storage management needs grow. For example, you can:

- ◆ Upgrade the basic level of server functionality, add support for additional (or larger) autochangers, add support for more clients, or add optional software modules without the need to reinstall the server software.
- ◆ Add special NetWorker Module client software to back up databases and other non-file-system data.
- ◆ Add support for remote storage nodes to control backup devices, while the data management tasks remain centralized on a controlling NetWorker server.
- ◆ Add the NetWorker License Manager (NLM) software to administer all of your network's EMC software licenses from a single server.

Optional software additions

Optional additions to the NetWorker software include:

- ◆ NetWorker Autochanger Module
- ◆ NetWorker Silo Software Module
- ◆ NetWorker Archive Module
- ◆ NetWorker Database Modules (for backing up several types of databases)
- ◆ NetWorker SNMP (Simple Network Management Protocol)
- ◆ NDMP (Network Data Management Protocol) support
- ◆ Cluster support
- ◆ NetWorker License Manager
- ◆ Advanced reporting capability

NetWorker user interfaces

The NetWorker application consists of these user interfaces:

- ◆ [“NetWorker Management Console interface” on page 35](#)
- ◆ [“NetWorker client interface” on page 44](#)
- ◆ [“NetWorker character-based interface” on page 47](#)
- ◆ [“NetWorker command-line interface” on page 47](#)

NetWorker Management Console interface

The interface for NetWorker Management Console, also called the NetWorker Console, consists of two main windows:

- ◆ Console window
- ◆ Administration window

Console window

When NetWorker software is started, the Console window appears as shown in [Figure 2 on page 36](#).

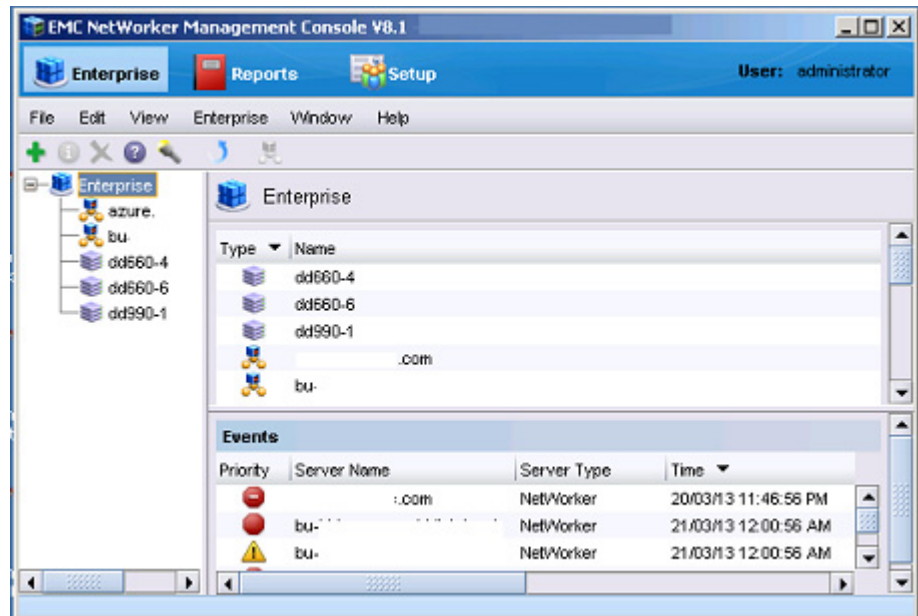





Figure 2 Console window

The Console window is the first point of access for NetWorker tasks. [Table 1 on page 36](#) lists the task-based windows that can be opened from the Console window taskbar.

Table 1 Windows opened from the Console window

Button	Window	Description
	Enterprise	Select a NetWorker server to manage and monitor the server and its backup clients. The Enterprise window allows you to open the Administration window for a NetWorker server.
	Reports	Configure and view Console reports. Chapter 15, “Enterprise reporting and events monitoring.” provides information about reports.
	Setup	Control administrative functions: <ul style="list-style-type: none"> • User management — Add, edit, and delete Console user accounts, restrict user views of servers. “NMC server authentication” on page 504 provides information about user management. • License management — Manage NetWorker licenses. The new for 7.6 Service Pack 1 NetWorker <i>Licensing Guide</i> provides information about license management.
	Events	View important messages about all NetWorker servers that have been added as Enterprise applications, as well as the Console server, and Avamar server. “Events” on page 456 provides information about managed events.

Administration window

NetWorker servers are managed through the Administration window as shown in [Figure 3](#) on page 37.

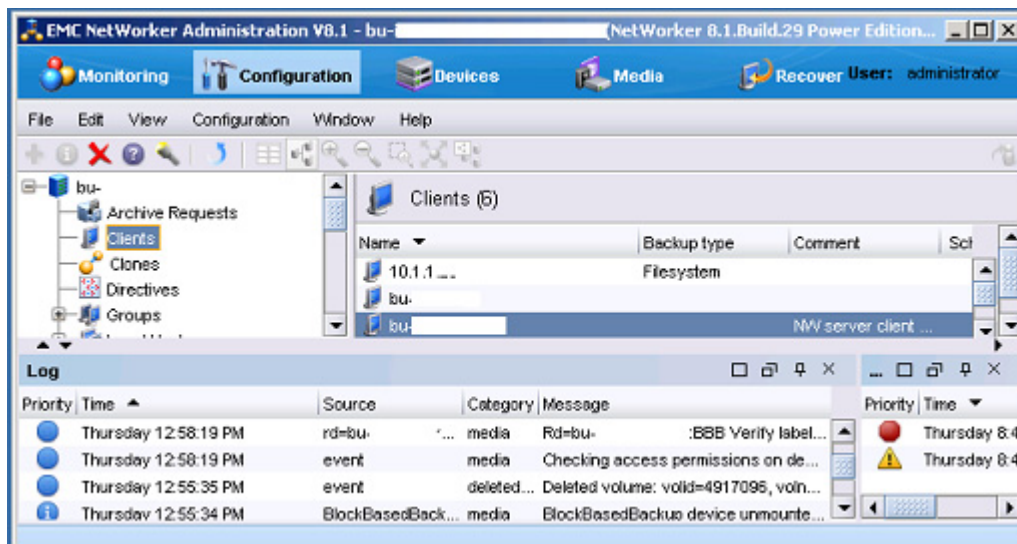







Figure 3 Administration window

You can toggle between the **Administration** window and the Console window.

[Table 2 on page 37](#) lists the windows that can be launched from the **Administration** window taskbar.

Table 2 Windows launched from the Administration window

Button	Window	Description
	Monitoring	Monitor various activities related to the NetWorker server. For example, you can monitor the progress of a scheduled backup and view any alerts. A portion of the Monitoring window appears at the bottom of the Administration window at all times, providing monitoring information on Log Messages and Alerts.
	Configuration	Manage the NetWorker server and its resources such as clients, backup schedules, and policies. For example, you can create a backup schedule, add NetWorker clients, and apply the backup schedule to several NetWorker clients.
	Devices	Add, configure, and operate single or multiple devices, libraries, and silos for the NetWorker server.
	Media	Manage activities and resources related to backup volumes. For example, you can mount a backup volume or create a template for labeling backup volumes.
	Recover	Configure recover configurations and schedule recover jobs for NetWorker hosts from a centralized location, the Console server.

Starting Console for the first time

These steps assume that the NetWorker software is installed and that all of the software and hardware requirements have been met on the computer that will access Console. The *NetWorker Installation Guide* provides more information.

To open Console for the first time:

1. From a supported a web browser session, type the URL of the Console server:

```
http://server_name:http_service_port
```

where:

- *server_name* is the name of the Console server.
- *http_service_port* is the port for the embedded HTTP server. The default HTTP port is 9000.

For example:

```
http://houston:9000
```

2. On the **Welcome** window, click **Start**.
3. On the **Security Warning** window, click **Start** to install and run **NetWorker Console**.
4. On the **Licensing Agreement** window, select **Accept**.
5. If you did not install the appropriate JRE version on the system, a prompt to install JRE appears. Follow the on screen instructions to install JRE.
6. On the **Welcome to the Console Configuration Wizard** window, click **Next**.
7. On the **Set Administrator password** window, type the NMC password, then click **Next**.
8. On the **Set Database Backup Server** window, specify the name of the NetWorker server that will backup the Console server database, and then click **Next**.
9. On the **Add NetWorker servers** window, specify the names of the NetWorker server that the Console server will manage, one name per line. Leave the default options **Capture Events** and **Gather Reporting Data** enabled.

Consider the following:

- Enable the **Capture Events** option to allow the Console server to monitor and record alerts for events that occur on the NetWorker server.
 - Enable the **Gather Reporting Data** option to allow the Console server to automatically collect data about the NetWorker server and generate reports. The *NetWorker Administration Guide* on the EMC Online Support Site describes on how to run reports and the reports that are available.
10. Click **Finish**.

The **Console** window and the **Getting Started** window appear.

1. In the **Enterprise window**, right click the NetWorker server and select **Launch Application**.

Start Console after the first time

After Console has been started the first time, you can restart it by using any of these methods:

- ◆ Point the browser to the same URL. [“Starting Console for the first time” on page 38](#) provides information.
- ◆ Double-click the **NetWorker Console** product name in the **Java Web Start Application Manager**.
- ◆ Double-click the desktop button, if one was set up through the Java Web Start Application Manager.

Opening the Administration Window

To add and select a NetWorker server and open the Administration window:

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 your 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 is launched as a separate application.

Sorting tables

Console software’s organization and display of tabular information can be changed. Tables can be sorted by column heading, and then by alphabetic or numeric order within those columns.

To rearrange information in a table:

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.

Example 1 Sorting managed events

John wants to see all the managed events about servers that were unreachable by the Console server.

1. From the Console window, John clicks **Events**.
2. He clicks and drags the **Message** column until it is over the **Priority** column, where he drops it.

3. He clicks the **Message** column heading so that a down-arrow appears.

Now he can scan down the list of messages until he finds three servers, all with the message Unable to connect to server.

John could also generate a **Managed Event Details** report to get the same information, which could be printed or exported for use in another application. [Chapter 15, “Enterprise reporting and events monitoring”](#) provides more information about reports.

Sorting selected rows in a table

To sort selected rows in a table:

1. From the **Edit** menu, select **Find** or press **Ctrl + F** to view the **Find** panel.
2. Click the rows to be selected or select rows by using the Find criteria.
3. Select **Sort Selected**.

Selected rows will be sorted to the top of the table. This 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.

Sorting multiple columns in a table

To sort by two or more columns in a table:

1. Click the column to be used as the last sort key.
2. Click the column to be used as the next-to-last sort key and so on until the primary column is selected.

For example, given a large table of events, you can select the **Time** column as the tertiary sort key, the **Category** column for the secondary sort key, and the **Server** name as the primary sort key. The resulting display would list the servers in alphabetical order, and the events for each server would be grouped by category and would display in chronological order.

Displaying columns in a table

To 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 on a table header and selecting **Add Column** from the drop-down.

Multiple resource editing

Within NMC's Configuration window, you can edit an attribute for multiple resources at the same time. For example, if you want all clients within a group to have their backup schedule changed from the default to “Full Every Friday”, do the following:

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 will change when the cursor is in the column, as shown in [Figure 4 on page 41](#).

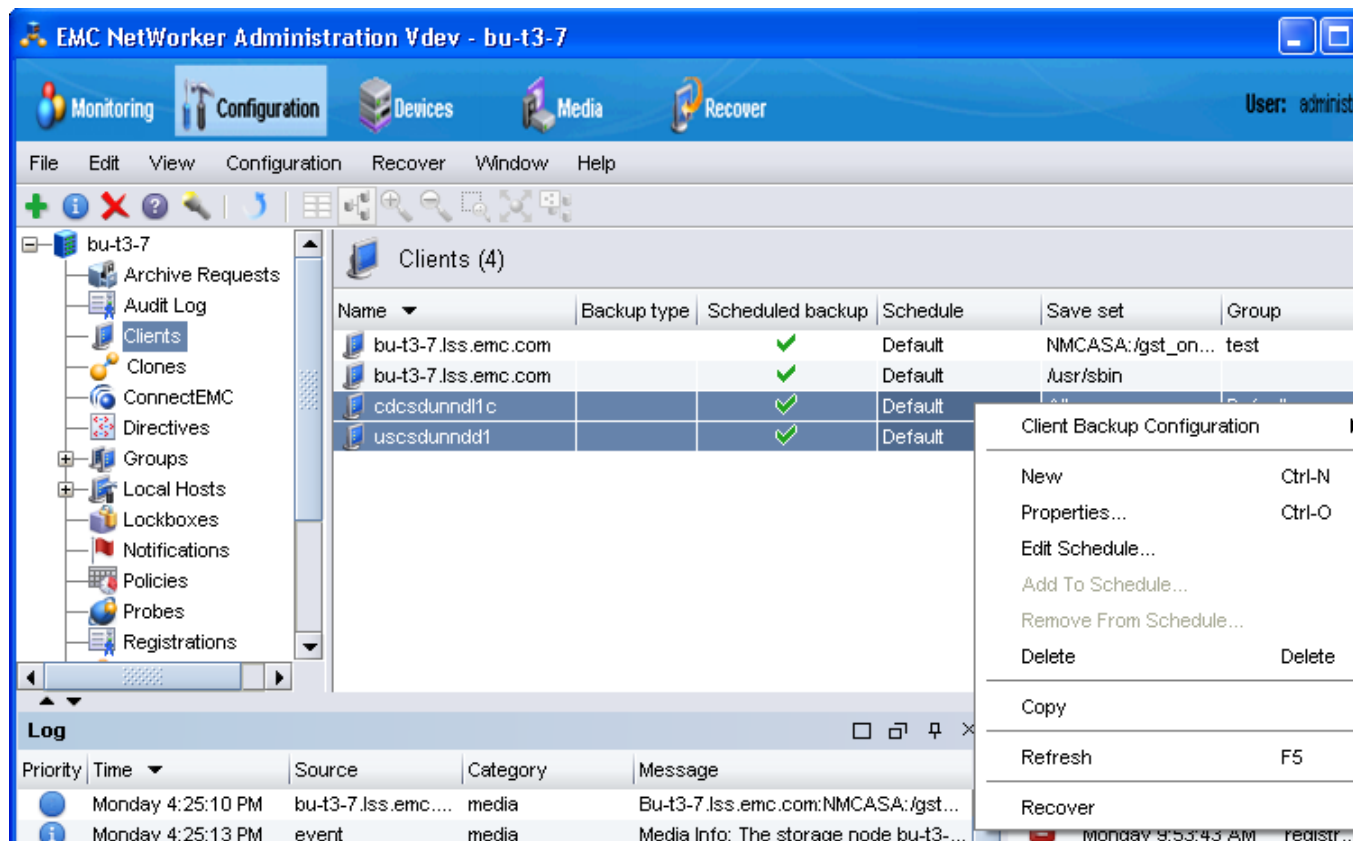


Figure 4 Multiple resource editing in NMC

- Right-click in that column and select from the options available in the drop-down. 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:

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

Drag-and-drop functionality

Drag-and-drop functionality is available in the Console and Administration interfaces to perform the following tasks:

- ◆ “[Drag-and-drop between resource types in the Console window](#)” on page 41
- ◆ “[Client and group management in the Administration window](#)” on page 42
- ◆ “[Library operations in the Devices window](#)” on page 43
- ◆ “[Copy and paste tabular information to operating system clipboard](#)” on page 43

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, holding down 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, hold down the Ctrl key, and drag the contents to a new folder. Select a collection of folders and/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.

NOTICE

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 Configuration window to do the following:

- ◆ Copy selected clients to a new NetWorker group. Expand a group in the directory tree and select. Drag-and-drop the client objects in the Client Summary table to a new group in the directory tree.
- ◆ Copy selected clients from one NetWorker group to a new group. Select a group in the directory tree and move clients from the Client Summary table to another NetWorker group.
- ◆ Move selected clients to a new NetWorker group. Expand a group in the directory tree and select one or more clients. Drag-and-drop the client objects in the Client Summary table to a new group in the directory tree. This will remove the client objects from the initial NetWorker group.
- ◆ Change a group's current schedule to a selected schedule. Select a group in the directory tree to display the group objects. Drag-and-drop a schedule in the Schedule Summary table to a different group in the directory tree.

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 by selecting multiple slots in the Slots table and dragging them anywhere in to the Devices table.
- ◆ Unmount a volume from a selected device and deposit it back in its designated slot. To unmount a volume, select an individual device from the Devices table and drag it anywhere in the Slots table. The volume image will appear in the corresponding slot.
- ◆ Unmount multiple volumes from a selected device and deposit them back in their designated slot. To unmount multiple volumes, select the devices from the Devices table and drag them anywhere in the Slots table. The volumes will appear 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.

NOTICE

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 and EMC AlphaStor® only)

Setting user interaction preferences

Depending on the window button that was selected (see [Table 2 on page 37](#)), 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.

To set user preferences:

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

NetWorker client interface

The client interface is where users can recover data and perform manual backup and archive operations. Manual operations are not scheduled. Instead, they are performed when a user wants to back up or archive one or more files on the NetWorker client immediately. Scheduled backup and archive operations are set up through the Console interface. For information about the Console interface, see [“NetWorker Management Console interface” on page 35](#).

Windows client interface

The NetWorker User program shown in [Figure 5 on page 44](#) is the Windows client interface.




Figure 5 NetWorker User program

Starting the NetWorker User program on Windows

To start the NetWorker User program, perform one of the following:

- ◆ Click the Windows Start button and select **Programs>EMC NetWorker>NetWorker User**.
- ◆ From the Administration window, click **Start** on the main menu, select **“NetWorker User...”** If the NetWorker Module for Microsoft Applications is installed on the client computer, this operation starts the NetWorker Module for Microsoft Applications instead.

The NetWorker client package must be installed on the host where you start the NetWorker User program. Otherwise, you will 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. [Table 3 on page 45](#) 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.


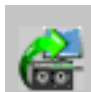


Table 3 NetWorker User Groups requirements

Logged on	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. [Table 4 on page 45](#) describes the function of each button.

Table 4 NetWorker User toolbar functions

Button	Name	Function
	Backup	Starts an 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 just backed up are the same as those currently on the disk.

Browse window

A browse window in the NetWorker User program appears when you select one of the following:

- ◆ A toolbar button
- ◆ A Backup, Recover, Archive, Verify, or Local Directive command from the NetWorker User File menu

The browse window, shown in [Figure 6 on page 46](#) displays the directory tree of the file system that is being browsed.

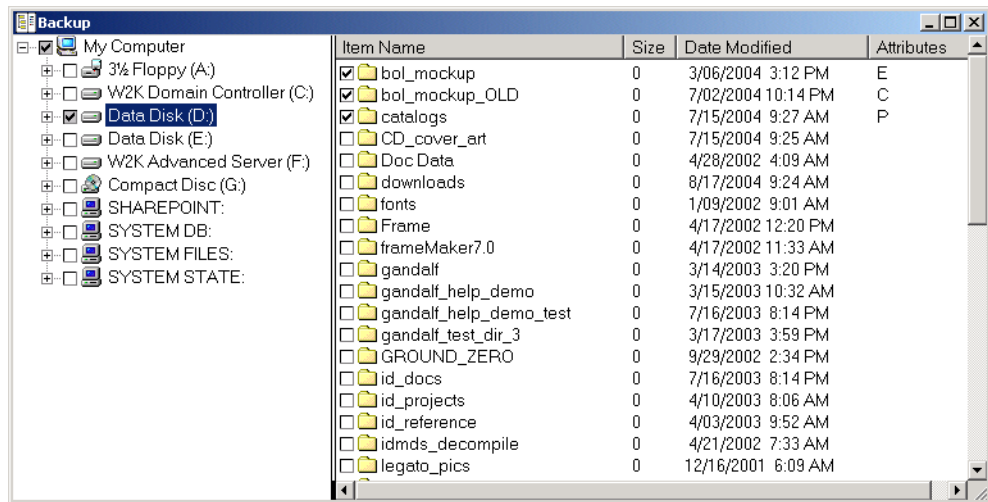


Figure 6 Example of the browse window

NOTICE

When you mark a disk volume or directory for an operation, all of its nested subdirectories and files are also marked.

A checkmark 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. [“Encrypting backup data” on page 108](#) provides more information.
- ◆ **E** — The item is marked for password-protection and encryption, using the PW2 ASM. [“Encrypting backup data” on page 108](#) provides more information.
- ◆ **C** — The item is marked for compression. [“Compressing backup data” on page 109](#) provides more information.

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.

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. [“Task 6: Create a backup Client resource” on page 64](#) provides information about creating a Client resource.

To connect to a NetWorker server:

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**.

UNIX client interfaces

On UNIX, use command line utilities to perform manual backups (the **save** command), archiving (**nsrarchive**) and recovery operations (**recover**). For more information on these commands, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

NetWorker character-based interface

Use the NetWorker character-based interface (**nsradmin**) to perform configuration and management tasks for a NetWorker server.

To start the **nsradmin** interface, type this command:

```
nsradmin
```

For more information about **nsradmin**, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

NetWorker command-line interface

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

Common NetWorker tasks


This section identifies some of the most common tasks performed with the NetWorker Console. These include:

- ◆ [“How to add a new host” on page 47](#)
- ◆ [“How to configure devices” on page 48](#)
- ◆ [“How to label media” on page 49](#)
- ◆ [“How to schedule a backup” on page 50](#)
- ◆ [“How to view failed backups” on page 50](#)
- ◆ [“How to perform a manual backup” on page 50](#)

[“NetWorker Management Console interface” on page 35](#) has information about starting the NetWorker Console.

How to add a new host

To add a new host in the NetWorker Console:

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, DNS name, or WINS name of the NetWorker server and click **Next**.
6. In the **Select Host** Type window, select **NetWorker** and click **Next**.
7. In the **Manage NetWorker** window, leave the default options **Capture Events** and **Gather Reporting Data** enabled.
 - Enable the **Capture Events** option to allow the Console server to monitor and record alerts for events that occur on the NetWorker server.
 - Enable the **Gather Reporting Data** option to allow the Console server to automatically collect data about the NetWorker server and generate reports. The *NetWorker Administration Guide* on the EMC Online Support Site describes on how to run reports and the reports that are available.
8. Click **Finish**.




How to configure devices

Configure one of the following devices to test the NetWorker software.

- ◆ [“Stand-alone tape device” on page 48](#)
- ◆ [“Stand-alone advanced file type device” on page 48](#)
- ◆ [“Autochanger or silo” on page 49](#)

Stand-alone tape device




To configure a stand-alone tape device:

1. Log in to the Console as a NetWorker Administrator.
2. Click the **Enterprise** button  on the taskbar.
3. Highlight a host in the navigation tree, right-click **NetWorker** and select **Launch Application**. The **NetWorker Administration** window appears.
4. Click the **Configuration** button  on the taskbar.
5. Click the **Devices** button  on the taskbar.
6. In the navigation tree view, right-click a host and select **Scan for Devices**.
7. From the right pane, select the new device.
8. From the **Devices** menu, select **Devices > Device Operations > Label**.
9. In the **Label** window, verify the information and click **OK**.

Stand-alone advanced file type device




To configure a stand-alone file or advanced file device:

1. Log in to the Console as a NetWorker Administrator.

2. Click the **Enterprise** button  on the taskbar.
3. Highlight a host in the navigation tree, right-click **NetWorker** and select **Launch Application**. The **NetWorker Administration** window appears.
4. Click the **Configuration** button  on the taskbar.
1. Click the **Devices** button  on the taskbar.
2. In the select Device Type window, select Advanced File Type Device (AFTD), then click **Next**.
3. From the **File** menu, select **New Device Wizard**.
4. In the **Select Storage Node** window, click **Next**.
5. In the **Select the Device Path** window, select an empty folder or create a new folder on the NetWorker server, then click **Next**.
6. In the **Configure Attributes** window, specify a name for the new device, for example: myaftd, then click **Next**.
7. In the **Label and Mount Devices** window, click **Next**.
8. In the **Review Configuration Settings** window, click **Configure**.
9. Click **Finish**.



Autochanger or silo


To configure a new library resource:

1. Log in to the Console as a NetWorker Administrator.
2. Click the **Enterprise** button  on the taskbar.
3. Highlight a host in the navigation tree, right-click **NetWorker** and select **Launch Application**. The **NetWorker Administration** window appears.
4. Click the **Configuration** button  on the taskbar.
1. Click the **Devices** button  on the taskbar.
2. From the left pane, select **Storage Nodes**.
3. Right-click the storage node for the device and select **Configure All Libraries**.
4. Fill in the requested information and click **Start Configuration**
5. Click **Finish**.

How to label media

To label tapes from the NetWorker Console:



1. Log in to the Console as a NetWorker Administrator.
2. Click the **Enterprise** button  on the taskbar.
3. Highlight a host in the navigation tree, right-click **NetWorker** and select **Launch Application**. The **NetWorker Administration** window appears.
4. Click the **Configuration** button  on the taskbar.

5. Click the **Devices** button  on the taskbar.
6. In the navigation tree view, expand **Libraries** and highlight a library, or select **Devices**.
7. In the **Device list**, right-click a device and select **Label**.

[Chapter 6, “Media Management”](#) provides information about labeling tapes.

How to schedule a backup




To schedule backups from the NetWorker Console:

1. Log in to the Console as a NetWorker Administrator.
2. Click the **Enterprise** button  on the taskbar.
3. Highlight a host in the navigation tree, right-click **NetWorker** and select **Launch Application**. The **NetWorker Administration** window appears.
4. Click the **Configuration** button  on the taskbar.
5. Define schedules, groups, and clients.

[“Planning a scheduled backup” on page 58](#) provides information about scheduling backups.

How to view failed backups

To see whether any backups have failed:

1. Log in to the Console as a NetWorker Administrator.
2. Click the **Enterprise** button  on the taskbar.
3. Highlight a host in the navigation tree, right-click **NetWorker** and select **Launch Application**. The **NetWorker Administration** window appears.
4. Click the **Configuration** button  on the taskbar.
5. Click **Monitoring**  .
6. Select **Groups** in the docking panel.

[Chapter 16, “NetWorker server events reporting and monitoring”](#) provides information about viewing backup status information.

How to perform 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” on page 50](#)
- ◆ [“Performing a manual backup on UNIX” on page 51](#)

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.

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* on the EMC Online Support Site describes how to use the **nsr_render_log** program.

The location of the **daemon.raw** 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 [Chapter 28, “Troubleshooting,”](#)
- 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 system 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.

NetWorker services

This section provides information about the main services and programs for the NetWorker server, NetWorker Storage Node, NetWorker Client and NetWorker Management Console server. It also describes how to start and stop these services.

For more information about:

- ◆ Main NetWorker services, The *EMC NetWorker Command Reference Guide* or the UNIX man pages provides more information.
- ◆ Service port requirements when configuring a firewall. [Appendix B, “Firewall Support”](#) provides more information.

Services and programs on the NetWorker server

[Table 5 on page 52](#) describes the main services and programs that provide NetWorker software functionality.

Table 5 Services or programs on the NetWorker server

Service or program	Function
nsrexecd	<ul style="list-style-type: none"> • Authenticates and processes the NetWorker server remote execution requests and executes the save and savefs programs on the client.
nsrd	<ul style="list-style-type: none"> • Is 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, nsrd will spawn an agent process, ansrd.
nsrmmdbd	<ul style="list-style-type: none"> • Provides media database management services to the local nsrd and nsrmmmd services and records entries in the media database. This is the media management database service.
nsrjobd	<ul style="list-style-type: none"> • Monitors NetWorker activity during a backup or recovery operation.
nsrindexd	<ul style="list-style-type: none"> • Provides a method for inserting entries into the client file index that is based on information passed by the save program.
nsrmmgd	<ul style="list-style-type: none"> • Manages media library operations. • Provides an RPC-based service that manages all jukebox operations on behalf of the nsrd service. • The nsrd service starts only one instance of nsrmmgd on the NetWorker server as needed.
nsrlogd	<ul style="list-style-type: none"> • Provides support for the NetWorker audit log service and is configured to run on the NetWorker server by default.
nsrcpd	<ul style="list-style-type: none"> • Started automatically when there is a request to start up the remote client software installation service (client push). • Allows users to distribute and upgrade client software from a centralized software repository across a network.

Services and programs on the NetWorker client

[Table 6 on page 52](#) describes the main service on the NetWorker client.

Table 6 Services or programs on the NetWorker client

Service or program	Function
nsrexecd	<ul style="list-style-type: none"> • Authenticates and processes the NetWorker server remote execution requests and executes the save and savefs programs on the client.

Services and programs on the NetWorker storage node

[Table 7 on page 53](#) describes the main services or programs on the NetWorker Storage Node.

Table 7 Services or programs on the NetWorker storage node

Service	Function
nsrexecd	<ul style="list-style-type: none"> Authenticates and processes the NetWorker server remote execution requests and executes the save and savefs programs on the storage node.
nsrmmmd	<ul style="list-style-type: none"> Provides device support, generates mount requests, multiplexes save set data during a multiclient backup, and de-multiplexes recover data. It writes the data sent by save to storage media. Forwards storage information to nsrmmdbd for recording in the NetWorker server media database.
nsrsnmd	<ul style="list-style-type: none"> Provides an RPC-based service to manage all of the device operations that the nsrmmmd process handles on behalf of the nsrd process on the NetWorker server. Ensures that the necessary device operations are actually performed when needed by nsrd. Is automatically invoked by nsrd as required. Only one nsrsnmd service is run on each storage node with configured and enabled devices.
nsrlcpd	<ul style="list-style-type: none"> Provides a uniform library interface to the NetWorker media management daemon, nsrmmgd. Manages the library subsystem media, slot, drive and port resources providing control to move and access the resources within the library subsystems. One nsrlcpd is started for each configured jukebox.

Services and programs on the NetWorker Management Console server

[Table 8 on page 53](#) describes the main services or programs on the NetWorker Management Console (NMC) server.

Table 8 Services or programs on the NetWorker Management Console server

Service or program	Function
nsrexecd	Authenticates and processes the NetWorker server remote execution requests and executes the save and savefs programs on the client.
gstd	Known as the Generic Services Toolkit (GST), controls other services provided by the Console server.
httpd	Starts the NMC console on the client through a web browser.
dbsrv12	A database server that manages information pertaining to console management. For example, Console reports.
gstsnptrapd	<ul style="list-style-type: none"> 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. “Enabling or disabling the Capture Events option” on page 457 provide more information on how to configure SNMP Traps for a NMC managed Data Domain system.

Stopping and starting the Console server

Refer to the following sections for the steps to stop and start the NetWorker NMC server service:

- ◆ [“Stop the Console server on Windows” on page 54](#)
- ◆ [“Start the Console server on Windows” on page 54](#)
- ◆ [“Stop the Console server on UNIX” on page 54](#)
- ◆ [“Start the Console server on UNIX” on page 54](#)

Stop the Console server on Windows

To stop the Console server:

1. Log in as a Windows Administrator and right-click **My Computer** and select **Manage**.
2. Expand **Services and Applications** and select **Services**.
3. Right-click **EMC GST Service (gstd)**, then select **Stop**.

Start the Console server on Windows

To start the Console server:

1. Log in as a Windows Administrator and right-click **My Computer** and select **Manage**.
2. Expand **Services and Applications** and select **Services**.
3. Verify that the NetWorker client is running.

The NetWorker Remote Exec Service (**nsrexecd**) 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 (gstd)**, then select **Start**.

Stop the Console server on UNIX

To stop the Console server:

1. Log in as root.
2. Type one of the following commands:
 - Solaris and Linux: `/etc/init.d/gst stop`
 - AIX: `/etc/rc.gst stop`

Start the Console server on UNIX

To start the Console server:

1. Log in as root.
2. Verify that the NetWorker client is running.

For example, type the following command:

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

If the client is running, a message similar to this appears:

```
root240 1 0 ? 0:04 /usr/sbin/nsrexecd -s mysrvr
```

If the client is not running, start it. [“Start a NetWorker host on UNIX” on page 56](#) provides information about starting the client.

3. Start the Console server by typing one of the following commands:
 - Solaris and Linux: `/etc/init.d/gst start`
 - AIX: `/etc/rc.gst start`

Stopping and starting 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. In NetWorker 8.0 and later, new attributes have been introduced to configure a NetWorker server to not accept any new backup or recover sessions in preparation of a NetWorker daemon shutdown or server reboot. [“Restrict backup and recover access to the NetWorker server” on page 568](#) further information around how to prevent the NetWorker server from accepting new backup and recover sessions.

Refer to the following sections for the steps to stop and start the services:

- ◆ [“Stop a NetWorker host on Windows” on page 55](#)
- ◆ [“Start a NetWorker host on Windows” on page 55](#)
- ◆ [“Stop a NetWorker host on UNIX” on page 56](#)
- ◆ [“Start a NetWorker host on UNIX” on page 56](#)
- ◆ [“Stop a NetWorker host on Mac OS X” on page 56](#)
- ◆ [“Start a NetWorker host on Mac OS X” on page 56](#)

Stop a NetWorker host on Windows

To stop a host server, client, or storage node:

1. Log in as a Windows Administrator.
2. Right-click **My Computer** and select **Manage**.
3. Expand **Services and Applications** and select **Services**.
4. Right-click **NetWorker Remote Exec Service (nsrexecd)** and select **Stop**.

Start a NetWorker host on Windows

To start a host server, client, or storage node:

1. Log in as a Windows Administrator.
2. Right-click **My Computer** and select **Manage**.
3. Expand **Services and Applications** and select **Services**.
4. Start the appropriate service.
 - NetWorker server: Right-click the **NetWorker Backup and Recover Server** service (**nsrd**) and select **Start**.
 - NetWorker client or storage node: Right-click the **NetWorker Remote Exec Service** (**nsrexecd**) and select **Start**.

Stop a NetWorker host on UNIX

To stop the NetWorker services, log in as root and type the following command:

```
nsr_shutdown
```

Start a NetWorker host on UNIX

To start NetWorker services, log in as root and type the appropriate startup command listed in [Table 9 on page 56](#).

Table 9 NetWorker startup commands

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

Stop a NetWorker host on Mac OS X

To stop the NetWorker host:

1. Log in as a Mac Administrator.
2. Open the Mac OS X Terminal application utility.
3. Stop the service by typing the following command:

```
# SystemStarter stop NetWorker
```

Start a NetWorker host on Mac OS X

To start the NetWorker host:

1. Log in as a Mac Administrator.
2. Open the Mac OS X Terminal application utility.
3. Start the client by typing the following command:

```
# SystemStarter start NetWorker
```


CHAPTER 2

Backing Up Data

This chapter covers these topics:

◆ Scheduled backups.....	58
◆ Save sets	66
◆ Manual backups	70
◆ Verifying backed-up data.....	76
◆ Synthetic full backups.....	76
◆ Virtual synthetic full backups (for Data Domain systems).....	88
◆ Enable parallel save streams.....	93
◆ Probe-based backups	94
◆ Client Direct backups	95
◆ Checkpoint restart backups.....	95
◆ Deduplication backups	107
◆ Encrypting backup data.....	108
◆ Compressing backup data.....	109
◆ Special data handling for NetWorker clients on Windows	109
◆ Backing up Console server management data	110
◆ Backing up Windows mount points	112
◆ Backing up the Windows Content Index Server	114
◆ Backing up Windows DHCP and WINS databases.....	115
◆ Windows backup and recovery notes.....	116
◆ Customizing the backup command.....	119
◆ Considerations for backing up raw partitions.....	128
◆ Backing up a mapped drive.....	128
◆ Backing up access control lists.....	129
◆ Backing up BOOT/BCD Data on Windows	129
◆ Support for backing up renamed directories.....	129
◆ Backing up only client file indexes and the bootstrap	130

Scheduled backups

The NetWorker server backs up client data regularly by using scheduled backups. They are preferred over the [“Manual backups” on page 70](#) because the backups occur automatically, and data can be recovered more easily. You can also start scheduled backups at any time.

This section explains how to plan and create scheduled backups, including:

- ◆ [“Planning a scheduled backup” on page 58](#)
- ◆ [“Setting up a scheduled backup” on page 59](#)

Planning a scheduled backup

This section uses a scenario with the requirements of an accounting department to highlight tasks to consider when planning a scheduled backup.

Example 2 Planning scheduled backups for the accounting computers

Company XYZ wants to ensure that all of the computers in the Accounting department are backed up according to the requirements listed in [Table 10 on page 58](#). This table also maps each requirement to specific NetWorker features.

Table 10 Accounting department backup requirements

Requirement	NetWorker feature	More information
Backups occur at the same time.	Backup Schedule Backup Group	“Task 1: Set up a schedule for backups” on page 60 “Task 2: Set up a group for backup clients” on page 61
Accounting backups for the past 3 months are available immediately.	Browse Policy	“Task 3: Set up policies for quick access and long term storage” on page 62
Accounting backups for the past 7 years are available, though not necessarily immediately.	Retention Policy	“Task 3: Set up policies for quick access and long term storage” on page 62
Backups are routed to volumes that can be identified as Accounting backup volumes.	Label Template Pools	“Task 4: Set up a label template to identify volumes” on page 62 (if Match Bar Code Labels attribute is not used for the Library resource) “Task 7: Set up a pool to sort backup data” on page 66
To avoid unnecessary backups, do not back up files with a .tmp extension.	Directives	“Task 5: Set up directives for special processing” on page 63
The same files and folders are backed up on each accounting computer.	Client resource	“Task 6: Create a backup Client resource” on page 64
Non-accounting data need only be recoverable for one year.	Browse Policy Retention Policy Client resource	“Task 3: Set up policies for quick access and long term storage” on page 62 “Task 6: Create a backup Client resource” on page 64

Setting up a scheduled backup

You can create a scheduled backup quickly by using the Client Backup Configuration Wizard or through manual configuration in the Console.

Using the Client Backup Configuration Wizard

The Client Backup Configuration Wizard provides the ability to:

- ◆ Create Client resources for scheduled backups.
- ◆ Create Group resources.
- ◆ Add new clients to existing backup groups.
- ◆ Modify existing client configurations.

The wizard supports NetWorker servers and clients in a stand-alone or cluster environment. The NetWorker *Cluster Integration Guide* provides details about creating a client resource for virtual client backups.

NOTICE

The Client Backup Configuration wizard cannot be used to configure a NetWorker NDMP client or clients for NetWare.

Client Backup Configuration Wizard requirements

This section contains requirements or constraints specific to the use of the Client Backup Configuration Wizard.

- ◆ The wizard user must:
 - Have NetWorker server and client privileges, or have root (UNIX) or Administrator (Windows) privileges.
 - Have Configure NetWorker privileges on the NetWorker server where the scheduled backup is to be configured.
- ◆ The NetWorker server's host must be listed in the servers file on the client machine that is being configured for a scheduled backup.
- ◆ Communication between the Console server, NetWorker client host, and NetWorker server must use **nsrauth** strong authentication.
- ◆ The Console server, NetWorker client host, and NetWorker server must be using NetWorker 7.5 or later.
- ◆ Multiple wizard hosts cannot access the same client machine simultaneously.

[“Client wizard issues” on page 809](#) discusses known issues with the Client wizard.

Accessing the Client Backup Configuration Wizard

To access the Client Backup Configuration Wizard:

1. From the **Administration** window, click **Configuration**.
2. In the **Configuration** window, click **Clients**.

3. Add a client or modify an existing client:
 - To add a new client, select **Configuration** menu > **Client Backup Configuration** > **New**.
 - To modify an existing client, select the client and then select **Configuration** menu > **Client Backup Configuration** > **Modify**.

The wizard opens. If the wizard fails to open, ensure that all prerequisites in [“Client Backup Configuration Wizard requirements” on page 59](#) are met. Also check the NetWorker daemon log for additional details. [“Viewing log files” on page 803](#) provides more information.

Manually creating a scheduled backup in the Console

To exercise more control over scheduled backups than is possible by using the Client Backup Configuration wizard, complete these tasks:

- ◆ [“Task 1: Set up a schedule for backups” on page 60](#)
- ◆ [“Task 2: Set up a group for backup clients” on page 61](#)
- ◆ [“Task 3: Set up policies for quick access and long term storage” on page 62](#)
- ◆ [“Task 4: Set up a label template to identify volumes” on page 62](#)
- ◆ [“Task 5: Set up directives for special processing” on page 63](#)
- ◆ [“Task 6: Create a backup Client resource” on page 64](#)
- ◆ [“Task 7: Set up a pool to sort backup data” on page 66](#)

You do not have permissions to make configuration selections if the following error message appears while completing any task in this section:

```
user user_name needs to be on administrator's list
```

[“Managing server access” on page 558](#) provides information about getting permissions.

NOTICE

[Appendix F, “MAC OS X Support”](#) provides information about backing up NetWorker clients on Mac OS X

Task 1: Set up a schedule for backups

A schedule can be applied to each client backup. [Chapter 7, “Backup Groups and Schedules”](#) provides information about schedules

To create a schedule for backups:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Schedules**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a name for the schedule.

5. From the **Period** attribute, select **Week** or **Month**.
 - Select **Week** to create a weekly backup schedule. For example, if a full backup for a Friday is selected, every Friday will have a full backup.
 - Select **Month** to create a monthly schedule. For example, if a full backup for the first of the month is selected, every month will have a full backup on the first of the month.
6. Select a backup level for each day in the weekly or monthly period:
 - a. Select a day.
 - b. Right-click and from the **Set Level** menu, select a backup level.
7. If required, select an override backup level for any day. An override occurs once only for the selected day:
 - a. Select a day.
 - b. Right-click and from the **Override Level** menu, select a backup level.
8. Click **OK**.

Task 2: Set up a group for backup clients

A backup group specifies the time of day when a backup occurs. Creating groups for backup clients enables you to:

- ◆ Balance backup loads to reduce the impact on your network and the NetWorker server.
- ◆ Sort data to specific backup volumes. To sort data, groups are used in conjunction with backup *pools*.

[Chapter 7, “Backup Groups and Schedules”](#) provides information about groups.

To create a group:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Groups**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a name for the group.
5. In the **Comment** attribute, type a description of the group.
6. For the **Start Time** attribute, type a new time, unless it is appropriate to maintain the default time of 3:33 A.M. Ensure that start times for different groups are far enough apart so that one group has completed backing up before the next group starts.
7. For the **Autostart** attribute, select **Enabled**.
8. In the **Printer** attribute, type the name of the printer on which bootstrap save set information will be printed. For information about setting up a printer on Microsoft Windows systems, see [“Using nsrlpr to print notifications” on page 487](#).

9. Click the **Advanced** tab.
10. For the **Client Retries** attribute, change the number of retries, if necessary. This value specifies the number of times the NetWorker software attempts to back up a failed client.
11. Click **OK**.

Task 3: Set up policies for quick access and long term storage

Backup clients specify two policies: a browse policy and a retention policy.

- ◆ A browse policy determines how long backup data will be available for quick recovery.
- ◆ A retention policy determines how long backup data will be available for recovery, though not necessarily quickly. For example:
 - If it is likely that accounting data would need to be recovered within the past year, a browse policy of one year would be appropriate.
 - If the same accounting data had to be recoverable for up to seven years even though the likelihood of needing to recover it was low, a retention policy of seven years would be appropriate.

[“About browse and retention policies” on page 276](#) provides information about browse and retention policies.

To create a policy:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Policies**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a name for the policy. Choose a name that reflects the length of time for which the policy specifies. For example, if the policy is for 15 months, type 15 months.
5. In the **Comment** attribute, type a comment about the policy.
6. In the **Number of Periods** attribute, type the number of periods applied to the policy. For example, if you choose months for the **Period** attribute and 3 for the **Number of Periods** attribute, then the policy lasts for 3 months (one quarter).
7. From the **Period** list, select a **period**. This attribute works in conjunction with the **Number of Periods** attribute. For example, a quarterly policy is configured in terms of the number of months (3). A week as seven days beginning on Sunday, a month is the calendar month, and a year is the calendar year.
8. Click **OK**.

Task 4: Set 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. [Chapter 10, “Sorting Backup Data”](#) provides more information about label templates.

To create a label template:

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**.

Task 5: Set up directives for special processing

Directives are optional instructions that control how files and directories are processed during backup and recovery. For instance, one could use a directive to skip all temporary files (*.tmp) during backup.

Other common uses for directives include adding password-protection and data compression to scheduled backups. [Chapter 9, “Directives”](#) provides information about directives.

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. [“Preconfigured global directive resources” on page 294](#) provides more information.

To create a directive:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Directives**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a name for the directive.

5. In the **Comment** attribute, type a description for the directive.
6. In the **Directive** attribute, type the directive instructions. For example, to skip all files on C:\ that have a .tmp extension, type:

```
<< "C:\" >>
skip: *
```

7. Click **OK**.

Task 6: Create a backup Client resource

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 included in a scheduled backup. A Client resource also specifies information about the backup, such as the backup schedule, the backup group, browse policies, and retention policies.

A single NetWorker client computer can have several Client resources, although clients with the same save set cannot be in the same group. For instance, suppose the accounting data on a computer should be backed up according to a different schedule than the operating system files on the same computer. To accomplish this, one could create two Client resources on each computer: one for accounting data and another for operating system data.

Another common reason to create multiple Client resources for the same computer is to back up large client file systems more efficiently. For instance, one could create two Client resources: one for each file system on a computer. Each Client resource could be scheduled to back up separately.

[“Multiple clients from the same computer” on page 622](#) provides information about multiple Client resources.

To create a Client resource:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Clients**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type the hostname of the NetWorker client computer.
5. In the **Comment** attribute, 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. From the **Browse Policy** attribute, select a browse policy from the list. The browse policy determines how long backed-up data is available for quick access.
7. From the **Retention Policy** attribute, select a retention policy from the list. The retention policy determines how long backed-up data is available though not necessarily quickly.
8. Select the **Scheduled Backups** attribute.
9. The **Client Direct** attribute, which is selected by default, enables the client to send backup data directly to the storage device, bypassing the storage node.

[“Client Direct backups” on page 95](#) provides details.

10. Select the **Block based backup** attribute to perform backups at the block level.

[“Block Based Backup and Recovery” on page 627](#) provides information about performing block based backup and recovery.

11. From the **Directive** attribute, select a directive from the list, if applicable.
12. In the **Save Set** attribute, type the name of the files or directories to be backed up. Place multiple entries on separate lines. For example, to back up a log file directory named C:\log and all of the data under the directory named D:\accounting, the entries would look similar to:

```
C:\log
D:\accounting
```

Type **ALL** to back up *all* client data. For Microsoft Windows operating systems, the WINDOWS ROLES AND FEATURES save sets that determine the client system’s state should be backed up on a regular basis.

More information can be found in the following locations:

- [“Save sets” on page 66](#) provides information on using the ALL save set and the SYSTEM, VSS SYSTEM, or WINDOWS ROLES AND FEATURES save sets.
- [“The ALL save set” on page 68](#) describes how the components of the ALL save set can differ significantly between the various supported Windows operating systems.
- [“Backing up a mapped drive” on page 128](#) provides information on backing up a mapped drive.
- [“Backing up Windows mount points” on page 112](#) provides information on backing up mount points and nested mount points.
- [“Scheduled backups of non-ASCII files or directories” on page 623](#) provides information on using non-English paths in the **Save Set** attribute.

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. [“Preconfigured global directive resources” on page 294](#) provides information.

13. From the **Group** attribute, select a group from the list.
14. From the **Schedule** attribute, select a schedule from the list.
15. Select the **Backup renamed directories** attribute to back up the files and subdirectories of a renamed directory even if only the name of the directory has changed.

If this attribute is selected, and a directory is renamed, all files and subdirectories under that directory will be backed up during the next scheduled full or non-full backup. [“Support for backing up renamed directories” on page 129](#) provides more information about this feature.
16. Select **Globals (2 of 2)**, in the **Owner notification** attribute, specify the command to send a backup completion email to email recipients. [“Owner notifications” on page 492](#) describes how to configure **Owner notifications**.
17. Click **OK**. The client is now set up for scheduled backups.

To determine whether a client is enabled for scheduled backups, locate the client entry in the right pane and look for a check mark under the Scheduled backup column.

Task 7: Set up a pool to sort backup data

A backup pool is a collection of volumes to which backup data is written. Use pools to sort backup volumes so that the volumes are easy to locate when they are needed. [Chapter 10, “Sorting Backup Data”](#) provides more information about pools.

To create a backup pool:

1. From the **Administration** window, click **Media**.
2. In the expanded left pane, select **Media Pools**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a name for the pool. A pool is associated with a label template. Use a name that clearly associates the pool with the corresponding label template.
5. In the **Comment** attribute, type a description of the pool.
6. Select the **Enabled** attribute.
7. For the **Pool Type** attribute, select **Backup**.
8. In the **Label Template** attribute, select the matching label template.
9. Modify the attribute to use to direct specific backup data to the volumes belonging to this pool. Data can be sorted by group, backup clients, save sets, and backup levels. [“Using media pools” on page 304](#) provides more information about sorting criteria.
10. Click **OK**.

Save sets

A Client resource identifies the client data to be backed up. The collection of data items backed up during a backup session between the NetWorker server and the 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.

NOTICE

A save set is defined when a Client resource is created. [“Task 6: Create a backup Client resource” on page 64](#) provides information about creating a Client resource.

Scheduling predefined save sets for backup

In addition to entering files or file systems in the Save Set attribute of the Client resource, you can also type the names of predefined save sets when configuring the NetWorker client.

For Microsoft Windows Server 2003 that has VSS disabled, and for Windows XP Professional, these predefined save sets are available:

- ◆ ALL
- ◆ SYSTEM STATE:
- ◆ SYSTEM DB:
- ◆ SYSTEM FILES:

For Windows Server 2003 with VSS enabled (default setting), and for Windows Vista, these predefined save sets are available:

- ◆ ALL
- ◆ VSS USER DATA: (Windows Server 2003 only)
- ◆ VSS OTHER: (Windows Server 2003 only)

For Windows Server 2008, Windows Server 2008 R2, Windows 7, Windows Server 2012, Windows 8, Windows Server 2012 R2 and Windows 8.1 these predefined save sets are available:

- ◆ DISASTER_RECOVERY:\

Consider the following:

- ◆ The DISASTER_RECOVERY:\ save set backs up the system state critical volumes for an operating system, and except for some applications such as clusters, Active Directory, DFS-R, and Windows Server Failover Cluster, can completely replace the VSS SYSTEM savesets.
- ◆ Checkpoint restart backup for the Windows DISASTER_RECOVERY:\ save set is not supported. If a Client with a DISASTER_RECOVERY:\ save set is checkpoint-enabled, it is quietly ignored for the Disaster Recovery save sets. The save set will be marked with a cb flag instead of a k flag, indicating that the checkpoint is not considered for Disaster Recovery.
- ◆ The save set ALL includes all predefined save sets, including the DISASTER_RECOVERY:\ save set. See the preceding bullet regarding checkpoint-enabled backups and the DISASTER_RECOVERY:\ save set.
- ◆ VSS writer files Windows Server 2003, and associated with the system state or with applications, are skipped during the regular file system backup. These files are backed up after the regular file system backup. Files that are associated with the system state are backed up under the VSS SYSTEM savesets. Files that are associated with applications are backed up under their corresponding VSS savesets. When performing a recovery, system state files can be found under the VSS system savesets, which include the VSS SYSTEM BOOT:, VSS SYSTEM SERVICES:, and VSS SYSTEM FILESET: savesets.
- ◆ To properly protect NetWorker client computers, all of the SYSTEM or VSS SYSTEM save sets must be backed up and recovered simultaneously. Failure to do so will yield unpredictable results.

For information about:

- ◆ SYSTEM, VSS SYSTEM, SHAREPOINT, and ASR save sets, see [Appendix A, “SYSTEM, VSS SYSTEM, and WINDOWS ROLES AND FEATURES Save Sets.”](#)

- ◆ Support for ASR, see [Chapter 24, “Recovery Support for Windows XP and 2003 Automated System Recovery.”](#)
- ◆ Support for Disaster Recovery on Windows Server 2008. Windows Server 2008 R2, Windows 7, Windows Server 2012, Windows 8, Windows Server 2012 R2 and Windows 8.1 see [Chapter 25, “Windows Bare Metal Recovery.”](#)

The ALL save set

The ALL save set is the default save set used when a client is created: [Table 11 on page 68](#) provides a list of the save sets that are included in the ALL save set.

Table 11 Components in the ALL save set (1 of 2)

Operating system	Files
UNIX	<ul style="list-style-type: none"> • When the backup starts, the savefs process will read 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. Only currently mounted file systems that are configured in the files mentioned above are backed up. • 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 filesystems are backed up. • If the save set name includes a symbolic link, a save set recovery is not supported.
Windows 2003	<p>DISASTER_RECOVERY:\ (included in Full backup only) VSS SYSTEM BOOT VSS SYSTEM FILESET VSS SYSTEM SERVICES All local physical drives</p> <p>You can revert the definition of the ALL save set to be equivalent to a Windows 2003 host with VSS enabled by using the VSS:DISASTER_RECOVERY=off keyword in the Save Operations attribute. Chapter 25, “Windows Bare Metal Recovery,” provides more information about using the ALL save set with Windows 2008, 2008 R2, Windows 7, Windows Server 2012, Windows 8.1 and Windows Server 2012 R2.</p> <p>Notice: When you use the ALL save set with synthetic full and virtual synthetic full backups, the non-critical volumes save successfully. However, critical volumes including DISASTER_RECOVERY:\ are not backed up. The nsrconsolidate() command is able to perform a synthetic full for regular volumes but not critical volumes. The client then runs a level full backup for the DISASTER_RECOVERY:\ volume.</p>

Table 11 Components in the ALL save set (2 of 2)

Operating system	Files
Windows Server 2003 with VSS enabled (default setting)	VSS SYSTEM BOOT VSS SYSTEM FILESET VSS SYSTEM SERVICES VSS USER DATA (Windows Server 2003 only) VSS OTHER (Windows Server 2003 only) All local, physical drives
Windows XP Professional, Windows Server 2003 with VSS disabled	SYSTEM STATE SYSTEM DB SYSTEM FILES SHAREPOINT All local, physical drives
Windows Server 2008, Windows Server 2008 R2, Windows 7, Windows Server 2012, Windows 8, Windows Server 2012 R2, Windows 8.1	WINDOWS ROLES AND FEATURES
Mac OS X	All local and mounted volumes

File system objects not included in the ALL save set

The ALL save set excludes certain named directories, file systems, and files. A list of the file system object names that the ALL save set does not back up includes:

- hsfs
- proc
- fd
- cachefs
- lofs
- mntfs
- ctfs
- objfs
- sharefs
- nfs2
- nfs3
- nfs3perf
- profs
- nfs4
- nfs
- brfs
- msdos
- dfs
- autofs
- iso9060
- udf
- sysfs
- debugfs
- subfs
- usbdevfs
- binfmt_misc
- usbfs
- devpts
- smbfs
- cifs
- swap
- tmp
- tmpfs
- nucfs
- nucam
- fdfs
- xx
- none

[“Using the save set ALL to back up particular file systems” on page 69](#) describes how to use the special save set ALL syntax to back up these file system objects.

Using the save set ALL to back up particular file systems

Use special keywords with the save set ALL to define the file systems to include in a client backup. [Table 12 on page 70](#) provides a list of the special ALL save sets and the backup behavior.

NOTICE

When the All save set is used for a back up, the Networker software creates a temporary file similar to a directive under each drive. The file is named by using the format, *<drive guid>.txt*. The file lists the files that are excluded from the backup. The file is temporary and is automatically deleted when the backup completes.

Table 12 Special ALL save sets

Special ALL save set syntax	Backup behavior
<i>all-filesystem</i>	<ul style="list-style-type: none"> • Only backup locally mounted file systems of a particular type. For example, the all-zfs save set backs up all locally mounted zfs file systems on a Solaris host. • Other save sets include: all-ntfs and all-ext3. • File systems that are normally skipped such as NFS are still skipped. • The <i>NetWorker Software Compatibility Guide</i> provides a list of the supported file system for each operating system.
all-mounts	<ul style="list-style-type: none"> • On UNIX clients, backup all of the currently mounted file systems. File systems that are normally skipped such as NFS are still skipped. • On Windows clients, the all-mounts save set is equivalent to the save set ALL.
all-local	<ul style="list-style-type: none"> • 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 save set is equivalent to the save set ALL.
all-global	<ul style="list-style-type: none"> • 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 save set is equivalent to the save set ALL.

Manual backups

Manual backups enable users to make quick backups of a few files. Unlike scheduled backups, manual backups do not:

- ◆ Generate bootstrap files
- ◆ Back up indexes

This may present recovery problems if the indexes are recovered after a disaster, but before a scheduled backup has backed up the latest indexes. For this reason, scheduled backups are the preferred backup method. However, indexes can be saved manually by using the **savegrp** program. [“Performing a manual backup from the command prompt” on page 73](#) provides more information.

On Microsoft Windows, manual backups can be performed by using the graphical NetWorker User program. On UNIX and Linux, manual backups can be performed from only the command line.

NOTICE

You can also start a scheduled backup manually. [“Starting a group immediately” on page 466](#) provides information about starting a scheduled backup group manually.

Performing a manual backup on Windows

NOTICE

If performing a NetWorker User backup on a NetWorker server, see [“Excluding file type devices from a manual backup on Windows” on page 71](#).

The NetWorker User program cannot be used to back up deduplication data. Deduplication data must be backed up by using scheduled backups or from the command line.

To start a manual backup on Windows:

1. In the NetWorker **User** program, click **Backup**. [Chapter 1, “Overview”](#) provides general information about the NetWorker **User** program

NOTICE

There are considerations to be aware of when performing a manual backup of SYSTEM or VSS SYSTEM save sets. [“Manual backups of the SYSTEM and VSS SYSTEM save sets” on page 72](#) provides more information.

2. In the left pane of the **Backup** window, click the appropriate directory folder.
3. Select each directory or file, and click **Mark**. To clear an item, click **Unmark**.
4. Click **Start** to begin the manual backup. The **Backup Status** dialog box displays the progress of the backup.

When the backup finishes, a message similar to this appears:

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

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.

NOTICE

Certain kinds 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.

Excluding file type devices from a manual backup on Windows

When performing a NetWorker User backup on a NetWorker server or storage node that is backing up to a local file type device, do *not* include the local file type device in the backup. If the local file type devices are included, the backup file will grow until there is no more disk space. The following procedure must be performed before selecting any files for backup or archiving, or before performing any activities from the Operation menu of the NetWorker User program.

To ensure that file type devices are excluded from NetWorker User backups, create a local directive on the NetWorker server as follows:

1. Start the NetWorker **User** program.
2. From the **Options** menu, select **Local Backup Directives**.
3. Click the filename of any file device to unmark it.
4. From the **File** menu, select **Save Directive**. This creates a directive file named `networkr.cfg`. [“Local directives within the NetWorker User program” on page 293](#) provides more information about the `networkr.cfg` file. [“File type devices” on page 166](#) provides information about file type devices.

Manual backups of the SYSTEM and VSS SYSTEM save sets

This section discusses manual backups of the SYSTEM or VSS SYSTEM save sets. These save sets are used to back up Windows system files. [Appendix A, “SYSTEM, VSS SYSTEM, and WINDOWS ROLES AND FEATURES Save Sets”](#) provides more information.

NOTICE

To back up and recover SYSTEM and VSS SYSTEM save sets, you must have local Windows Administrator privileges.

Manual backups of the SYSTEM and VSS SYSTEM save sets from the NetWorker User Program

In the NetWorker User Backup window, each of the SYSTEM or VSS SYSTEM save sets appear as a distinct node in the left pane. Expanding any of these nodes reveals its components in the right pane, as shown in [Figure 7 on page 72](#) and [Figure 8 on page 73](#).

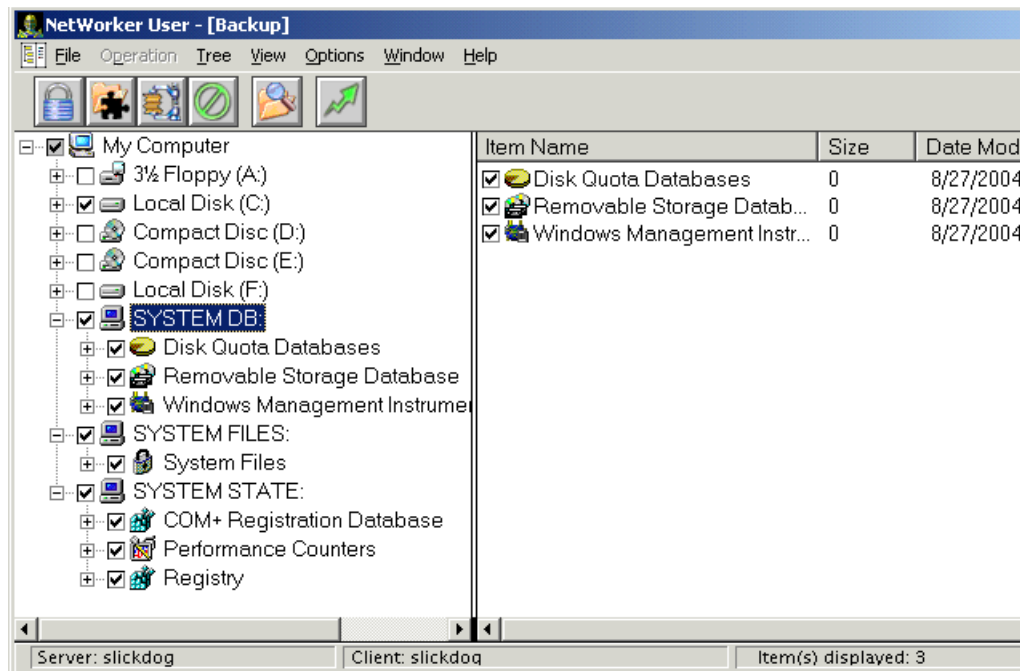


Figure 7 SYSTEM save sets in the NetWorker User backup window

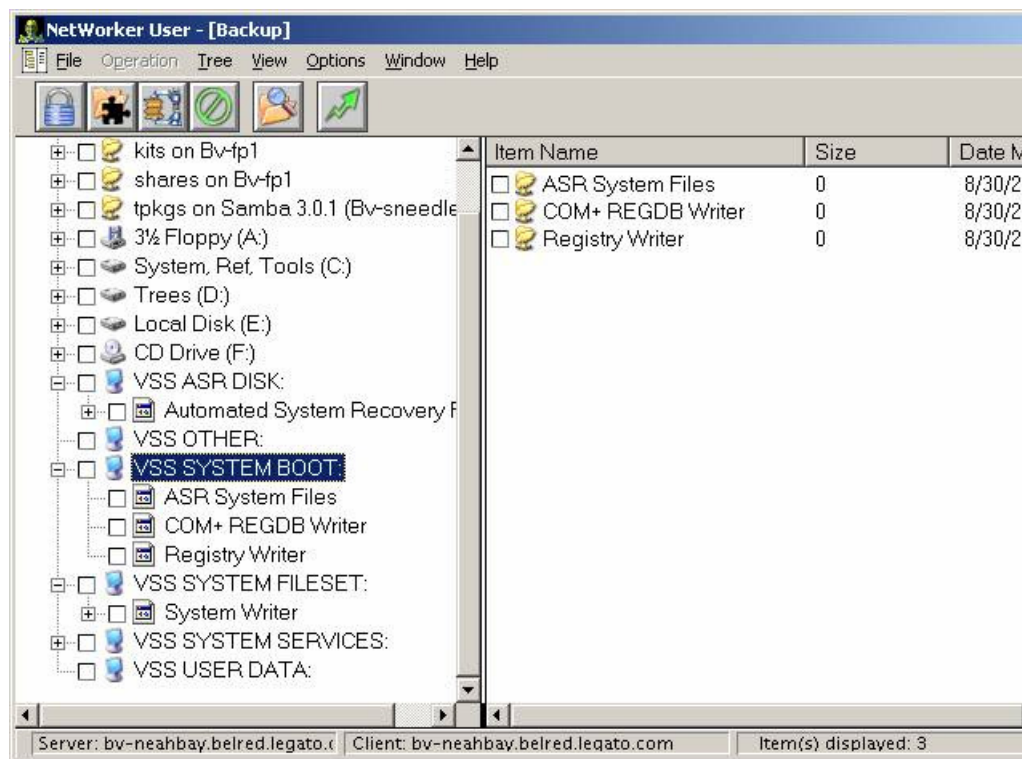


Figure 8 VSS SYSTEM save sets in the NetWorker User backup window

The nodes that appear in the right pane reflect save set components that were eligible for backup at the time the NetWorker User program was started. After the NetWorker User Backup window is opened, the display does not dynamically refresh to reflect save set components that have since become eligible or ineligible for backup. However, all the eligible save set components are included in a backup when the backup operation starts, including those components that become available after the NetWorker User program starts.

NOTICE

The NetWorker User program's special handling features (password-protect, encrypt, or compress) cannot be used when any of the SYSTEM or VSS SYSTEM save sets are marked for backup.

Performing a manual backup from the command prompt

A manual backup can also be performed from the command prompt by using the **save** command. For example, to back up myfile to the server *jupiter*, type:

```
save -s jupiter myfile
```

If you do not specify the **-s** option with the **save** command, the file(s) will be backed up to the NetWorker server defined in the `/nsr/res/servers` file that comes first in alphabetical order.

NOTICE

You can also manually back up the bootstrap and indexes for a group by using the **savegrp** program with the **-O** option and a group name. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provides information about **save** and **savegrp**.

Manual backups of the SYSTEM and VSS SYSTEM save sets

A manual backup of the SYSTEM and VSS SYSTEM save sets can also be performed from the command prompt.

With VSS disabled

This section describes how to back up SYSTEM save sets on a NetWorker client that is running:

- ◆ Windows Server 2003 with VSS disabled
- ◆ Windows XP Professional

To back up all components of the Windows system state:

```
save [-s NetWorker_server_name] "SYSTEM STATE:"
save [-s NetWorker_server_name] "SYSTEM FILES:"
```

To back up all components of the SYSTEM DB, SHAREPOINT, and ASR save sets:

```
save [-s NetWorker_server_name] "SYSTEM DB:"
save [-s NetWorker_server_name] "SHAREPOINT:"
save [-s NetWorker_server_name] "ASR:"
```

With VSS enabled (default setting)

This section describes how to back up VSS SYSTEM save sets on a NetWorker client that is running Microsoft Windows Server 2003 with VSS enabled:

To back up all components of the Windows system state:

```
save [-s NetWorker_server_name] "VSS SYSTEM BOOT:"
save [-s NetWorker_server_name] "VSS SYSTEM FILESET:"
```

To back up all components of the VSS SYSTEM SERVICES, VSS USER DATA, VSS OTHER, and VSS ASR DISK (Windows 2003 only) save sets:

```
save [-s NetWorker_server_name] "VSS SYSTEM SERVICES:"
save [-s NetWorker_server_name] "VSS USER DATA:"
save [-s NetWorker_server_name] "VSS OTHER:"
save [-s NetWorker_server_name] "VSS ASR DISK:"
```

Requirements and limitations

When backing up SYSTEM or VSS SYSTEM save sets from the command line, these requirements and limitations apply:

- ◆ Do not select individual components of any of the SYSTEM or VSS SYSTEM save sets for backup.
- ◆ A maximum of one SYSTEM or VSS SYSTEM save set can be included in the same **save** command.

- ◆ File system directories cannot be specified in the same **save** command.
- ◆ A maximum of one SYSTEM or the VSS SYSTEM save set can be specified in an input file.

NOTICE

An input file is specified in a **save** command with the **-I** option.

File system directories cannot be specified in an input file.

Examples of invalid command line entries include:

```
save -s servername "SYSTEM DB:" "SYSTEM STATE:"
save -s servername D:\letters "SYSTEM DB:"
save -s servername -I D:\list.txt
```

where **list.txt** is an input file. Examples of invalid input files include:

- ◆ The following input file is invalid because it includes a file system and a VSS SYSTEM save set:
D:\letters
VSS SYSTEM BOOT:
- ◆ The following input value file is invalid because it includes multiple VSS SYSTEM save sets:
VSS SYSTEM BOOT:
VSS SYSTEM SERVICES:

Examples of valid command line entries include:

```
save -s servername "VSS SYSTEM BOOT:"
save -s servername "VSS SYSTEM SERVICES:"
```

NOTICE

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 additional information. VSS backup error messages are also written to the log file (*networkr.raw*) on the local client.

Backing up multiple SYSTEM save sets

To back up multiple SYSTEM or VSS SYSTEM save sets in one operation, choose one of the following options:

- ◆ In the NetWorker Administration window, edit the Client resource to include multiple SYSTEM or VSS SYSTEM save sets. Alternatively, select the default save set ALL, which will include all SYSTEM and VSS SYSTEM save sets.
- ◆ In the NetWorker User program, mark all of the required SYSTEM or VSS SYSTEM save sets as well as any other required save sets, and then complete the backup.

Verifying backed-up data

NOTICE

This feature is not available on UNIX clients.

Use the NetWorker Verify feature to ensure that backup data on the NetWorker server matches the data on the local disk. The Verify feature compares the file types, file modification times, file sizes, and file contents. It does *not* verify other system attributes, such as read-only, archive, hidden, system, compressed, and file access control list (ACL). The NetWorker server alerts you to any changes that have occurred to your data since the backup. Verification also determines whether a hardware failure kept the NetWorker server from completing a successful backup. The Verify feature provides a way to test the ability to recover data.

NOTICE

To verify files, you must have Windows Administrator privileges for the computer.

To verify backup data:

1. In the NetWorker **User** program, select **Verify Files** from the **Operation** menu.
2. Select the data items to be verified.
3. Click **Start**.

You can monitor the data verification progress in the Verify Files Status window. After the verification is complete, the Verify Status dialog box shows any data discrepancies.

Synthetic full backups

A synthetic full backup combines a full backup and subsequent incremental backups to form a new full backup which is called a synthetic full backup. A synthetic full is equivalent to a traditional full backup and can be used in all of the same ways as a traditional full backup. Although the synthetic full backup method is quite different from the traditional method of creating a full backup, the result is exactly the same.

The synthetic full save set includes data that was backed up between the full backup and the last incremental backup. After a synthetic full backup is performed, the next synthetic full backup combines the previous synthetic full backup and subsequent incremental backups.

The full backups and incremental backups must be created with NetWorker 8.0 and later. Synthetic full backups cannot be created using full, level, or incremental backups that were created with NetWorker versions prior to NetWorker 8.0. This feature supports filesystem backup only. [“Limitations of synthetic full backups” on page 80](#) provides more details.

Using synthetic full backups 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. The save sets that result from a synthetic full backup are equivalent

to a traditional full backup of these same clients as of the time of the last incremental backup that was used in the creation of the synthetic full backup. Each synthetic full backup forms the basis upon which the next synthetic full is created.

Synthetic full backup is supported on Avamar deduplication nodes by using the **nsrconsolidate** command. The synthetic full backup rehydrates (reverts) the deduplicated data to its original non-deduplicated state. [“Performing manual synthetic full backups using nsrconsolidate” on page 85](#) provides details.

NOTICE

NetWorker 8.0 and later does not support the consolidate backup level. During an upgrade from a previous release of the NetWorker server software to version 8.0 or later, the update process changes the value in the Level attribute from Consolidate to NULL.

NetWorker 8.1 SP1 and later does not support the synth_full backup level in the Level attribute of a Group resource. During an upgrade from a previous release of the NetWorker server software to version 8.1 SP1 or later, the update process changes the value in the Level attribute from synth_full to NULL. To perform a synthetic full backup, in the Level attribute of the Group resource, select incr_synth_full. The synth_full backup level is not displayed in the Level dropdown list in the Group resource.

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.

Differences between a synthetic full backup and a traditional backup

When traditional full backups are performed, data from the clients is sent over the network to the storage nodes. In some cases, this can have a negative effect on client network performance.

In contrast, to create a synthetic full, the NetWorker software:

1. Analyzes an existing full backup along with subsequent incremental backups.
2. Extracts from each of these the most current versions of files in the backup set.
3. Streams them into a new full backup.

Synthesizing the new full backup does not involve the client machines and localizes the network traffic to the NetWorker server and storage nodes.

NOTICE

When you use the ALL save set with synthetic full and virtual synthetic full backups, the non-critical volumes save successfully. However, critical volumes including DISASTER_RECOVERY:\ are not backed up. The nsrconsolidate() command is unable to process the DISASTER_RECOVERY:\ saveset. The client then runs a level full backup for the DISASTER_RECOVERY:\ saveset. More information on the ALL save set is available in [“The ALL save set” on page 68](#).

When to use synthetic full backups

Synthetic full backups can be used on any eligible client. However, synthetic full backups, provide the most benefit in the following cases:

- ◆ If 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 small.
- ◆ Large backups over the network are cost-prohibitive.

Synthetic full backups involve only the NetWorker server and storage node. If all of the data is located on a few storage nodes, then the network overhead for creating the synthetic full 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 than a full backup because incremental backups are combined into a synthetic full backups. Used indiscriminately, synthetic full backups might impact the performance of the storage node.

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.

How a synthetic full backup is created

Figure 9 on page 78 illustrates how a synthetic full backup is created.




Figure 9 Synthetic full backups

In this example, 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 new full backup which is called 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 can not be assigned the exact same save set time. For example, if T4 is created at 1334389404, then T4+Delta is created at 1334389405, a difference of one second. The synthetic full save set will only include files covered by save sets up to T4 at 1334389404. The incremental backup following the synthetic full backup

at 1334389405 will include 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 given time.

The synthetic full save set includes data that was backed up between the full backup and the last incremental backup of the client from a point-in-time after the last successful incremental backup in the chain that was recorded in the media database. After a synthetic full backup is performed, the next synthetic full backup combines the previous synthetic full backup and subsequent incremental backups.

Note that the synthetic full backup is based upon the last successful incremental backup in the chain of incremental backups. If T3 and T4 are removed from the chain of incremental backups, then the synthetic full backup will be taken at T2; not T4.

Synthetic full requirements

Before performing a synthetic full backup, ensure that the following requirements are met:

- ◆ A full or synthetic full backup was created with NetWorker 8.0 or later before subsequent synthetic full backups are run. If a full backup does not exist, the synthetic full backup will fail with the following error:

```
Synthetic full operation must include one full save
```

To resolve this issue, perform a full backup and incremental backups of the save sets which will be participating in synthetic full backup.

- ◆ All incremental backups participating in the synthetic full backup are in the media database. If one of the incremental backups are missing, the synthetic full backup might fail with the following errors:

```
nsrconsolidate: info, Anchor saveset time time (machine_name:path)
not found in media database.
```

```
savegrp: Check criteria for machine_name:path returned: Anchor
saveset time time (machine_name:path) not found in media database.
```

To resolve this issue, perform an incremental synthetic full level backup. This backup level creates an incremental backup of the client from a point-in-time after the last successful incremental backup in the chain of incremental backups that was recorded in the media database. A synthetic full backup is then automatically run.

- ◆ All save sets participating in the construction of a synthetic full save set must:
 - Retain the same client name during the incremental and full backups that combine to form the synthetic full backup.
 - Retain the same save set name during the incremental and full backups that combine to form the synthetic full backup.
 - Be browsable in the online index. If one of the save sets to be used in the synthetic full backup is no longer browsable, the synthetic full backup will fail with the following error:

```
Save set saveset invalid for consolidate: no longer browsable.
```

- ◆ Have the **backup renamed directories** attribute enabled (default setting) for all clients that participate in the synthetic full backup. If the **backup renamed directories** attribute is not enabled during any of the full or incremental backups that participate in the synthetic full backup, the synthetic full backup might fail with the following error message:

```
Save set saveset invalid for consolidate: backup renamed directory
index lookup information missing from index
```

To resolve this issue:

1. Enable **backup renamed directories** for all clients that are participating in the synthetic full backup. The Backup renamed directories attribute is found in the Client resource.
 2. Perform a full backup.
 3. Perform at least one incremental backup.
 4. Perform a synthetic full backup.
- ◆ Have a client resource created for the NetWorker storage node that will be used for the synthetic full backup. A client connection license for this storage node is not used if the storage node is not backed up.
 - ◆ Have two available attached devices; one for reading and one for writing. Devices that support concurrent read/write access can be shared for reading and writing, if the pool restrictions and session limits allow for the use of the device.

Consider the following:

- Use Advanced File Devices and DataDomain devices to store all of your backups to a single device.
- A synthetic full backup can not be saved to a file device or tape volume that contains backups which will be used to create the synthetic full backups.
- ◆ Use the following criteria to determine the recovery storage node:
 - If the required volume is not mounted, the recovery storage node is selected based on the setting in the client's recovery storage node attribute.
 - If the required volume is already mounted, the storage node where the volume is mounted will be selected for recovering the data.

Limitations of synthetic full backups

The synthetic full operation is resource intensive. In order to manage resource usage, it is best practice to perform synthetic full operations outside of the normal backup window.

Do not perform a synthetic full backup, if:

- ◆ The backup type is VSS.
- ◆ The backup type is NDMP, SCSI, or VCB.

- ◆ The save set belongs to a snapshot group.
- ◆ The save sets contain backups of raw disk file partitions.
- ◆ The save sets contain database systems such as Microsoft Exchange and Oracle.
- ◆ The backup command with save is not used.
- ◆ For UNIX clients, include the forward slash to designate root "/" when specifying a save set name. Otherwise, the synthetic full backup will fail. For example, if /tmp is misspelled as tmp in the save set list of the Client resource, the backup will fail.
- ◆ For Windows clients, include the backslash "\ " when specifying a drive letter in a save set name. Otherwise, the synthetic full backup will fail. For example, if D: is typed in the save set list of the Client resource, instead of D:\, the synthetic full backup will fail.

“[Synthetic full backups](#)” on page 728 describe how the synthetic full backup feature works with Windows offline disaster recovery.

NOTICE

Synthetic full backups can free network bandwidth and Client resources by reducing the backup window. However, because incremental backups are combined into a synthetic full, a synthetic full backup might take longer to run on the storage node than a comparable full backup. Without proper planning, synthetic full backups might impact the performance of the storage node.

A synthetic full operation is resource intensive. In order to manage resource usage, it is best practice to perform synthetic full operations outside of the normal backup window.

Recommended devices for synthetic full backups

Although synthetic full backups can be directed to any device that can be used in a traditional full backup, there are special considerations to account for when conducting synthetic full backups. Since synthetic full backups involve concurrent recover and save operations, it is strongly recommended that the synthetic full be directed to devices which 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.

Other devices such as tape drives, VTLs, and basic file devices can be used as the destination for synthetic full backups but careful preparation is required if the backup is to succeed. The backup must be configured so that the destination volume does not contain any of the source save sets that are used for the synthetic full backup. Also, for tape media, ensure that there is 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.

Synthetic full scheduling considerations

A synthetic full operation is resource intensive because it concurrently performs both recover and save operations. As a result, perform synthetic full operations outside of the regular backup window. You can do this by creating separate groups for synthetic full operations. 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 the current backup schedule performs a full backup once a week, the full backup can be replaced with an incremental backup followed by a synthetic full backup without increasing the backup space usage.

A suggested change for using synthetic full backup for the following schedule is:

```
Current schedule:    Sunday : Full,    Mon-Sat  : Incr
New Schedule: First Sunday : Full, Mon-Sat : Incr, 2nd - 5th Sun : incr
                  followed by synth_full.
```

Synthetic full and backup levels

NetWorker 8.0, NetWorker 8.0 SP1, and NetWorker 8.1 support two backup levels:

- ◆ synth_full, defined in a Schedule resource.
- ◆ incr_synth_full, defined in Level attribute of a Group or the Schedule resource.

Prior to NetWorker 8.0, performing a backup at a particular level created a backup of that same level. However, synthetic full backups do not follow this convention. Performing a synthetic full backup creates a full level backup. The [“Schedules” on page 260](#) and [“Backup levels” on page 267](#) sections provide more information.

In NetWorker 8.1 SP1 and later, synth_full is no longer an available attribute in the Group Properties dialog. If you want to create a synth_full backup you can create one outside of the Group backup window by configuring a backup with a different schedule.

[Table 13 on page 82](#) shows the level backups that are created by synthetic full and incremental synthetic full backups.

Table 13 Synthetic full backup levels

Backup level	Reported backup level
synth_full	full
incr_synth_full	incr full

Incremental Synthetic Full

Synthetic full backups only cover the period in time up to the last incremental backup that was used in the synthesizing process. To backup data that has changed since that last incremental backup, you must perform an incremental synthetic full backup. During an incremental synthetic full backup, 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 and then performs the synthetic full backup.

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.

These sections provide more details:

- ◆ [“Configuring synthetic full backups from the Administration window” on page 83](#)
- ◆ [“Performing manual incremental synthetic full backup by using the savegrp program” on page 85](#)
- ◆ [“Performing manual synthetic full backups using nsrconsolidate” on page 85](#)

Configuring synthetic full backups from the Administration window

To configure a scheduled synthetic full backup:

1. Ensure the NetWorker server, storage node, and clients are at NetWorker 8.0 or later.
2. For each storage node that participates the synthetic full backup, ensure that a client resource has been created for that storage node.
3. For each NetWorker client that participates in the synthetic full backup:
 - a. Ensure that the **Backup renamed directories** attribute is enabled.

NOTICE

New client instances created with NetWorker 8.0 have the **Backup renamed directories** attribute enabled by default.

- b. In the **Save Set** attribute, type the name of the files or directories to be backed up:
 - For the Microsoft Windows operating systems, even though most file systems are case-independent, you must use the same path name case that the Windows file system uses. The NetWorker software’s cross-platform indexing system is case-sensitive. It is best practice to always specify the Windows drive letter in upper case.
 - Place multiple entries on separate lines. For example, to back up a log file directory named C:\Docs\CustomerLogs and all of the data under the directory named D:\accounting, the entries would look similar to:


```
C:\Docs\CustomerLogs
D:\accounting
```
 - On Windows, when specifying a drive name, always use the backslash “\” otherwise, the synthetic full backup will fail. For example, to specify the D: drive, type **D:**.
 - On UNIX, include the forward slash to designate root “/” when specifying a save set name. Otherwise, the synthetic full backup will fail. For example, if /tmp is misspelled as tmp in the save set list of the Client resource, the backup will fail.

4. Create a **Group** for the synthetic full backups, as described in [“Task 2: Set up a group for backup clients” on page 61](#):
 - Ensure that the groups that are used for the synthetic full backup operations contain only the save sets that are compatible with synthetic full backup operations. [“Limitations of synthetic full backups” on page 80](#) provides details. The save sets that are incompatible with synthetic full feature are reported as failed save sets in the savegrp log file.
 - Configure Windows clients within a dedicated group; not mixed with UNIX clients.
 - If multiple groups have been configured to run concurrently, set the Parallelism attribute in the client resource that was created for the NetWorker server to 40. You can find the Parallelism attribute in the Globals (1 of 2) tab of the Client property dialog box. This limits the number of concurrent synthetic full operations to 20. The Parallelism setting should be divided by 2 to formulate the number of concurrently running synthetic full operations.

NOTICE

The number of concurrent synthetic full operations in a data zone should be limited to 20.

The optimal number of concurrent synthetic full operations operation depends on the following:

- Configuration of the NetWorker server
- Size of the save sets, number of clients
- Number of savegrp instances that are concurrently running
- Number of other active operations, such as cloning that are concurrently running.

5. Open the schedule to which the synthetic full backup will be applied, and use either of the following:
 - Preconfigured Incr+Synthetic Full schedules:

Using one of these schedules forces the NetWorker software to perform an incremental backup prior to performing a synthetic full back. This backs up all of the data that has changed since last full backup and subsequent incrementals to create a synthetic full backup:

 - Incr+Synthetic Full 1st Friday of Month
 - Incr+Synthetic Full Every Friday
 - Incr+Synthetic Full 1st of Month
 - Incr+Synthetic Full Quarterly
 - Create a new schedule that uses the Synthetic Full+Incr backup level — Using this schedule will force the NetWorker software to perform an incremental backup of the save group immediately prior to performing a synthetic full.

Use this level when synthetic full backups fail because the incremental backup chain is broken or one or more incremental backups did not have Backup renamed directories attribute enabled. This level triggers an incremental to be created before the synthetic full backup. This mends the broken chain of incremental backups. [“Task 1: Set up a schedule for backups” on page 60](#) provides information about creating a schedule.

NOTICE

For Windows clients, ensure that the backup levels in the schedule are equal to or greater than the previous backup level. This prevents renamed files from being skipped when the synthetic full backup is created. For example, if mixed backup levels are run before a synthetic full backup is run, renamed files might not be included in the synthetic full backup since the level backup following the incremental backup might not include the renamed files.

Performing manual incremental synthetic full backup by using the **savegrp** program

The **savegrp** program can be run from the command line of the NetWorker server to perform an incremental synthetic full level backups of a particular group.

The **-l** option is used to indicate that the backup level is a synthetic full or incremental synthetic full. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information.

For example:

```
savegrp -N 20 -l inc_synth_full group_name
```

The **-l incr_synth_full** option initiates an incremental backup from the client system before running the synthetic full process. This ensures the time difference between T4 and T4+delta is minimized.

Performing manual synthetic full backups using **nsrconsolidate**

The **nsrconsolidate** program can be run from the command line of the NetWorker server to perform more granular synthetic full backup. It cannot be used to perform an incremental synthetic full level backup of a particular group.

nsrconsolidate can be used to define backup data to be included in the synthetic full backup by:

- ◆ client name and save set name
- ◆ ssid/cloneid
- ◆ time range

The **nsrconsolidate** command can be used to create synthetic full backups of data stored on an Avamar deduplication node. The deduplicated data is rehydrated (reverted) to its original non-deduplicated state.

NOTICE

For the Microsoft Windows operating systems, even though most file systems are case-independent, you must use the same path name case that the Windows file system uses when specifying a save set name, client name, file, or directory. The NetWorker software's cross-platform indexing system is case-sensitive. It is best practice to always specify the Windows drive letter in upper case.

When running the **nsrconsolidate** command, it is best practice to run fewer **nsrconsolidate** commands that include many save sets than to run multiple **nsrconsolidate** commands with a fewer number of save sets. This helps **nsrconsolidate** to manage the number of concurrent synthetic full operations and reduce resource usage.

The optimal number of concurrent **nsrconsolidate** operations depend on the following:

- ◆ NetWorker server configuration
- ◆ Size of the save sets
- ◆ Number of clients
- ◆ Number of savegrp instances that are running concurrently
- ◆ Number of other active operations, such as cloning that are running concurrently

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information.

Checking the integrity of a synthetic full backup

To check the integrity of a synthetic full backup:

1. Open the group that contains the clients that are participating in the synthetic full backup.
2. In the **Advanced** tab of the properties of the group:
 - If not already enabled, enable the **Verify synthetic full** option to verify the integrity of the new indexes entries that are created for the synthetic full backup in the client file index.
 - If not already enabled, enable the **Revert to full when synthetic full fails** option to force the group to perform a full backup on the save set if the synthetic full backup fails.

NOTICE

If the **clones** option is enabled for the group, all backups will be cloned including the synthetic full backup.

If a synthetic full backup fails because the incremental backup chain is broken or one or more incremental backups did not have **Backup renamed directories** attribute enabled, you can change the level attribute in the Group resource to **incr_synth_full**. This level triggers an incremental to be created before the synthetic full backup. This mends the broken chain of incremental backups.

To change the level of a group:

1. Open the group that contains the clients that are participating in the synthetic full backup.
2. In the **Advanced** tab of the properties of the group, select **incr_synth_full** from the **level** attribute.
3. Click **Ok**.

Running the **savegroup** mends the broken chain of incremental backups by triggering a new incremental backup to be created prior to the next synthetic full backup.

Checkpoint restart considerations with synthetic full backups

Backups that are performed during a checkpoint restart might be included in a synthetic full backup if the standard [“Synthetic full requirements” on page 79](#) are met. [“Checkpoint restart backups” on page 95](#) provides more information on how to configure and use checkpoint restart backups.

Reporting and synthetic full backups

When defining the backup statistics or backup status report in the Console:

- ◆ For the Save Sets Details report, the value **Synthetic** in the **Type** column indicates that the backup is a synthetic full.
- ◆ For the Save Sets Details by client report, the value **Synthetic** in the **Type** column indicates that the backup is a synthetic full.

[“Preconfigured reports” on page 435](#) provides more information.

Running queries on synthetic full backups

To view information on the partial save sets:

1. In the **Administration** window, click **Media**. Media-related topics appear in the navigation tree.
2. Select **Save Sets**.
3. Select the **Query Save Set** tab.
4. Click **Select From** in the **Type** pane and then click the **Synthetic Full** checkbox.
5. Select the **Save Set List** tab to view the result of the save set query.
6. Review the **Type** column. **Synthetic Full** appears in the **Type** column if the save set is part of a synthetic full backup.

Monitoring

A new table called Synthetic Full Sessions is available in the Monitoring section of the console. It contains one row for each client and save set pair. [“Monitoring NetWorker server activities” on page 462](#) provides more information.

Directives

The following directives can be used when performing synthetic full backups:

- ◆ compressasm
- ◆ aes (encryption)

[“Compressing backup data” on page 109](#) and [“Encrypting backup data” on page 108](#) provide more information on using compressasm and aes.

When using directives with synthetic full backups, consider the following:

- ◆ If directives were applied to save sets during the full and incremental backups that participate in the synthetic full backup, the synthetic full backup will not remove those directives.
- ◆ If any directives including the compressasm or aes directives were applied to the full and incremental backups that participate in the synthetic full backup; these directives will not be applied again.
- ◆ If the target volume for the synthetic full backup is a Data Domain device, directives should not be set for the synthetic full backup.
- ◆ If an unsupported directive is specified during a synthetic full backup, the directive is ignored.

Review the **nsrconsolidate** syntax in the *EMC NetWorker Command Reference Guide* or the UNIX man pages for more information.

NOTICE

Directives cannot be used with virtual synthetic full backups, so this option must be disabled prior to running the backup. When applying directives to traditional synthetic full backups, ensure that the virtual synthetic full option is unchecked in NMC. The section [“Virtual synthetic full backups \(for Data Domain systems\)” on page 88](#) provides more information.

Directing data from a synthetic full backup to a dedicated pool

To direct consolidated save sets to a specific set of media:

1. Create a **Group** for consolidation backups, as described in [“Task 2: Set up a group for backup clients” on page 61](#).
2. Create a **Pool** for consolidation backups, as described in [“Creating a media pool” on page 311](#).
3. In the **Pool** resource, select the group that was created in [step 1 on page 88](#) as the value for the **Groups** attribute.
4. Edit the **Group** attribute in each **Client** resource that is to have consolidated backups, and assign to the group that you created in [step 1 on page 88](#).

Virtual synthetic full backups (for Data Domain systems)

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 new full backup. However, since the backup occurs on a Data Domain system using new DD Boost APIs, the backup does not require save set data to be sent over the wire, resulting in improved performance over synthetic full and traditional backups.

The following table compares traditional synthetic full and virtual synthetic full.

Table 14 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 DDR
Supports read/write for all types of volumes	Only Data Domain devices are supported, and the source and destination volumes must belong to the same DDR. As per the DDBoost API, there are no restrictions if the volumes belong to different M-Trees in the same DDR.
Client file index created by nsrrecovery	Client file index created by nsrconsolidate
Does not require DFA support	Requires DFA support
Supports Avamar save sets	Cannot use Avamar save sets for creating the virtual synthetic full.

VSF requirements

VSF requires the following:

- ◆ DDOS version 5.3 is the minimum supported DDOS version that supports NetWorker VSF backup and clone.
- ◆ DDBoost version 2.6; this version of DDBoost is shipped with NetWorker 8.1.
- ◆ The DDR must have virtual synthetics enabled.
- ◆ All of the constituent backups must be on the same DDR.
- ◆ Only non-Avamar clients are supported.

Note: If all the required Avamar deduplicated save sets are rehydrated to the same DDR prior to initiating a VSF, a VSF backup can be performed, but not recommended.

- ◆ The client must have the DFA attribute enabled. Also, a value is required in the volume location attribute for a Data Domain device. NetWorker updates the volume location attribute during the device mount operation. Before you update a storage node that uses Data Domain devices, unmount each device. Once the update completes, mount each device.
- ◆ The DDR must have **virtual-synthetics** enabled. When you log into the Data Domain system and run **ddboost option show**, ensure that the output indicates that virtual synthetics is enabled, as shown in the following:

```
data_domain_user@data_domain_host# ddboost option show
Option                               Value
distributed-segment-processing      enabled
virtual-synthetics                  enabled
fc                                   enabled
```

- ◆ If upgrading the NetWorker client from a pre-NetWorker 8.1 release to NetWorker 8.1, a Full level backup is required before performing a VSF backup. If the Full level backup is not performed prior to the VSF backup, file-by-file recovery fails.

If all requirements are met and synthetic full backups are enabled, VSF will automatically be performed. If NetWorker detects that one or more of the requirements are not met, then traditional synthetic full mode is used.

If the only missing requirement is that DDR does not have “virtual-synthetics” enabled, NetWorker does not fall back to traditional synthetic full and attempts to perform VSF. In this case, the VSF backup will fail with errors.

Enable/disable VSF in NMC

VSF is enabled by default and will be performed if all the requirements are met. If you do not want to use VSF, you can disable the option.

To disable VSF in NMC:

1. In the Administration window, click **View** and select **Diagnostic Mode**.
2. Click **Configuration**.
3. Select **Groups**.
4. Right-click on the save group for which VSF is to be disabled and select **Properties**.
5. Select the **Advanced** tab in the properties dialog.
6. In the **Operations** pane, uncheck **Perform virtual synthetic full**.




Figure 10 Enable/disable Perform virtually synthetic full in NMC

Performing VSF

Virtual Synthetics will be created in the same manner as traditional synthetic full backups by using one of the following commands:

- ◆ **savegrp -l**
- ◆ **nsrconsolidate**

Using the savegrp program to perform incremental synthetic full backup

“Performing manual incremental synthetic full backup by using the savegrp program” on page 85 describes how to perform VSF using the **savegrp** program. The procedure for using VSF is the same as traditional synthetic full.

Using the nsrconsolidate command to perform VSF

“Performing manual synthetic full backups using nsrconsolidate” on page 85 describes how to perform VSF using the **nsrconsolidate** command. The procedure for VSF is the same as traditional synthetic full.

Directives

Directives cannot be used with VSF backups as VSF is created by the Data Domain system.

Directives can still be applied to traditional synthetic full backups, if the VSF option is disabled in NMC.

[“Directives” on page 87](#) provides more information.

Multiple storage node distribution

The VSF save sets backed up from NetWorker Data Domain devices can be distributed across multiple storage nodes if:

- ◆ All of the save sets being used in the VSF are stored on a single Data Domain storage system, and
- ◆ The VSF requirements identified in the section [“VSF requirements” on page 89](#) are met.

Also, save sets located in different mtrees on the Data Domain system can be used in a VSF.

Concurrent operations

The concurrent volume of Virtual Synthetics that a DDR can handle depends on the DDR model and the capacity of the NetWorker host. The following scenarios have been tested and verified to work:

- ◆ Concurrent VSF
- ◆ VSF concurrent with clone
- ◆ VSF concurrent with Clone-controlled replication

Avamar rehydrated save sets

If all the required Avamar deduplicated save sets are rehydrated to the same DDR prior to initiating a VSF, a VSF backup can be performed. Rehydrate Avamar save sets using the `nsrconsolidate` command, as specified in [“Using the nsrconsolidate command to perform VSF” on page 90](#).

Note: Rehydrating Avamar save sets for the purpose of performing a VSF backup is not recommended. If a VSF backup of Avamar save sets is performed, specify the following `nsrconsolidate` options to properly rehydrate Avamar savesets prior to VSF:

```
nsrconsolidate -R -i on -S ssid1 ssid2
```

Validating VSF

You can validate VSF backups using any of the following options:

- ◆ `mminfo` commands
- ◆ NMC
- ◆ `savegrp` logs

Using mminfo to validate VSF

Run the following commands to validate VSF backups:

- ◆ **mminfo -aS** — shows detailed information about Synthetic Full backups, including information about the save sets used to form the synthetic full.
- ◆ **mminfo -q syntheticfull -c {client} -N {save set_name}** — queries all synthetic full save sets for the specified client and save set names
- ◆ **mminfo -q rehydrated** — queries rehydrated Avamar save sets
- ◆ **mminfo -avot -q dedupe** — queries all deduplicated Avamar savesets

Using NMC to validate VSF

In NMC, query synthetic full backups by navigating to **Media > Save Sets** in the NetWorker Administration window. Select the **synthetic full** checkbox to return information for all synthetic full save sets when you click **Save Set List** in this window. “[Running queries on synthetic full backups](#)” on page 87 provides more information.

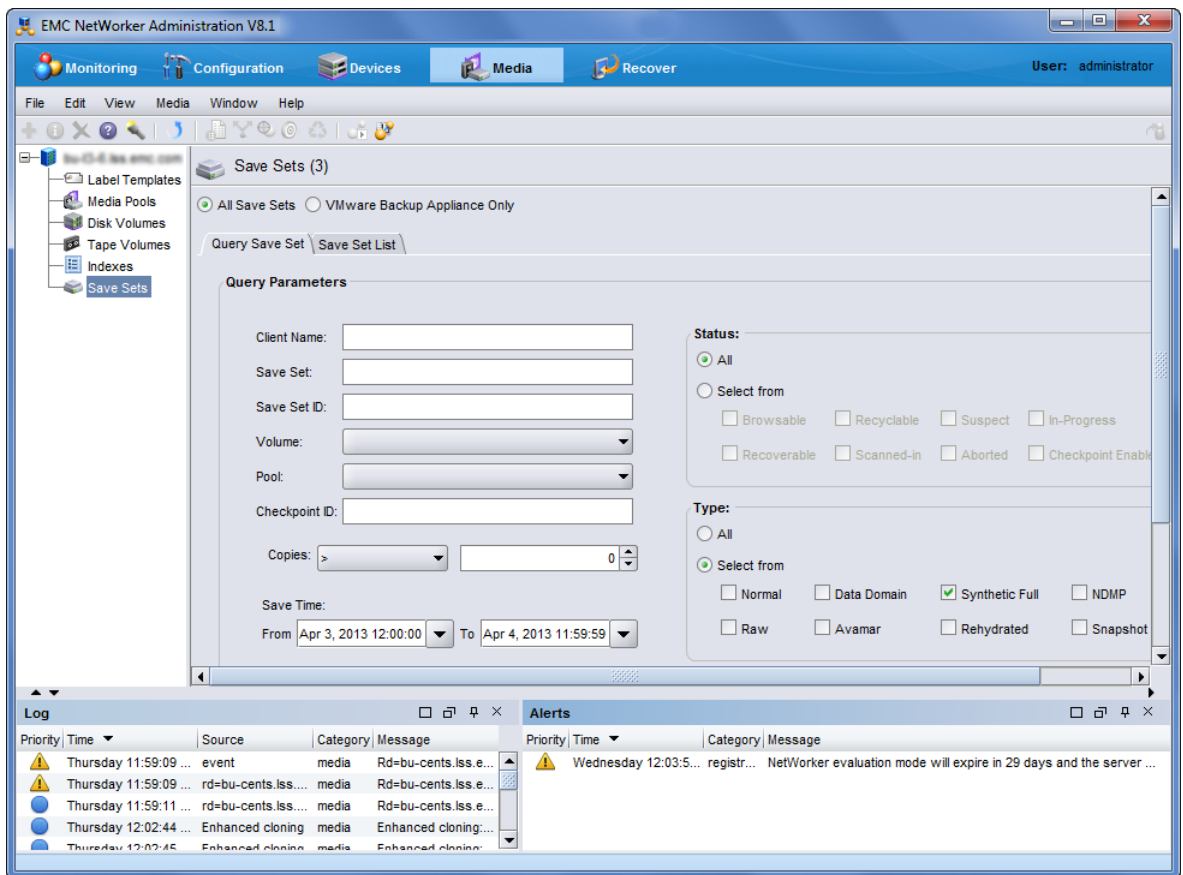


Figure 11 Save Sets query in NMC

Viewing the savegrp logs to validate VSF

Use savegroup logs to determine when VSF backups were attempted. The following excerpts from various savegrp log files illustrate the types of messages NetWorker displays when performing VSF, traditional synthetic full, and falling back from VSF to SF:

```
1707:97864:nsrconsolidate: Unable to perform virtual synthetic full
operation since at least one source save set is an Avamar save set.
Falling back to non-virtual synthetic full mode

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
```

VSF limitations

The following limitations apply to using VSF:

- ◆ All DD devices on a storage node must be unmounted prior to upgrading to NetWorker 8.1 if planning to use VSF. If this is not performed, then a new full backup will be required after upgrading and prior to performing any VSF backups.
- ◆ DD Archivers must be running DD OS 5.3 or later.
- ◆ Cloning fails if the clone device is on a DDR that does not have VSF enabled.

Enable parallel save streams

In NetWorker 8.1 and later, the parallel save streams (PSS) feature provides the ability for Client resource save sets (UNIX and Linux) to be backed up by multiple parallel save streams to one or more destination backup devices. PSS is used for the scheduled, file-based backup of file systems. PSS is disabled by default.

To enable PSS:

1. In **Configuration**, select **Clients**.
2. Right-click the client for which to enable PSS, and select **Modify Client Properties....**
3. In **Globals (1 of 2)**, click the check box in the **Optimizations** section to enable **Parallel save streams per save set**.
4. Specify save set **All** or **/paths** in the **Save Set** attribute of the client for which to enable PSS.

Note: Currently there is no PSS support for Synthetic Full, Checkpoint Restart, or Avamar deduplication backups. Also, both the NetWorker server and client must be at NetWorker 8.1 or later in order to use the PSS functionality for full and incremental level backups, recovery, and cloning.

The *NetWorker Performance Optimization Planning Guide* provides complete details on PSS requirements and performance benefits.

Troubleshooting PSS

To troubleshoot PSS with the guidance of EMC technical support:

1. Perform the following:

- Modify the Client resource backup command attribute:

```
save -v -D7
```

- On the client, perform the following command:

```
touch /nsr/debug/mbsdfopen
```

2. Enable the -v verbose option for scheduled backups:

```
NMC > Group Properties > Advanced > Options > Verbose > checkbox
```

3. Run the save group, by using any of the following methods:

- NMC:

```
NMC > Group Properties > Setup > Setup > Autostart > Start now
```

- From the server command line:

```
savegrp -v 2>&1 | tee /nsr/tmp/savegrp.log
```

- Wait for the normal automated scheduled group run.

4. After the save group finishes, collect the following log files for EMC technical support:

- Client:

```
/nsr/tmp/save-mbs-* log files
```

- Server:

```
/nsr/logs/daemon.raw log file  
/nsr/logs/sg/<group name>/* log files  
/nsr/tmp/savegrp.log log file
```

The *NetWorker Performance Optimization Planning Guide* provides complete details on PSS requirements and performance benefits.

Probe-based backups

The NetWorker server schedules probe-based backups that are based on user-defined events for clients and NetWorker modules, in addition to time-based events.

To run probe-based backups, perform the following tasks:

1. Create a user-defined client probe. [“Creating a client probe” on page 610](#) provides detailed information.
2. Associate a probe with a Client resource. [“Associating a probe with a Client resource” on page 611](#) provides detailed information.

3. Create a probe group. [“Preconfigured groups” on page 250](#) provides detailed information.
4. Run the probe group the same as you would a save group. [“Scheduled backups” on page 58](#), or [“Manual backups” on page 70](#) provides detailed information.

NOTICE

Running the **savegrp** program with the **-g** attribute bypasses probing when running a probe group backup from the command line.

Client Direct backups

The NetWorker 8.0 client software enables clients with network access to AFTD or DD Boost storage devices to send their backup data directly to the devices, bypassing the NetWorker storage node. The storage node manages the devices for the NetWorker clients, but does not handle the backup data. The **Client Direct** feature reduces bandwidth usage and bottlenecks at the storage node, and provides highly efficient backup data transmission.

Destination devices must specify their complete paths in their **Device Access Information** attribute. If the **Client Direct** backup is not available, a traditional storage node backup is performed instead. The **Client Direct** feature is enabled by default, but can be disabled on each client by the **Client Direct** attribute.

[“Considerations for Client Direct clients” on page 176](#) provides details.

Checkpoint restart backups

The checkpoint restart feature allows a failed backup operation to restart at a known good point, prior to the point-of-failure during the backup. 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.

This section includes the following information about the checkpoint restart feature:

- ◆ [“Checkpoint restart usage” on page 96](#)
- ◆ [“About partial non-NDMP save sets” on page 97](#)
- ◆ [“Configuring checkpoint enabled clients” on page 98](#)
- ◆ [“Restarting a checkpoint-enabled backup” on page 100](#)
- ◆ [“Monitoring checkpoint-enabled backups” on page 101](#)
- ◆ [“Reporting checkpoint-enabled backups” on page 102](#)
- ◆ [“Recovering checkpoint restart data” on page 106](#)
- ◆ [“Cloning and scanning partial savesets” on page 107](#)
- ◆ [“Cloud backup devices and partial savesets” on page 107](#)

Checkpoint restart usage

Backup failures occur for various reasons. The most common reasons include: hardware failures, loss of network connectivity, and primary storage software failures. 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, 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 encountered during a scheduled backup operation might require a re-run 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 6 tapes, if a failure occurs while writing to the last tape, the previous 5 tapes representing 9 hours of backup time may need to be re-run. As data sets continue to increase in size, so does the impact of backup failures.

Checkpoint-enabled clients provide the following enhancements:

- ◆ Failed save sets are marked as partial; not as aborted.
- ◆ Restarted save sets have a new SSID and savetime.
- ◆ Partial non-NDMP save sets are indexed.
- ◆ For partial NDMP savesets, only the first saveset has an index associated with it. The index covers all of the files in all of the partial save sets that make up a complete backup.
- ◆ Partial save sets are not removed from the index, the media databases, and media such as AFTD.

Support and considerations

The following considerations apply to checkpoint restart configurations:

- ◆ Checkpoint restart is not enabled by default.
- ◆ Checkpoint restart does not support Client Direct backups 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.
- ◆ Starting in NetWorker 8.0, Checkpoint restart does not support Avamar deduplication backups.
- ◆ Checkpoint restart does not support NDMP NetApp clients for releases earlier than NetWorker 8.0 server and client software. [Chapter 21, “NetWorker support for NDMP,”](#) provides details.
- ◆ Checkpoint restart does not support NDMP Isilon clients for releases earlier than the NetWorker 8.1 SP1 server and client software. [Chapter 21, “NetWorker support for NDMP,”](#) provides details.
- ◆ Checkpoint restart for non-NDMP clients is not supported for releases earlier than NetWorker 7.6 SP1 server and client software.

- ◆ 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. “Save sets” on page 66 provides details.
- ◆ For Client Direct backups to AFTDs, checkpoints are not made less than 15 seconds apart. Checkpoints are always made after larger files that require more than 15 seconds to back up.
- ◆ The checkpoint restart option is ignored for index and bootstrap save sets.
- ◆ Checkpoint-enabled might impact the backup speed. This depends on the datazone environment and configuration.

Using the checkpoint restart feature might increase the size of the index. This increase might occur because additional index records are created for the valid recoverable data. These partial save sets should not be manually removed from the index.

About partial non-NDMP save sets

The backup sequence of partial save sets is not the same as for complete backups. Each partial save set provides protection for a part of the filesystem, but the completeness and consistency of the coverage of the whole filesystem 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 filesystem that is different from the state of the filesystem at any given point-in-time. The resulting filesystem is not guaranteed to be consistent. It can be different than at any other point-in-time.

Files and directories are backed up in alphabetical order. If a crash occurs, subsequent backups continue from the last point alphabetically from where they were in progress. Previously backed up files or directories are not reviewed for modifications. If a file or directory that is earlier alphabetically was modified or added, it will not be backed up.

Example

A backup is interrupted while saving a directory and restarted after the directory contents have changed. As a result, different files are saved than the original filesystem entry.

For example:

1. A save set contains */disk1/dir*. The files include *file_a*, *file_c* and *file_d*.
2. A point of interruption occurs in the backup of the save set when *file_d* is being backed up.
3. The first partial save set contains *file_a* and *file_c*.
4. Before the checkpoint restart is initiated for the save set, *file_b* is added to the filesystem.
5. The second partial save set contains *file_d* and */disk1/dir*.

NOTICE

/disk1/dir contains *file_a*, *file_b*, *file_c*, and *file_d*.

6. Notice that *file_b* has not been backed up.

About partial NDMP save sets

For checkpoint enabled backups of an NDMP client, the NetApp filer creates a snapshot of the filesystem prior to the start of the backup. The save set is generated from the snapshot. If the NDMP backup is interrupted and later restarted, the partial save sets are generated from the snapshot. As a result, the partial backups provide an image of the filesystem from the point-in-time that the snapshot is taken.

Configuring checkpoint enabled clients

NetWorker clients can be configured to allow an interrupted backup to restart from the point-of-failure.

Consider the following:

- ◆ For non-NDMP save sets the backup can be restarted at the directory or file level from the point of failure. The level is defined by the Checkpoint Granularity attribute for the client.
- ◆ For NDMP save sets, checkpoints are made at regular time intervals during the backup. The interval is set by an environment variable defined in the Application Information attribute for the client. The backup is restarted from the last successful checkpoint.
- ◆ The checkpoint restart feature is not enabled by default.
- ◆ Configuring a client as checkpoint enabled might impact the backup speed. This is dependent upon the data zone environment and its configuration.
- ◆ Ensure that all NetWorker clients are configured with the same name. There should not be both short and fully qualified domain name (FQDN) client resources. [“Hostname aliases” on page 813](#) provides more information.

To configure a client for checkpoint enabled backups:

1. From the **Administration** window, click **Configuration**.
2. Set the **Checkpoint enabled** attribute:
 - a. In the expanded left pane, select **Clients**.
 - b. Right-click the client to be enabled.
 - c. Select **Properties**. The **Properties** dialog box appears.
 - d. Click the **General** tab.
 - e. Click the **Checkpoint enabled** checkbox.
3. For the **Checkpoint granularity attribute**:
 - a. Select whether to restart the backup **by directory** or **by file**. This value is not applicable to NDMP clients and will be ignored.

Restart by directory is the default. 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.

NOTICE

Use the **restarting by file** option only for save sets with few large files. Committing every file to the index and the media database is time consuming. This might lead to performance degradation during a backup that contains many small files.

- b. Click **OK**.
4. For NDMP NetApp clients only, define the interval at which checkpoints are written during the backup using the **CHECKPOINT_INTERVAL_IN_BYTES** variable. This variable is added to the **Application Information** attribute located under the **Apps and Module** tab.

The value defined for **CHECKPOINT_INTERVAL_IN_BYTES**:

- Is in bytes by default.
For example: **CHECKPOINT_INTERVAL_IN_BYTES=1000000**
- Can be defined using different multipliers. Acceptable multipliers include: KB, MB, GB, TB, kb, mb, gb, and tb.
For example: **CHECKPOINT_INTERVAL_IN_BYTES=1GB**
- Is automatically rounded up to a multiple of the tape blocksize.

5. Select the group to which the checkpoint enabled client belongs:
 - a. In the expanded left pane, select **Groups**.
 - b. Click the **Advanced** tab.
 - c. If required, change the **Client Retries** attribute to a value greater than 0. This value specifies the number of times the NetWorker software attempts to back up a failed client.
 - d. Click **OK**.

Restarting a checkpoint-enabled backup

Restarting a checkpoint-enabled backup of a partial save set is faster than restarting the backup of a save set from the beginning. This depends on how much data was saved in the previous backup.

There are two ways in which a checkpoint-enabled backup can be initiated:

- ◆ Manually through a group restart operation.
- ◆ Automatically by setting the Client retries attribute in the Group resource of the checkpoint-enabled client.

NOTICE

Changing the name of a save set will cause checkpoint restart to fail to find a match against a previous run, and the restart will revert to a complete backup. Additionally, do not modify browse or retention policies for the client in between checkpoint restarts, as an expired partial save set may leave gaps in the backup set.

Manually restarting a checkpoint enabled backup by using the Group Restart attribute

When a **Group** is restarted, the NetWorker software determines which save sets were not completed within the backup:

- ◆ If the client is Checkpoint-enabled, the incomplete save sets are checkpoint restarted and continue from the point at which they were stopped.
- ◆ If the client has not been checkpoint-enabled, the incomplete save sets are backed up again in full.

To manually initiate a group restart:

1. Ensure that the client has been checkpoint-enabled. [“Configuring checkpoint enabled clients” on page 98](#) provides detailed information.
2. From the **Administration** window, click **Monitoring**.
3. Click **Groups** in the docking panel.
4. Right-click the group to which the checkpoint enabled client belongs, then select **Restart**.
5. Click **Yes** to confirm the restart.

Setting the Client retries attribute

If the NetWorker server fails to connect to a client, the Client retries attribute specifies the number of times that the server will reattempt the connection to the client before the backup is considered a failure. The Client retries mechanism is the same for checkpoint-enabled clients and non-checkpoint enabled clients, with the exception that a partial save sets is created when there is a failure for a checkpoint-enabled client.

If the NetWorker software detects that the backup fails and the number of client retries is not exceeded, the NetWorker software will checkpoint-restart the backup, immediately after the failure. This operation takes into account the Group restart window and will not restart the backup if the defined backup window has expired.

Example 1:

There are 6 clients in a group, each with 3 save sets. The Client retries attribute for the group is 1. One save set fails and it is checkpoint restarted immediately. The remaining save sets in the group continue to backup. 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 that is defined in the Client retries attribute.

When all of the save set backup attempts in the group complete, the group completion report:

- ◆ Provides a list of the successful save sets.
- ◆ Reports that the failed partial save set is unsuccessful.
- ◆ Reports that the group has failed.

Example 2:

There are 6 clients in a group, each with 3 save sets. The Client retries attribute for the group is 2. One save set fails and it is checkpoint restarted immediately. The remaining save sets continue to backup. The partial save set fails a the second time and it is checkpoint restarted immediately. This time, the partial save set succeeds.

When all of the save set backup attempts in the group are complete, the group completion report:

- ◆ Provides a list of the successful save sets.
- ◆ Reports that the two partial save sets are successful.
- ◆ Reports that the group completed successfully.

To increase the number of times each client in a group is retried before the backup attempt is considered unsuccessful, change the value in the Client retries attribute in the Group resource. [“Configuring checkpoint enabled clients” on page 98](#) describes how to set the Client retries attribute.

Monitoring checkpoint-enabled backups

To view detailed information about a checkpoint-enabled backup:

1. From the **Administration** window, select **Monitoring > Groups**.
2. Right-click the group to which the checkpoint enabled client belongs, then select **Show Details**.
3. View the detailed information related to the group backups:
 - If the partial save set is in the work list for the group, the save set appears in the **Waiting to Run** section.
 - If the partial save set is running, the save set appears in the **Currently Running** section.
 - If the entire partial save sets sequence of the savegroup is complete, the save set appears in the **Completed Successfully** section.
 - If the entire partial save sets sequence of the savegroup is not complete, the save set appears in the **Failed** section.

NOTICE

If any messages are generated, the **Show Messages** button is enabled. Click **Show Messages** to view the messages.

4. Click **OK** to close the **Group Details** window.

Reporting checkpoint-enabled backups

The `daemon.raw` file on the NetWorker server contains details about groups that are run with checkpoint-enabled clients. When a group backup complete, the savegroup completion report also reports the status of each client backup.

The following sections describe how to review group information for checkpoint-enabled clients:

- ◆ [“Reviewing the savegroup reports for checkpoint-enabled client backups” on page 102](#)
- ◆ [“Reviewing the NetWorker server `daemon.raw` file for the status of a checkpoint-enabled backup” on page 103](#)
- ◆ [“Querying the media database for partial save sets” on page 104](#)

Reviewing the savegroup reports for checkpoint-enabled client backups

Consider the following when reviewing the savegrp completion report for a savegroup that contains a checkpoint-enabled client.

- ◆ When a checkpoint-enabled client backup attempt fails:

- The savegroup status is reported as a failure:

```
nsrd info, savegroup failure alert: test Completed/Aborted, Total
1 client(s), 0 Clients disabled, 0 Hostname(s) Unresolved, 0
Failed, 0 Succeeded, 1 CPR Failed, 0 CPR Succeeded, 0 BMR Failed,
0 BMR Succeeded.
```

```
nsrd info, savegroup alert: <group_name>aborted, Total 1
client(s), 1 CPR Failed. Please see group completion details for
more information.
```

- The failed save sets are reported in the **Unsuccessful Save Set** status section:

```
* cprclient.emc.com:/usr/sbin, number of checkpoint enabled
save sets 1
* cprclient.emc.com:/usr/sbin 86705:save: Successfully
established DFA session with adv_file device for save-set ID
'4078798790' (bu-t3-7.lss.emc.com:/usr/sbin).
* cprclient.emc.com:/usr/sbin (interrupted), exiting
* cprclient.emc.com:/usr/sbin aborted
```

- ◆ When a checkpoint-enabled client backup succeeds:

- The savegroup status is reported as a success:

```
NetWorker savegroup: (notice) test completed, Total 1 client(s),
1 CPR Succeeded. Please see group completion details for more
information.
```

- The total number of partial save sets that make up the checkpoint save sets is displayed in the **Save Set Summary** section:

```
client_name:save_set, number of checkpoint enabled savesets x
```

- The failed save sets are reported in the **Successful Save Set** status section:

```
* cprclient.emc.com:savefs savefs cprclient.emc.com: succeeded.
bu-t3-7.lss.emc.com:/usr/sbin, number of checkpoint enabled
savesets 2
* cprclient.emc.com:/usr/sbin 86705:save: Successfully
established DFA session with adv_file device for save-set ID
'4062021648' (bu-t3-7.lss.emc.com:/usr/sbin).
```

Reviewing the NetWorker server daemon.raw file for 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.

Typical messages include:

- ◆ “nsrd info, Savegroup Info: group_name:client_name checkpoint enabled, mode: mode. (severity 0, message 71193)” on page 103
- ◆ “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” on page 103
- ◆ “savegrp group_name checkpoint restartable saveset client_name:save_set failed and will not be restarted” on page 103
- ◆ “savegrp group_name: checkpoint restartable saveset client_name:save_set completed without interruption” on page 104

nsrd info, Savegroup Info: *group_name:client_name* checkpoint enabled, mode: *mode*. (severity 0, message 71193)

This message is reported when a savegroup is started. 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 will occur 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 group backup.

Querying the media database for partial save sets

The savegrp completion report does not provide detailed information about partial save sets that might be necessary to perform a recovery.

There are two ways to query the NetWorker server for partial save set information:

- ◆ [“Querying partial save sets from the Console” on page 104](#)
- ◆ [“Querying partial save sets by using the mminfo command” on page 105](#)

Querying partial save sets from the Console

To view information about the partial save sets:

1. In the **Administration** window, click **Media**. Media-related topics appear in the navigation tree.
2. Select **Save Sets**. The following tabs appear in the **Save Sets** window:
 - **Query Save Set**
 - **Save Set List**
3. Select the **Query Save Set** tab, to query:
 - All partial save sets, select **Checkpoint Enabled**.
 - All partial save sets with the same Checkpoint ID, in the **Checkpoint ID** field, type the **Checkpoint ID** of the partial save set on which you want to perform the query.
4. Select the **Save Set List** tab to view the result of the save set query:
 - The **Checkpoint ID** column displays the partial save set **Checkpoint ID** and its Sequence ID. The **Checkpoint ID** is listed first followed by the **Sequence ID**, which is encased within brackets.
 - Sort the **Checkpoint ID** column to view the complete sequence of partial save sets.
 - The **Status** column displays the status of the partial save sets:
 - A **Checkpoint browsable** status indicates that the save sets can be browsed for recover.
 - A **Checkpoint aborted** status indicates that the backup of the partial save set was stopped or aborted. A save set recover is used to recover the partial save set.

Consider the following:

- When a checkpoint-enabled backup completes successfully, the status of the last partial save set is Checkpoint browsable.
- When a checkpoint-enabled backup completes successfully, on the first backup attempt, the save set status is Checkpoint browsable. Only one Sequence id is associated with the Checkpoint ID. The Sequence id is 1. If the Sequence id is 2, the first partial save set in the checkpoint-enabled backup is missing.

NOTICE

If no partial save sets are found that match the query, ensure that the backup of the partial save sets was started within the **Save Time** period. To change the values for the Save Time attribute, open the Save Set Query tab and select a date and time from the Save Time calendar.

Querying partial save sets by using the `mminfo` command

By default, the `mminfo` command output only displays the browsable save sets. The first and intermediate partial save sets are not displayed. Only complete checkpoint-enabled save sets or final partial save sets are displayed.

Use the `mminfo` command with specific queries to display more information about checkpoint-enabled save sets.

The following new media attributes support the Checkpoint Restart feature:

- ◆ `checkpoint_id` — Displays the checkpoint restart id of the partial save set in the `chkpt_id` column.
- ◆ `checkpoint_seq` — Displays the partial save set sequence id in the `chkpt_seq` column.
- ◆ `checkpoint-restart` — This flag attribute is used to only display checkpoint restart enabled save sets.

In addition, several media sumflags are used with the Checkpoint Restart feature:

- ◆ `k` — Indicates this is a checkpoint enabled save set.
- ◆ `a` — The first and all intermediate partial save sets of a checkpoint sequence will have aborted status.
- ◆ `b` — The last partial or complete save set of a checkpoint sequence will be marked browseable.

Displaying checkpoint enabled save sets

To display all checkpoint enabled save sets, type:

```
# mminfo -q 'checkpoint-restart' -r 'client,nsavetime,ssid(11),
sumflags(3),name,checkpoint_id,checkpoint_seq'
```

```
client  save time      ssid            ssflags  filename  chkpt_id  chkpt_seq
plapew  1251910303  4204700319    cak      /space    1251910303  1
```

client	save time	ssid	ssflags	filename	chkpt_id	chkpt_seq
plapew	1251910327	4187923127	cbk	/space	1251910303	2
plapew	1251910710	4087260214	cak	/space	1251910710	1
plapew	1251910725	4070483013	cbk	/space	1251910710	2

Displaying all partial save sets for the checkpoint id

To display all partial save sets for the **checkpoint id**, type:

```
mminfo -q "checkpoint_id=1251910303"
```

volume	client	date	size	level	name
plapew.001	plapew	09/02/09	17 MB	full	/space
plapew.001	plapew	09/02/09	799 MB	full	/space

Recovering checkpoint restart data

This section outlines how to recover data in the following two scenarios:

- ◆ [“Recovering data from the complete sequence of partial backups that comprise the original save set” on page 106](#)
- ◆ [“Recovering data from a partial save set” on page 107](#)

Recovering data from the complete sequence of partial backups that comprise the original save set

File-by-file recover is available only if there is a complete sequence of partial save sets that span the original save set. The directory structure is saved after the files are saved. If the directory structure is not available, then the browser cannot access the files.

To restore data from the complete sequence of partial save sets that make up the original save set:

1. Perform a query for all partial save sets. [“Querying partial save sets from the Console” on page 104](#) provides detailed information.
2. Use one of the following programs to restore the data:
 - For Windows - The NetWorker User program
 - For UNIX - The recover program

NOTICE

If the sequence of partial save sets is incomplete and does not make up the original save set, use the save set recovery procedure to recover the data from the partial save set. [“Recovering data from a partial save set” on page 107](#) provides detailed information.

Recovering data from a partial save set

Recovering from a partial NDMP save set differs from a non-NDMP save set:

- ◆ A save set recovery of a partial NDMP save set, recovers all of the partial save sets in the checkpoint sequence. The data in a partial save set cannot be recovered independently of the other partial savesets in the checkpoint sequence.
- ◆ A save set recovery of a non-NDMP save set enables you to recover data from a partial save set rather than by browsing and selecting individual files for recovery. However, the partial save set only contains successfully backed up files and not the entire set of data. An incomplete set of partial save sets cannot be browsed.

The procedure to restore data from partial save sets is the same as recovering by save set selection. [“Overview of NetWorker recovery methods” on page 372](#) provides detailed information on performing data recovery.

Use the **nsrinfo** command to display the contents of a partial save set. The **nsrinfo** man page or the *NetWorker 8.0 Command Reference Guide* provides detailed information about the **nsrinfo** command.

Cloning and scanning partial savesets

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.

Cloud backup devices and partial savesets

By default, the **CheckPoint restart** feature does not support cloud backup devices because partial save sets are not retained on the cloud backup devices.

Workaround

When the cloud backup device is used as a backup device for a Checkpoint restart operation, on the Server Properties menu, enable the Keep Incomplete Backups attribute.

If the Keep Incomplete Backups attribute is *not* enabled, the NetWorker software will not keep the partial savesets.

Deduplication backups

The *NetWorker Avamar Integration Guide* provides more information about Avamar deduplication backups.

The *NetWorker Data Domain Deduplication Devices Integration Guide* provides information on DD Boost deduplication backups using Data Domain storage systems.

Encrypting backup data

Backup and archive data on UNIX and Windows hosts can be encrypted with the **aes** Application Specific Module (ASM). The **aes** ASM provides 256-bit data encryption. Backup data is encrypted based on a user-defined pass phrase. If no pass phrase is specified, data is encrypted with a default pass phrase.

Microsoft Encrypting File System (EFS) and NetWorker AES encryption

Do not use NetWorker AES data encryption when backing up files that are encrypted by using the Microsoft Windows Encrypting File System (EFS). Even though the backup will be reported as successful, recovery of the file will fail and the following message will be written to the NetWorker log file: recover: **Error recovering <filename>. The RPC call completed before all pipes were processed.**

When EFS encrypted files are backed up, they are transmitted and stored on backup volumes in their encrypted format. When they are recovered, they are also recovered in their encrypted format. [“Encrypting file system” on page 886](#) provides more information about EFS encryption.

Set the Datazone pass phrase for a NetWorker server

To set the Datazone pass phrase:

1. From the **Administration** window, click **Configuration**.
2. Select the server name.
3. From the **File** menu, select **Properties**.
4. Click the **Configuration** tab and type a pass phrase in the **Datazone pass phrase** attribute.
5. Click **OK**.

NOTICE

By default, the current Datazone pass phrase is automatically used to recover password-protected files. If the current Datazone pass phrase was created after a password-protected backup was performed, you must provide the password that was in effect when a file was originally backed up. Keep password changes to a minimum.

Apply AES data encryption to clients in the datazone

To implement AES data encryption, apply the Encryption global directive to individual clients by using the Directives attribute of the Client resource. [“Editing a client” on page 606](#) describes how to edit a Client resource.

Compressing backup data

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.

NOTICE

Both compression and password-protection cannot be selected.

Applying compression to a scheduled backup

To apply compression to a scheduled backup, use the **compressasm** ASM in a local or global directive. Alternatively, use a preconfigured directive that is appropriate for the client computer's operating system. [“Directives” on page 289](#) provides information about creating and applying directives.

Applying compression to a manual backup

To compress data for a manual backup on Windows:

1. Open the **NetWorker User** program and click **Backup**.
2. Mark the data to be compressed.
3. From the **File** menu, select **Special Handling**. [“Special data handling for NetWorker clients on Windows” on page 109](#) provides more information.

To compress data for a manual backup on UNIX, you must use the **compressasm** in a local directive file. [“Directives” on page 289](#) provides more information.

Special data handling for NetWorker clients on Windows

To select directories and files for password-protection, encryption, and compression, or to clear a selected data item:

1. In the **NetWorker User** program, click **Backup**.
2. In the **Backup** window, select each data item to be designated for compression, encryption, or password-protection. If you select a disk volume or directory for an operation, the special handling will be applied to all of its nested subdirectories and files.
3. From the **File** menu, select **Special Handling**.
4. Select an option and click **OK**. You can also right-click the file and select a special handling options.

Depending on which special handling options were selected, the letter **P** (password-protection), **E** (encryption), or **C** (compression) appears next to the folder or filename.

To remove special handling, select a data item and click **Remove**.

Backing up Console server management data

To protect the Console server management data, such as report information, perform regular backups of the Console server database. The Console server database remains available during the backup.

NOTICE

The **savepsm** backup command backs up the Console server database and also backs up the database credential file and the authentication configuration files in a separate save set named `CONSOLE_BACKUP_FILES`.

Scheduling backups for the Console server database

If a NetWorker server was specified during the setup of the Console server in the Console Configuration Wizard, a Client resource was created to back up the Console server database on a scheduled basis. If a Client resource was created, modify the Client resource with respect to the backup schedule, browse and retention policies, and so on.

If a Client resource was *not* created to back up the Console server database, create a Client resource:

1. From the **Console** window, click **Setup**.
2. From the **Setup** menu, select **Set Database Backup Server**.
3. In the **NetWorker server** attribute, type the name of the NetWorker server that will back up the Console server database.
4. Select the **Create client** resource.
5. In the **Client name** attribute, type the name of the Console server.
6. Click **OK**.

A Client resource is created with the following attributes:

- **Name** attribute: the name of the Console server computer.
- **Save Set** attribute:

`NMCASA:/gst_on_server_name/lgto_gst`

where `<server_name>` is the short name of the host where the Console server component was installed.

- **Backup Command** attribute: **savepsm** (for a Windows Console server) or **savepsm.sh** (for a UNIX Console server).

One can also specify a NetWorker server to back up the Console server database through the Console Configuration Wizard. [“Accessing the Console Configuration Wizard” on page 535](#) provides more information.

[“Setting up a scheduled backup” on page 59](#) provides information about creating and tailoring a Client resource.

NOTICE

Only full, incremental, and skip backup levels are supported. All other backup levels (1-9) are mapped to an incremental backup.

Performing a manual backup of the Console database

Before performing a manual backup on UNIX, set the appropriate library path environment variable to:

- ◆ *Console_install_dir*/sybasa/lib64
- ◆ *Console_install_dir*/sybasa/lib on Linux

The environment variable to set varies by platform:

- ◆ Solaris/Linux: LD_LIBRARY_PATH
- ◆ AIX: LIBPATH

To perform a manual backup of the Console server database, type:

```
savepsm -I "Console_install_dir" save options
```

where *Console_install_dir* is the installation directory for the Console server.

For example:

- ◆ On Solaris, the default installation directory is /opt/LGTONmc
- ◆ On Linux/AIX, the default installation directory is /opt/lgtonmc
- ◆ On Windows, the default installation directory is C:\Program Files\EMC NetWorker\Management\GST

Installation directory paths that have spaces must be enclosed in quotations. For example:

```
savepsm -I "C:\Program Files\EMC NetWorker\Management\GST" save options
```

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provides information about the **savepsm** command.

Managing the size of the Console database log file

The Console database transaction log files are automatically truncated whenever a scheduled or a manual backup of the Console server database is performed.

To truncate the transaction log file without performing a backup, type:

```
savepsm -I "<Console_install_dir>" -T
```

Consideration for managing the Console database log file

If the transaction log file is truncated manually, the next Console server database backup that is performed after truncation *must* be a full backup. The next backup can be either a scheduled backup or a manual backup.

To ensure that the next scheduled backup of the Console server database is a full backup:

1. Open the backup schedule for the Client resource that is associated with the Console server database.
2. If necessary, perform a manual override on the next scheduled backup so that it is a full backup. [“Editing a schedule” on page 266](#) describes how to edit a backup schedule.

When performing a manual backup, a full backup is performed by default. [“Performing a manual backup of the Console database” on page 111](#) provides information about manual backups.

Backing up Windows mount points

A *volume mount point* (or mount point) is an NTFS file system feature of Windows 2008, 2008 R2, Server 2003, and Windows XP Professional.

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.

NOTICE

NetWorker backup and recovery of mount points require special handling, as explained in this section.

Including mount points in scheduled backups

To include mount points and their data in scheduled backups, you must specify the host volume, and each mount point. For example, to back up a single mount point on drive D: and all of its data, include this in the client’s Save Set attribute:

```
D:\mountpoint_name
```

To include nested mount points in scheduled backups, you can either specify save set All, 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:, include these in the client’s Save Set attribute:

```
D:\mountpoint_name1
```

```
D:\mountpoint_name1\mountpoint_name2
```

```
D:\mountpoint_name1\mountpoint_name2\mountpoint_name3
```

[“Scheduled backups” on page 58](#) provides for more information about setting up a scheduled backup. [“Directory specifications” on page 297](#) provides information about including mount points and nested mount points in a backup directive.

Performing a manual backup of a mount point and its data

To back up a mount point and its data:

1. Start the NetWorker **User** program.
2. Click **Backup**.
3. In the **Backup** window, expand the host drive that contains the mount point to back up, for example, drive D:\.
4. Under D:\, select the *mountpoint_name*.
5. Expand the *mountpoint_name* and verify that all data beneath it is selected for backup.
6. Click **Start**.

“Manual backups” on page 70 provides information about performing manual backups.

Performing a manual backup of nested mount points and their data

To perform a manual backup of nested mount points and their data, perform successive backup operations for each nested mount point and its data.

To back up three nested mount points and their data on drive D:\, for example:

1. Start the NetWorker **User** program.
2. Back up the top-level mount point and its data:
 - a. Click **Backup**.
 - b. In the **Backup** window, expand drive D:\ and mark *mountpoint_name1*.

NOTICE

When you mark a mount point for backup, all files, directories, and nested mount points beneath it are marked by default. Before starting the backup, make sure only *mountpoint_name1* and the files and directories beneath it are marked. You must unmark any mount points nested beneath *mountpoint_name1*.

- c. Click **Start** to run the backup.
3. Back up the second mount point and its data:
 - a. Click **Backup**.
 - b. In the **Backup** window, expand D:\ and *mountpoint1*.
 - c. Select *mountpoint_name2* and its data.

NOTICE

Be sure to clear (unmark) any mount points nested beneath *mountpoint_name2* before starting the backup.

- d. Click **Start** to run the backup.

4. Back up the third mount point and its data:
 - a. Click **Backup**.
 - b. In the **Backup** window, expand D:\, then expand *mountpoint_name1*, then expand *mountpoint_name2*.
 - c. Select *mountpoint_name3* and its data.
 - d. Click **Start**.

Backing up the Windows Content Index Server

The Windows Content Index Server (CIS) indexes the full textual contents and property values of files and documents 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.

Backing up CIS on Windows

The backup and recovery of the CIS occurs as part of the SYSTEM DB save set. If VSS is enabled, the CIS is automatically regenerated upon system reboot.

NOTICE

Back up the SYSTEM STATE and the SYSTEM DB save sets whenever a CIS database is created, moved, or renamed.

Before a CIS backup, the NetWorker software performs the following:

1. Pauses any CIS catalogs that are to be backed up.
2. Backs up all files that belong to those catalogs.
3. Turns the catalogs on again when the backup is finished. A catalog can still be queried when it is paused, so no indexing functionality is lost during the CIS backup.

The CIS deletes the catalog folder during a backup and restores it as part of a recovery operation.

Troubleshooting problems with a CIS backup

To troubleshoot a problem with a CIS backup:

- ◆ Ensure that the catalog folder is named **catalog.wci**.
- ◆ Restart the CIS.
- ◆ Ensure that the CIS was installed correctly.
- ◆ Pause or stop the catalogs, and then try the backup again.

Directing NetWorker software to skip CIS catalog backups

To skip backing up the *catalog.wci* folder entirely, create a Directive resource in the NetWorker Console by typing the following:

```
[intentionally leave first line blank in this directive]
<< / >>
+skip: *.wci
```

[“Directives” on page 289](#) provides information about directives.

Backing up Windows DHCP and WINS databases

Windows Dynamic Host Configuration Protocol (DHCP) and Windows Internet Naming Service (WINS) databases are not included in the NetWorker SYSTEM DB save set. However, you can use the procedures in this section to configure NetWorker software to protect these databases.

NOTICE

If VSS is enabled, the DHCP and WINS databases are automatically included when performing a backup of save set All. The procedures in this section are optional. [“Scheduled backups” on page 58](#) provides information about scheduling and recovering backups.

Backing up a DHCP database

To back up a DHCP database, ensure that this directory is included in the save sets for the NetWorker client that is the DHCP server, type:

```
%SystemRoot%\System32\dhcp
```

Backing up a WINS database

To back up a WINS database:

1. Use Microsoft WINS administrative tools to configure an automated backup of the WINS database to a local drive on the WINS server.
2. Ensure that the location chosen in [step 1](#) is included in the save sets for the NetWorker client that is the WINS server.

NOTICE

The Microsoft documentation provides information about the Microsoft WINS administrative tools.

Windows backup and recovery notes

This section contains notes about the backup and recovery of data on Windows clients, including:

- ◆ [“Enabling short filename support” on page 116](#)
- ◆ [“Enabling hard link support” on page 116](#)
- ◆ [“Failed backup and recovery attempts” on page 117](#)
- ◆ [“Granting full permissions for backup of Disk Quota database” on page 117](#)
- ◆ [“Native VHD volume support” on page 117](#)
- ◆ [“Recovery and case-sensitivity” on page 118](#)
- ◆ [“Security settings for logging operations performed by backup operator” on page 118](#)

Enabling short filename support

On Windows Server 2003 and Windows XP Professional, NetWorker software provides backup and recovery support for 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.s

To improve performance, support for short filenames is disabled by default.

To enable support for short filenames on a NetWorker client:

1. From the **Administration interface** window, click **Configuration**.
2. In the left pane, click **Clients**.
3. Right-click the appropriate client and select **Properties**.
4. Click the **Globals (2 of 2)** tab.
5. Select the **Short filenames** attribute.
6. Click **OK**.

Enabling hard link support

The NetWorker server backs up and recovers files with hard links. However, the hard links of files created by using a Portable Operating System Interface (POSIX) application are not preserved during recovery.

To improve performance, support for hard links is disabled by default.

To enable support for hard links on a NetWorker client:

1. From the **Administration interface** window, click **Configuration**.
2. In the left pane, click **Clients**.
3. Right-click the appropriate client and select **Properties**.

4. Click the **Globals (2 of 2)** tab.
5. Select the **Hard links** attribute.
6. Click **OK**.

Failed backup and recovery attempts

The NetWorker log file, located in <install_path>\logs\networkr.raw, contains a record of every file that was part of an attempted manual backup or recovery that was performed from the NetWorker User program. This file is overwritten with the next manual backup or recovery. If the file contains information that should be saved, you should rename the file or export the information by using the `nsr_render_log` program. [“Viewing log files” on page 803](#) provides information about viewing log files with the `nsr_render_log` program.

Granting full permissions for backup of Disk Quota database

NetWorker software backs up and recovers the Windows disk quota database as a component of the SYSTEM DB or VSS SYSTEM SERVICES save set. For any NetWorker client that uses the Windows Disk Quota feature, during SYSTEM DB or VSS SYSTEM SERVICES backup, NetWorker software creates temporary files to store the disk quota database settings in the root directory of each drive on the client.

If the permission settings for a local drive do not allow full control to the local system account, the disk quota database backup fails and an error message, similar to this, appears:

```
Failed to write to quota file, 0x80070005
```

To grant full permissions to the local system account:

1. Log in with administrator privileges to the NetWorker client host computer.
2. Using **Windows Explorer**, perform these steps for each local drive:
 - a. Right-click a drive icon.
 - b. In the **Properties** dialog box, select the **Security** tab.
 - c. Make sure the permissions settings allow full control to the system account.

NOTICE

By default, everyone has full permissions. If that setting has been changed such that the system account does not have full permissions, you must grant full permissions to the system account in order for the disk quota database to be backed up. For more information on setting permissions, refer to the Microsoft Windows documentation.

Native VHD volume support

NetWorker supports mounted, native VHD volumes on Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2, Windows Server 2012 and Windows Server 2012 R2. The VHD or virtual hard disk, is used as a mounted volume on designated hardware without any other parent operating system, virtual machine, or hypervisor. The volume can be used as a boot volume or as a data volume.

NetWorker has the following restrictions when using mounted, native VHD volumes:

- ◆ Native VHD volumes are not included in the ALL save set. You should configure a separate scheduled client to backup any native VHD volumes.
- ◆ Native VHD volumes should not be used as critical volumes if the volume that contains the native VHD is also a critical volume. This situation creates a conflict during a Windows Bare Metal Recovery backup.

Recovery and case-sensitivity

The NetWorker server is case-sensitive with regard to backup and recovery, although Windows file systems are not case-sensitive. This may result in the creation of multiple files with the same name but different cases.

For example, if you back up a file that is named temp.txt, delete it, then create a new file named Temp.txt, and then recover the old file, you will have two identical files in the directory -- one named temp.txt and the other named Temp.txt.

To avoid this problem, disable POSIX compliance by setting this system environment variable:

NSR_DISABLE_POSIX_CREATE=YES

The Windows online help contains detailed instructions about setting system environment variables.

Security settings for 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.

To enable NetWorker logging for Backup Operators, modify the security settings on the `<NetWorker_install_path>\logs` directory.

For example:

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**.

NetWorker logs operations performed by members of the Windows Backup Operators group.

Customizing the backup command

You can customize client backups by creating additional programs (scripts) that affect the way the NetWorker server backs up client file system data. For example, one can create a program that performs the following:

1. Shuts down a mail server or database *before* the NetWorker server performs a backup operation.
2. Restarts the mail server or database *after* the backup has completed.
3. Prints a message, such as, “Backup started at 3:33 A.M.”
4. Runs the backup.
5. Prints a message, such as “Backup completed at 6:30 A.M.”

You can customize a client’s scheduled backups in one of two ways:

- By creating a script that invokes the **save** program as part of its instructions. When the client is backed up, the customized program is invoked instead of the standard **save** program. [“Using the save command with a customized backup script” on page 119](#) provides more information.
- By typing **savenpc** in the Backup Command attribute of the Client resource. This way, the client backup invokes the **savenpc** program instead of the **save** program. The first time the client is backed up, **savenpc** creates a default backup program file, which you can then customize for future backups of the client. [“Using the savenpc command with a customized backup program” on page 124](#) provides more information.

Using the save command with a customized backup script

Additional processing instructions can be specified by entering the name of a custom script in the Backup Command attribute in the Client resource. The script is run instead of the default save program, when scheduled backups are initiated. The instructions in the script run separately for each save set backed up for the client.

When you use the save program, a new instance of the customized script is invoked for each save set listed in the client’s Save Set attribute, rather than just once for that client (as happens with **savenpc**). If you specify a save set value of All, the program is run for each file system on the client. As a result, if you create a Client resource with a customized backup for a database, a command to shut down the database is run for each save set that is listed.

NOTICE

After the creation of a customized backup script for a client, try backing up the client immediately. Any configuration or network problems that could prevent a backup should become apparent during this test.

The syntax used in the backup script or batch file must adhere to these criteria:

- ◆ The program name must begin with either the prefix *save* or *nsr* and must not exceed 64 characters.
- ◆ The program must reside on the client in the same directory as the NetWorker **save** command.
- ◆ Always specify the full path of the **save** command in the script.
- ◆ The NetWorker **save** command must be used in the backup program to ensure that the data is properly backed up.
- ◆ All commands within the program file must complete successfully. Otherwise, the NetWorker server cannot complete the remaining instructions.
- ◆ On UNIX and Linux, when you invoke the NetWorker **save** command, invoke the command with these arguments: **save "\$@"**. Doing so enables the **save** command in the script to accept the arguments usually passed to it by the NetWorker **savefs** program during a routine backup operation.

Create a custom backup script by using the save program

To create a custom backup script by using the **save** program:

1. Use a text editor to create a script in the directory where the NetWorker **save** command resides.

NOTICE

For custom backup scripts that are to run on Windows clients, the script name must start with *save* or *nsr* and must end with the *.bat* extension.

Commands in this script must be placed in this order:

- a. Run a preprocessing command before each save set backup (optional).
 - b. Back up the data by using the NetWorker **save** command (mandatory).
 - c. Run a postprocessing command after each save set backup (optional).
2. In the **Backup Command** attribute of the **Client resource**, type the name of the backup script.
 3. Back up the client to ensure that the newly created backup command works.

Example 3 The save Backup command on Windows

In this example, for each save set, the customized backup program runs pre-backup commands, runs the NetWorker **save** command, and then runs post-backup commands.

The backup program consists of these parts:

- ◆ **Pre-Backup Command:** Redirects the output of the Net start DOS command to create a *netstart.txt* file at the root of the C: drive and send all current computer Services Started information to this file.
- ◆ **Save:** Runs NetWorker commands required to start the backup process.

- ◆ **Post-Backup Command:** Redirects the output of the Set DOS command to a set.txt file at the root of the C: drive and send 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.

Also, you can check the batch file execution process in the Monitor Groups tab of the Administration window, or by viewing the savegrp log file in the <NetWorker_install_path>\logs directory.

[“Monitoring NetWorker server activities” on page 462](#) provides information about viewing execution file details from the **Monitor Groups** tab.

[“Viewing log files” on page 803](#) provides information about viewing log files.

This is an example backup 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

```

This information is displayed from the Monitor Groups tab, logged to the savegrp.log file after the backup process is done, and verifies the execution of the three processes.

```

--- 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=====
jupiter: index:jupiter level=full, 243 KB 00:00:00 23 files
jupiter: bootstrap level=full, 47 KB 00:00:00 7
files
* jupiter:bootstrap nsrldr: Either a printer isn't
* defined for printing the Bootstrap for
* this savegroup,
* jupiter:bootstrap 04/26/06 01:34:13 PM full
* 3901113601 3901113601 0
jupiter.001

```

Example 4 The save backup command on UNIX

This script backs up a ClearCase version object base (VOB). The script file must reside in the same directory as the NetWorker **save** command (for example, on a Solaris system, the **save** program is installed in the /usr/sbin directory). Type the name of the script into the Backup Command attribute of the Client resource that is used to back up the ClearCase VOB. As a result, this script is invoked instead of the usual **save** command during a scheduled backup.

NOTICE

Include the **save** command in the script and place the script in the same directory as the **save** program. Otherwise, the backup will fail.

This script locks a ClearCase 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
```

Controlling the custom backup script's exit status reporting

You can use the Job Control attribute in the client resource to control how end of job and exit status messages are determined for the custom script. This can be useful for debugging and reporting purposes.

By default, the programs **savegrp** and **nsrjobd** determine a custom script's success or failure based on the **save** program's completion (end of job). The following criteria apply:

- ◆ If the **save** job's completion status is *success*, then **savegrp** and **nsrjobd** report that the custom backup job succeeded.
- ◆ If the **save** job's completion status is *failure*, then **savegrp** and **nsrjobd** report that the custom backup job failed.

- ◆ If no completion status is received, the custom job's output is examined for `completed savetime=savetime` lines. If found and the savetime is non 0 (zero), the custom backup job is considered to have succeeded, otherwise it is considered to be failed.
- ◆ The exit code of the custom script's process is not taken into consideration.

The Job Control attribute in the Apps and Modules tab of the client resource enables you to change how end-of-job and success or failure messages are determined for a custom script. The Job Control attribute has three options that can be selected singly or in combination. The following table describes these options.

Table 15 Job Control selections

Job Control selections	Description	Uses
End on job end	A backup job is considered to be ended as soon as an end job message is received from the save command.	Use when you do not want to wait for the post processing commands of the script to end.
End on process exit	A backup job is considered to be ended as soon as the started process exits. Note: Background processes started by the backup command could still be running on the client.	Use when you want your custom script to start background processes and you do not want savegrp or nsrjobd to wait for their completion.
Use process exit code	Only the process exit code is used to determine the success or failure of the job. An exit code of 0 means success, otherwise the job is reported as failed.	Use when you want the script post processing command status to have an impact on the status of the save backup command without having to unset the <code>NSR_STD_MSG_FD</code> environment variable. Note: If your script invokes more than one NetWorker backup command such as save , you must still unset the <code>NSR_STD_MSG_FD</code> environment variable.
End on job end and End on process exit	Either event can trigger a job's end.	
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 job's success or failure.	

Using the **savenpc** command with a customized backup program

As an alternative to using the **save** program with a custom script, use the **savenpc** program. The **savenpc** program differs from using a custom script with the **save** program in that preprocessing and postprocessing commands run only once during the client backup, instead of once for each save set. This command can be useful if the client is running a database or other program that should be stopped before the client is backed up, and then restarted after the backup has completed. The options for the **savenpc** command are identical to those for the **save** command.

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provides For information about the **savenpc** command.

To run the **savenpc** program:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Clients**.
3. Create a new Client resource or select an existing Client for editing.
4. Select the **Apps & Modules** tab.
5. In the **Backup Command** attribute, type:

```
savenpc
```

6. Back up the client.

The first time a backup group with a client that uses **savenpc** runs, a standardized *group-name.res* file is created in one of the following:

- The */nsr/res* (UNIX)
- The *NetWorker_install_path\res* (Microsoft Windows) directory on the client

where *group-name* is the same as the name in the Group resource selected for that client. If the client belongs to multiple backup groups, a separate *group-name.res* file is created for each group to which the client belongs.

The initial *group-name.res* file contains **type**, **preprocessing**, **postprocessing**, **timeout**, and **abort precmd with group** attributes:

```
type: savenpc;
precmd: "echo hello";
pstcmd: "echo bye";
timeout: "12:00pm";
abort precmd with group: No;
```

NOTICE

The **abort precmd with group** attribute determines what will happen to the preprocessing command when the **savegroup** process aborts prematurely. By default, the preprocessing command process will not be killed if the **savegroup** process aborts prematurely. To kill the preprocessing command when the **savegroup** process aborts, set **abort precmd with group** to **Yes**.

When the *group-name.res* file exists, use a text editor to customize the file's attributes. These customized instructions are then applied the next time the client is backed up.

Before performing a save operation on the client, the modified **savenpc** program performs the following:

- ◆ Any preprocessing commands listed for the **precmd** attribute in the *group-name.res* file.
- ◆ The save by using the options specified for the **savenpc** command itself.
- ◆ Any postprocessing commands listed for the **pstcmd** attribute.

When editing a *group-name.res* file, these points apply:

- ◆ The command environment that is opened by the **savenpc** command to run a customized backup does not automatically inherit the system's default environment. Specifically, environment variables, including PATH, will either not exist or will be set to NULL. The environment must be built as part of the preprocessing (**precmd**) commands, especially the PATH variable. On UNIX clients, be sure to source the **.profile**, **.cshrc**, and other login scripts.
- ◆ The **save** command should not be specified in the *group-name.res* file. The **savenpc** program will automatically invoke the **save** command and back up the save sets specified in the **Save set** attribute for the client.
- ◆ To exclude the environment variables in the *group-name.res* file, include full pathnames for all commands and files.
- ◆ Resident commands, for which there is no executable file present, like **time** and **dir**, will not work as commands in your *group-name.res* file. The log reports that the executable file could not be found.
- ◆ On a Microsoft Windows client, do not use "@ECHO OFF" in the *group-name.res* file.
- ◆ To add more than one command sequence to the **precmd** and **postcmd** attributes, insert a comma (,) to separate the commands.
- ◆ A complete command-line for an attribute must end with a semicolon (;).
- ◆ Escape any backslash (\) characters in the *group-name.res* file. For example, the pathname C:\mydir\myprogram.exe must be written C:\\mydir\\myprogram.exe.

This is an example of a fully functional *group-name.res* file:

```
type: savenpc;
precmd: "V:\\usr\\sap\\PDB\\SYS\\exe\\run\\PDB-stop.cmd >
C:\\WINNT\\system32\\PDBStop.log 2>&1";
postcmd: "V:\\usr\\sap\\PDB\\SYS\\exe\\run\\PDB-start.cmd
C:\\WINNT\\system32\\PDBStart.log 2>&1";
timeout: "12:00pm";
```

It is not necessary to escape any backslash characters in scripts called from the *group-name.res*. To simplify the pathname issue, include all commands in a script or batch file, and then include that script's full pathname on the **precmd** or **postcmd** line.

- A line break is required after the semicolon that ends the last command in the *group-name.res* file.
- The following applies to text written to standard output:
 - Text written during preprocessing appears in the NetWorker completion notices. You can direct this output to a log file.
 - Text written during postprocessing is discarded. Consider redirecting this output to a log file so you can troubleshoot problems.

Timeout attribute

The **Timeout** attribute indicates when the postprocessing commands are to be run, regardless of whether all of the save sets have been backed up or not. The timeout entry must be specified in **nsr_getdate** format and must be enclosed in double quotes. For more information about **nsr_getdate**, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

If an invalid time is entered for the timeout, the timeout is not reached and an error message is not produced.

The Timeout attribute is optional. To disable the **Timeout** attribute, add a comment character (#) to the beginning of the line, for example:

```
# timeout: "12:00pm";
```

The value of the Timeout attribute may not be the exact time that postprocessing actually commences. The **savepnpc** program's **pstclntsave** subroutine uses a one-minute polling interval to check for the completion of preprocessing tasks. Therefore, the **savepnpc** log file may show that postprocessing was started up to 60 seconds after the designated timeout.

Customize the savepnpc command for multiple groups

To customize pre- and postcommand processing for multiple groups:

1. Copy existing *group-name.res* files.

- Microsoft Windows clients:

NetWorker_install_path\tmp\group-name.res to *NetWorker_install_path\res\your_new_group.res*

- UNIX clients:

/nsr/res/group-name.res to */nsr/res/your_new_group.res*

2. Edit the new *new_group-name.res* file.

If you do not have an existing *group-name.res* file, activate the group for **savepnpc** without the presence of this file. A default template will be created at one of the following:

- */nsr/res/your_new_group.res*
- *NetWorker_install_path\res\your_new_group.res*

You can then customize the template.

Message logging by the savepnpc command

Messages generated by **savepnpc** are written to the **savepnpc** log file, located in these locations on the NetWorker client:

- ◆ UNIX: */nsr/logs*
- ◆ Microsoft Windows: *NetWorker_install_path\logs*

The format of the **savenpc** log file is similar to:

- ◆ 04/03/07 13:56:43 preclntsave: All command(s) ran successfully.
- ◆ 04/03/07 13:57:43 preclntsave: All save sets on the worklist are done.

[“Viewing log files” on page 803](#) provides information about viewing log files.

Considerations for backing up 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 managed by an active database management system (DBMS), ensure that the partition is offline and the database manager is shutdown. For greater flexibility when backing up partitions that contain DBMS data, use a NetWorker Module application.

Backing up raw partitions on UNIX

To back up raw disk partitions on UNIX, use the **rawasm** directive. [“Precautions when using rawasm to back up UNIX raw partitions” on page 302](#) provides more information.

Backing up raw partitions on Windows

To back up a raw disk partition on Windows, specify the raw disk partition in a save set. Identify the raw partition as a physical drive or as a logical drive. For example:

```
save -s NetWorker_server_name \\.\PhysicalDrive0
save -s NetWorker_server_name \\.\C:
```

Backing up a mapped drive

To back up a mapped drive, follow these guidelines:

- ◆ To specify a drive to back up in either a scheduled or manual backup, do not specify the drive letter. Instead, specify the Universal Naming Convention (UNC) path.

For example, to specify the accounts directory on the server *jupiter*, type:

```
\\jupiter\accounts
```

- ◆ For scheduled backups:
 - Add the username required to access the UNC path to the **Remote User** attribute in the **Client** resource.
 - Add the password required to access the UNC path to the **Remote Password** attribute in the **Client** resource.

Backing up access control lists

The backup and restore of ACLs (Access Control Lists) and extended ACLs is fully supported. This support covers Linux, HP-UX, AIX, DEC, SOLARIS, OS/X, and Windows.

There is no special attribute or keyword that controls this support. When a file that has an associated ACL is backed up, the ACL is backed up along with the file data. When the file is recovered, any associated ACLs will also be recovered.

However, in order to recover files with associated ACLs, **ACL passthrough** must be checked in the **Recover** section in the **NetWorker Server Properties** window.

Backing up BOOT/BCD Data on Windows

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, 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's 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, 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 saveset name is GLOBALROOT/xxxxxx/.

The BCD can only be restored as part of a NetWorker Windows 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. [Chapter 25, "Windows Bare Metal Recovery,"](#) provides more information about planning and performing a NetWorker Windows offline BMR.

Support for backing up renamed directories

As of NetWorker 8.0, the option to back up renamed directories is enabled by default. The Backup renamed directories attribute can be disabled or enabled for each NetWorker client. ["Editing a client" on page 606](#) provides information about editing a Client resource.

Consider the following when deciding whether to disable or enable this feature:

- ◆ This feature must be enabled for NetWorker clients that use the synthetic full backup feature. ["Synthetic full backups" on page 76](#) provides more information.
- ◆ When this feature enabled and a renamed directory is encountered, a full backup is performed on all sub-directories and files under a renamed directory.

- ◆ When this feature is enabled and if a renamed directory is at some future date given its original name, files and subdirectories under that directory will not be eligible for backup until the files or subdirectories are updated or the next full backup occurs.
- ◆ When this feature is disabled, unchanged files and folders under the renamed directory will be skipped during a non-full level backup. This behavior can cause unexpected results during a recovery operation. If you attempt to recover data under a renamed directory from a date between the time that the directory was renamed and the next full level backup, it may appear that data is missing. For that recovery time period, any files or folders that were unchanged will not display under the renamed directory. Instead, they will be displayed under the old directory name.

Backing up only client file indexes and the bootstrap

You can set up a backup group to backup only the client file index information for those NetWorker clients that belong to the backup group. The bootstrap will also be backed up.

To backup only the client file indexes and the bootstrap:

1. Set up a backup group as described in [“Task 2: Set up a group for backup clients” on page 61](#)
2. In the backup group properties Advanced tab, select **Index only** from the **Options** attribute.
3. Click **OK**.

CHAPTER 3

Storage Nodes and Libraries

This chapter covers these topics:

◆ Storage nodes.....	132
◆ Configuring storage nodes	133
◆ Dedicated storage nodes	139
◆ Troubleshooting storage nodes	140
◆ Avamar deduplication nodes and replication nodes	140
◆ Devices and libraries.....	140
◆ Autodetection of libraries and tape devices.....	141
◆ Configuring libraries.....	144
◆ Library reconfiguration	146
◆ Specifying available library slots	147
◆ Miscellaneous library operations.....	148
◆ Tips for using libraries.....	150
◆ Library maintenance	152
◆ Deleting libraries.....	155
◆ Troubleshooting autoconfiguration failure	155
◆ Silo libraries	156
◆ Configuring silo libraries	158

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. [“NetWorker User Groups” on page 559](#) provides more information.

Requirements

To operate the NetWorker software with storage nodes, the following 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
 3. (Optional) EMC AlphaStor software.

This enables multiple NetWorker servers to share the storage node. The AlphaStor software must be used to manage the libraries, drives, and volumes on that node. The AlphaStor server is available on Solaris and Microsoft Windows only. AlphaStor DCP/LCP is available on all UNIX, Linux, and Windows platforms. Refer to the *EMC NetWorker Software Compatibility Guide* for information.

- ◆ On Windows systems, the Storage Node Option must be installed. This installs both the NetWorker client and storage node software.

Licensing

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

Configuring storage nodes

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

Configure a storage node

To configure a NetWorker storage node:

1. Ensure 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 a type:
 - **SCSI**
 - **NDMP**
 - **SILO** (or SILO with NDMP) The silo robot arm must not be detected by the NDMP storage node.
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 cancelled.
 - 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 **mmds for disabled devices**, select a nsrmmd (data mover) option (see note):
 - **Yes** to start nsrmmd processes for disabled devices.
 - **No** to *not* start nsrmmd processes for disabled devices.

- 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 **Target Sessions** is reached.
 - Unselected (static mode): NetWorker runs all available nsrmmd processes.

NOTICE

In environments where unattended firewall ports need to be restricted for security reasons, the storage node settings for **mmds for disabled devices** and **Dynamic nsrmmds unselected** (static mode) offer more control because they cause all available **nsrmmd** firewall ports to be attended by running **nsrmmd** processes.

[“Create and configure an AFTD” on page 167](#) provides details on device settings for session load balancing.

- 7. In **Remote Host**, if an NDMP tape library is used by this storage node, type the **Remote User** name and **Password**. Only one user is allowed per storage node.
- 8. Select the **Configuration** tab.
- 9. 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:
 - **Yes** for NetWorker to detect all LUNs (logical unit numbers). Detection can take a long time.
 - **No** (default) for NetWorker to stop searching at the first available LUN.
 - 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.

[“Persistent binding and naming” on page 200](#) provides details.
 - d. In **Skip SCSI targets**, list any SCSI targets to exclude from backup operations, one per line, if the storage node type is set to SCSI. The format is bus.target.lun where the target and LUN fields are optional. You can exclude a maximum of 63 targets.
- 10. In **Advanced Devices**, for AFTD or DD Boost devices, configure the settings:
 - 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 this attribute has no value, then the NetWorker server's storage node **Clone storage nodes** attribute is used, and if that has no value then the NetWorker client's **Storage nodes** attribute is used.

NOTICE

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. [“Specifying a clone from a volume shared by multiple devices” on page 356](#) provides details.

11. When finished, click **OK**.

The new storage node appears in the navigation tree.

Modifying the timeout attribute for storage node operations

An attribute named **Nsrmmmd 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.

Other attributes involved in storage node timeouts include:

- ◆ Nsrmmmd polling interval, which determines the number of minutes between storage node polls.
- ◆ Nsrmmmd restart interval, which determines the number of the minutes the NetWorker software waits before restarting the nsrmmmd process. A value of zero for the **Nsrmmmd restart interval** attributes indicates an immediate restart.

To modify these attributes:

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**.

Configure 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 a device's Properties.

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 to 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 these attributes:

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**.

Configure the client's storage node affinity list

The choice of which NetWorker servers and storage nodes receive a client's data—known as *storage node affinity*—is made by entering their hostnames in the Storage Nodes attribute located in the Client Properties, on the Globals (2 of 2) tab. The default setting for the Storage Nodes attribute on most Client resources is nsrserverhost (the host NetWorker server).

If the Client resource of a storage node computer is created after a remote device on the storage node has been created, the default setting of the Storage Nodes attribute is the storage node and the NetWorker server.

If you create a Client resource after you create a storage node, and you will configure the client to back up to that storage node, enter the name of the storage node in the Storage **Nodes** attribute of the Client resource *above* the default nsrserverhost. You can add Storage node names to the Storage Nodes attribute list at any time. The NetWorker software directs the client data to the first storage node in the 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 list based on the criteria specified in [“Storage node load balancing” on page 137](#).

To modify the Storage Nodes attribute:

1. In the server's **Administration interface**, click the **Configuration** button.
2. Select **Clients**, right-click the appropriate client and select **Properties**.
3. Select the **Globals (2 of 2)** tab.
4. Modify the **Storage Nodes** attribute as appropriate and click **OK**.

Storage node load balancing

Starting in NetWorker 8.1, a new feature, named Save session distribution, has been introduced that allows one to configure how save sessions are distributed among storage nodes.

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

The Save session distribution feature can be applied to all NetWorker clients globally or to selected clients only. This feature has two options:

- ◆ **max sessions**
Save sessions are distributed based on each storage node device's **max sessions** attribute. This is the default distribution method.
- ◆ **target sessions**
Save sessions are distributed based on each storage node device's **target sessions** attribute. This option is more likely to spread the backup load across multiple storage nodes, while the max sessions option is more likely to concentrate the backup load on fewer storage nodes.

[“Backup to Disk and Cloud” on page 161](#) provides more details on device target session and max session attributes.

When the max sessions option is selected, NetWorker client save sessions are distributed among eligible storage nodes as follows:

1. Identify the available storage nodes in the NetWorker client's Storage node affinity list.
2. Use 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 use the least loaded device.
4. When all devices on the first storage node are running at their max sessions level, continue to the next storage node and repeat the device selection process described previously in steps 2 and 3.
5. Continue until all available devices on all storage nodes in the client's storage node affinity list are in use.

When the target sessions attribute is selected, NetWorker client save sessions are distributed among eligible storage nodes as follows:

1. Identify the available storage nodes in the NetWorker client's Storage node affinity list.
2. Use 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. Continue with step 4 until all devices on all available storage nodes are running at their max session levels.

Specifying storage node load balancing

By default, NetWorker balances client backups across storage nodes 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. [“Overriding the save session distribution method for selected clients” on page 138](#) provides more information.

To apply storage node load balancing settings:

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 **Setup** tab.
5. Select a value from the **Save session distribution** list.

If you select **target sessions**, then all NetWorker clients will have their backups balanced across storage nodes based on device target session values. The Save session distribution attribute on each NetWorker client resource is ignored.

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.

6. 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 override this setting for selected clients.

To override the max sessions distribution method for a NetWorker client:

1. In the server’s **Administration interface**, click the **Configuration** button.
2. Select **Clients**.
3. Right-click the appropriate client and select **Properties**.
4. Select the **Globals (1 of 2)** tab.
5. Modify the **Storage Nodes** attribute as appropriate and click **OK**.
6. Select **target sessions** from the **Save session distribution** list.
7. Click **OK**.

Performance considerations for storage node load balancing

Be aware of the following performance considerations for storage node load balancing:

- ◆ Depending on how your backup environment is configured, there is a potential to shorten backup times by using the device target session option rather than the device maximum session 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.
- ◆ Each NetWorker client has a storage node affinity list. The Save sessions distribution feature can only distribute a NetWorker client's backup sessions to multiple Storage nodes if the client 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. [“Configure the client's storage node affinity list” on page 136](#) provides more information.

Bootstrap backup on a storage node

When the server's bootstrap save set is backed up, the data writes to a device that is local to the NetWorker server. A bootstrap cannot be backed up to a remote device, but a bootstrap can be cloned or staged to a remote device. When the **mmrecov** command is used to recover a bootstrap save set, the data must be recovered from a local device.

Staging bootstrap backups

Bootstrap backups can be directed to a disk device such as an AFTD or FTD device. However, if a bootstrap backup is staged to another device, the staging operation will complete and will be reported as complete even though the “recover space” operation will not be executed. This means that the staged bootstrap will remain on the original disk from which it was staged. Therefore, the original disk can be used to scan in the bootstrap data if the staged bootstrap is accidentally deleted. Also be aware that if the bootstrap data is not staged from the original disk, the data on the original disk will be subject to the same browse and retention policies as any other save set backup and will, therefore, be subject to deletion after the retention policy has expired.

This bootstrap information also applies to NDMP devices.

Dedicated storage nodes

All devices created on storage nodes (except servers) include the Dedicated Storage Node attribute. A dedicated storage node can back up only its own, local data.

Set this attribute when a device is created on a remote storage node. It is found in the device's Properties, on the Configuration tab. If the Dedicated Storage Node attribute is set to Yes, a Dedicated Storage Node License is required for the storage node. If, however, the Dedicated Storage Node attribute is set to No (the default value), a standard storage node license is required. The Dedicated Storage Node License also can be used for backing up virtual clients in a cluster.

NOTICE

A storage node host cannot mix storage node types. Either all devices on a storage node must be set up for a dedicated storage node, or all must be set up for a standard storage node.

In NetWorker release 7.6, NetWorker supports the installation of a dedicated storage node in a Solaris 10 local zone to backup directly to a physically attached device without sending data across the IP network. NetWorker allows sharing of a device between multiple dedicated storage nodes that are installed in multiple local zones of a single physical host, assuming all the storage nodes belong to a single NetWorker data zone.

Troubleshooting storage nodes

If a backup fails, this message might appear:

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

The problem could be related to storage node affinity.

Possible causes 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 client's list.
- ◆ 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.

Avamar deduplication nodes and replication nodes

Deduplication nodes and replication nodes exist on Avamar servers. Contact EMC Customer Support to configure these nodes on the Avamar server side. Once that has been done, you can create access to them from the NetWorker side.

The *NetWorker Avamar Integration Guide* provides information on how to create a NetWorker deduplication node.

Devices and libraries

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

SCSI 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.

In all cases, the robotic controller and associated tape devices are all controlled through a SCSI interface which is available on one or more storage hosts.

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.

Silo libraries

Silos libraries have a robotic controller that moves tape media between slots and devices. However silos do not use a SCSI interface to access and control the media movements. The movements are controlled by a separate host that receives requests over the network.

[“Silo libraries” on page 156](#) provides more information.

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 512. 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.

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:

- ◆ 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 899](#) 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 143](#) provides more information.

To scan for available devices:

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.

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.

To specify a specific network interface for device scan operations:

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**.

[“Server Network Interface attribute” on page 150](#) provides information on similar attributes that are available for other library and device operations.

Refreshing enterprise library views on request

To update enterprise library views on request:

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.

To change the update interval:

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**.

Configuring libraries

A library resource must be created for each library, including silos, on a storage node. Because a NetWorker server is also a storage node, this procedure applies to a NetWorker server and all of its storage nodes. You can configure a library either automatically with the configure all libraries wizard or manually with the user interface.

A storage node must be created before devices can be configured to be used by them. [“Storage nodes” on page 132](#) provides details. All scanning for devices is performed at the storage node level and can be performed 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 (the **inquire** or **sn** commands can be used to determine if a device returns a serial number).

Devices must be updated to the most recent firmware and drivers.

The following library types can be automatically configured:

- ◆ SCSI
- ◆ NDMP
- ◆ Silo (except DAS silo)

The following device types must be configured by using the **jbconfig** command:

- ◆ AlphaStor devices.
- ◆ IBM tape libraries controlled through the use of IBM’s tape driver. (This is because the device autodetection code uses the internal lus driver to control libraries.)
- ◆ Any library that does not return a serial number for the robotic arm or any of its tape devices.

Adding a library resource

To automatically configure a new library resource for a storage node:

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.

Configuring a virtual tape library (VTL)

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 *EMC NetWorker Licensing Guide* provides information about NetWorker licensing support for a Virtual Tape Library.

Queuing device resources for AlphaStor

Because the NetWorker software detects devices as virtual devices, users can request more devices than actually exist. The AlphaStor software queues these requests, and can prioritize them according to whether a tape is mounted for reading, or for writing. This allows AlphaStor users to prioritize recovery operations above backups or other operations that might compete for the same devices.

This feature requires AlphaStor release 3.1 or later.

Configure the AlphaStor library with the **jbconfig** command. For information about configuring resource queuing, refer to the *EMC AlphaStor Administration and Operator's Guide*. The **nsr_mount_request** man page describes the resource-queuing feature. Related attributes are also in the **nsr_pool** and **nsr_jbox** UNIX man pages and the *NetWorker Command Reference Guide*.

Library reconfiguration

Configure NetWorker privilege is required to reconfigure a library or to add or remove access paths to the devices in a library. This includes access paths that allow libraries to be shared.

Considerations when reconfiguring a library:

- ◆ The reconfiguration of stand-alone or file type devices is not supported. Instead, delete the stand-alone or file type device, and then create a new one.
- ◆ 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. [“Using the jbedit command to configure a library” on page 147](#) provides details.

Reconfiguring a library

To reconfigure a library:

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 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.

Using the `jbedit` command to configure a library

The `jbedit` (jukebox edit) program can be used as a fallback means of editing library configurations if the autoconfiguration program cannot be used. 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 privileges.

The `jbedit` program supports all direct-attached SCSI/SJI, SAN, and NDMP libraries. It does not currently support AlphaStor 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 [“Library reconfiguration” on page 146](#). The `jbedit` options provide selection lists that make it easy to find drives or devices to be added or deleted.

[Table 16 on page 147](#) lists the most commonly used `jbedit` program options.

Table 16 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 *EMC NetWorker Command Reference Guide* or the UNIX man page provides a detailed description of the `jbedit` command, its options, and associated diagnostic messages.

Specifying available library slots

The available slots feature controls which volumes the NetWorker server uses for backup. The server uses all of the 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.

Specify library slots

To configure which slots are available in a library:

1. Ensure that volumes have been placed in all the available slots of the library so that the NetWorker server can proceed 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 is a total of 64 sides and, therefore, 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.

Miscellaneous library operations

This section covers various additional library operation topics.

Sharing libraries among NetWorker hosts

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. [Figure 12 on page 149](#) shows how multiple NetWorker hosts can share library devices.




Figure 12 How library sharing works

Sleeping periods for library tasks

Library resources include attributes 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 [Table 17 on page 149](#).

Table 17 Library resource sleep attributes (1 of 2)

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 a an eject operation to complete.	60 seconds

Table 17 Library resource sleep attributes (2 of 2)

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 will allow a device with a volume to be idle before automatically unmounting it. For specific devices, this value can be overridden. “Automatic unmounting of volumes (idle device timeout)” on page 231 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.

Tips for using libraries

This section provides additional suggestions for using libraries effectively and reliably.

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.

“[Notifications](#)” on page 479 provides more information about notifications.

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.

Resetting a library

A library must be reset each time the library and the NetWorker software become out of sync. [Example 6, “Host crash requires user intervention”](#) provides details. A library reset can be done using either the Administration interface or the command prompt.

Reset a library in the Administration interface

To reset a library in the Administration interface:

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**.

Reset 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.

To make the NetWorker software aware of these new drives, execute **nsrjb -HE** to reset the library. The **-E** option reinitializes the library’s element status. Some libraries can keep track of 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 (**sjj** 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 *EMC NetWorker Command Reference Guide* or the UNIX man pages.

Using 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. [“Media pools” on page 304](#) provides more information.

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.

If it is necessary to use the library's front panel display to add and remove volumes, do the following:

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 233](#) has information about inventorying libraries.

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.

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

To set the cleaning attributes of a library:

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**.

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.

Delaying tape device cleaning

Occasionally it is necessary to delay the cleaning of a tape device which is scheduled for cleaning.

To set the value for cleaning delay:

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. To use this feature, select **enabled** for the **Cleaning** attribute of the **Device** resource so that automatic cleaning is enabled.

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. [“Common device interface” on page 199](#) provides more information about CDI.

Devices that are capable of TapeAlert perform constant self-diagnostics and communicate the diagnostic information via the **nsrmmmd** program to logs that can be viewed in the Monitoring task.

TapeAlert attributes are found in the device's Properties, on the Volume tab. Their respective descriptions are as follows:

- ◆ 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.

[Table 18 on page 154](#) describes the nature of the tape alert levels.

Table 18 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 **nsrmmmd** once the reported issue is handled.

Critical messages about hardware errors are not cleared by **nsrmmmd** because they might indicate intermittent hardware problems.

Deleting libraries

To delete a library from a storage node:

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.

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.

Troubleshooting autoconfiguration failure

Common symptoms of library autoconfiguration failure include the following:

- ◆ 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.

To troubleshoot this problem:

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.

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. Silos are controlled by silo management software, which is provided by the silo vendor and installed on a silo server. 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.

[“Media management in a silo” on page 237](#) provides information on silo-specific, media-management concerns.

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.

Installing a silo

To install a silo for use with NetWorker software:

1. Install the silo management software on the silo server.
2. Install the **STLI** on the NetWorker server, if required. 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 on the following models, because all the necessary software is installed when the NetWorker software is installed:

- IBM 3494 on Solaris and AIX
 - StorageTek on Solaris, AIX, and HP-UX
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 158](#) provides details.

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

The StorageTek (STK) silo management software uses either a program called ACSLS that runs on a UNIX system, or a program 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.

For tape drives, the name consists of four digits 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.

Ask the silo administrator for the drive names of the devices that the NetWorker server can use. There is no way to get this information from the NetWorker server. To connect to more than one drive, determine the SCSI IDs for each drive and properly match the IDs to the silo names. If the operating system device names and silo names are accidentally swapped, it is only possible to mount and unmount volumes. Volumes cannot be read or written to after they are mounted. To reconfigure the device names properly, use the Administration program to change the order of the device names in the STL Device Names attribute of the library’s Properties.

IBM 3494

The silo management software for the IBM 3494 names devices with an eight-digit number to identify the 3590 drives in the silo. Use the appropriate utility to obtain the device names, as follows:

- ◆ On an AIX system, the NetWorker software obtains the name of the device from the device driver and displays the device name as the default value.
- ◆ On a Solaris system, the IBM-supplied **mtlib** command (**mtlib -l library_name-D**) must be used to determine the names of all the devices in the 3494, if the silo name is configured by using the **jbconfig** command from the command-prompt, rather than through the configuration interface. Either ask the silo administrator which device is reserved for the NetWorker software, or test to determine which silo drive name matches with each Solaris device name.

Configuring silo libraries

To configure a new silo resource to a storage node automatically:

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.

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 this input (some of which is filled in by default):

- Library type (select **STL Silo**).
- Adjust the **Enable New Device** option, if necessary.
- Current server sharing policy. (Use maximal sharing with Dynamic Drive Sharing (DDS).)
- Storage nodes on which the libraries should be configured. You can select a storage node to see its details displayed; 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 name of the new storage node:

- a. Update storage node properties if required.
- b. Enter 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 entered for each one.
- c. Enter the **Hostname** of the silo controller.

- d. Enter the **Type of silo** controller. The default is ACSLS Silo (StorageTek).
 - e. (Optional) Use the **Test Silo Controller Connectivity** button to see whether the connection to a silo controller works. Use it once for each silo. An error message appears if the connection to a given silo fails.
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.

Using 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 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* 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.

To release a silo device:

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 [step 1](#) through [step 6](#) for each device to be released.

Cleaning silo devices

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 silo manufacturer's software documentation.

Environment variables for StorageTek silos

Environment variables must be set for StorageTek silos. [Table 19 on page 160](#) lists the environment variables to set.

Table 19 StorageTek environment variables

Silo model	Environment variables
StorageTek	<p>For UNIX systems:</p> <ul style="list-style-type: none"> • <code>CSI_HOSTNAME = name_of_ACSLT_system</code> <p>The following commands should also be running on the system and can be included in the NetWorker startup script:</p> <ul style="list-style-type: none"> • <code><binaries_path>/mini_el &</code> • <code><binaries_path>/ssi &</code> <p>For Windows systems:</p> <p>The LibAttach Configurator program is available from StorageTek. It creates an ssi process, and a link is available to start the mini_el process from Start > Programs > LibAttach menu tree.</p> <p>Once installed and configured, it starts on reboot.</p>

Set the environment variables for UNIX systems

To set environment variables for StorageTek silos for UNIX systems:

1. Create a Bourne shell script file `/nsr/nsrcc` 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.

CHAPTER 4

Backup to Disk and Cloud

This chapter covers these topics:

- ◆ Types of disk storage devices 162
- ◆ File type devices 166
- ◆ Advanced file type devices 167
- ◆ DD Boost devices 184
- ◆ Cloud devices 185

Types of disk storage devices

NetWorker software supports a variety of different backup to disk (B2D) methods. In all cases these methods use disk files that are created and managed as storage devices by the NetWorker software. These devices can reside on a computer's local disk or a network-attached disk.

The disk device types that NetWorker supports are FTD, AFTD, DD Boost, and cloud.

Disk-based devices that emulate other device types, such as virtual tape libraries (VTLs), are not covered in this chapter. [Chapter 5, "Backup to Tape and VTL,"](#) covers VTL devices.

The *EMC NetWorker Licensing Guide* provides information about NetWorker B2D and DD Boost licensing.

FTD

A file type device (FTD) is a very basic disk device type that has been available for many years. Its use and support is limited and is described in this chapter for legacy purposes only. ["File type devices" on page 166](#) provides details.

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 Microsoft Windows.

The Client Direct feature, also called direct file access (DFA), 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.

["Advanced file type devices" on page 167](#) provides AFTD configuration details.

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 the backup data is stored in a highly compressed and deduplicated format. DD Boost devices are accessed over a network by using the DD Boost API. DD Boost backups may be performed either through the NetWorker storage node workflow or the Client Direct file access (DFA) 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, bypassing the storage node and freeing up network bandwidth. The storage node manages the devices but does not handle the backup data workflow unless Client Direct workflow is not available.

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

DD Boost operations are not covered in this guide. The *EMC NetWorker Data Domain Deduplication Devices Integration Guide* provides details on DD Boost devices.

Cloud

Cloud devices are specific to cloud storage services, such as ATMOS. The cloud services are accessed through a private network. “[Cloud devices](#)” on page 185 provides details.

Example environment

Example 5 NetWorker backup to disk solutions

Figure 13 on page 163 shows various backup to disk options deployed in a mixed platform 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.
- ◆ Microsoft 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.

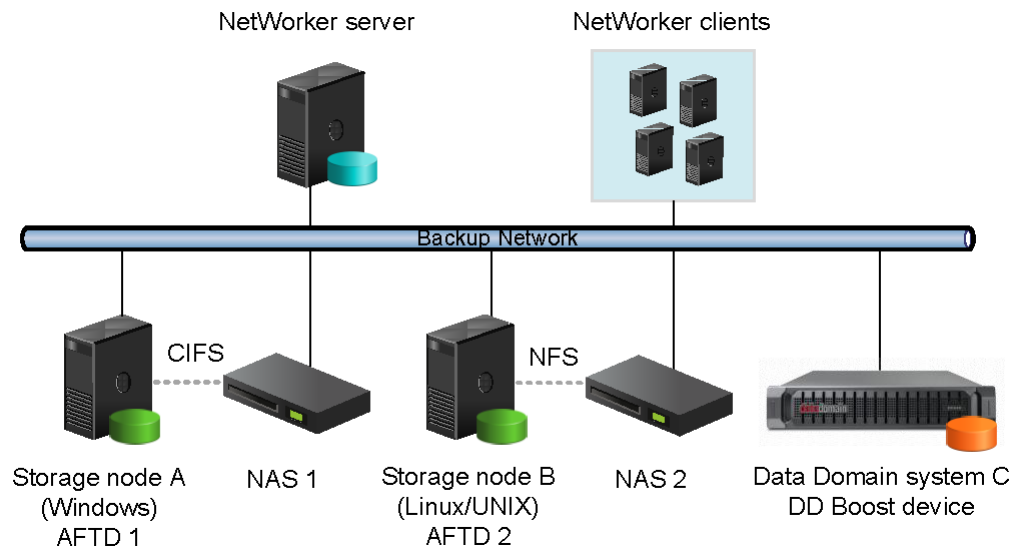


Figure 13 Example NetWorker DiskBackup configuration in a mixed backup environment.

Differences between FTDs, AFTDs and DD Boost devices

[Table 20 on page 164](#) lists the functional differences between traditional file type devices (FTDs), AFTDs, and DD Boost devices.

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

Table 20 Differences between disk devices (1 of 2)

Function or operation	File type device (FTD)	Advanced file type device (AFTD)	DD Boost device
Create a device	Device property window Select media type: file. UNIX/Linux storage node: local or NFS only. Windows storage node: local path only. CIFS is not supported for FTDs	<ul style="list-style-type: none"> 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.	<ul style="list-style-type: none"> 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 operations: “Recover savesets by using AFTD concurrent operations” on page 183 provides more information.	No	Yes	Yes
Reclaiming or recovering space	The nsrim program removes both aborted and expired save sets, once every 24 hours, after a savegroup is completed (if volume recycle is set to Auto).	<ul style="list-style-type: none"> Aborted save sets immediately removed. The nsrim program removes expired save sets, once every 24 hours, from the media database after a savegroup is completed (if volume recycle is set to Auto). Space on the AFTD is removed at the interval defined in the Reclaim Space Interval of the staging policy. 	<ul style="list-style-type: none"> Reclaims only data that is unique, not required by other existing backups. Aborted save sets are not immediately removed, but marked recyclable. This allows deduplication if the save set is restarted. The aborted save set is removed during the next Recover Space operation.
Volume default capacity for devices	If the file type device was used prior to 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.

Table 20 Differences between disk devices (2 of 2)

Function or operation	File type device (FTD)	Advanced file type device (AFTD)	DD Boost device
When file system or volume is full	<ul style="list-style-type: none"> • Waiting message 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. 	<ul style="list-style-type: none"> • Message displayed stating file system requires more space. • The nsrim program invoked to reclaim space for expired save set on AFTD. • Notification 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 completed on the same device.	No. Save sets that start on a DD Boost device must be completed on the same device.
Data format in device	EMC Open Tape Format (OTF).	Save stream (uasm) format (uses less space).	Deduplicated
Cloning operations	Save sets are cloned one at a time. Both automatic and manual cloning can begin only after all the save sets in a savegroup are backed up.	<p>Save sets are cloned one at a time. Automatic cloning begins after the save sets are backed up. Manual cloning of a save set can begin as soon as it has finished its backup.</p> <p>Two simultaneous clone operations can be run from the command prompt, provided no backup, recover, or stage operations run at the same time. The syntax is as follows:</p> <pre>nsrclone -S [ssid] cloneid1 nsrclone -S [ssid] cloneid2</pre>	The <i>EMC NetWorker Data Domain Deduplication Devices Integration Guide</i> provides details for advanced CCR cloning and replication with DD Boost devices.
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.	<p>Yes.</p> <p>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.</p> <hr/> <p>Note: NetWorker archive operations are not supported for Client Direct backup.</p> <hr/>	<p>Yes.</p> <p>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.</p> <hr/> <p>Note: NetWorker archive operations are not supported for Client Direct backup.</p> <hr/>

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.

Preventing full FTDs

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. [“Save set staging” on page 362](#) provides details.
- ◆ Review and, if required, modify the retention policy of the save sets.

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 *EMC Software Compatibility Guide* 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:

- ◆ 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.

Create and configure an AFTD

You can create an AFTD by using either the Device Wizard or the device properties window. Choose one of the following methods:

- ◆ [“Create an AFTD by using the Device Wizard” on page 168](#)
- ◆ [“Create an AFTD by using the Properties window \(Linux and UNIX\)” on page 170](#)
- ◆ [“Create an AFTD by using the Properties window \(Windows\)” on page 173](#)

Create an AFTD by using the Device Wizard

If you are creating an AFTD to use the client direct feature, see [“Considerations for Client Direct clients” on page 176](#) for information about specifying network path information when creating the AFTD.

To create one or more AFTDs by using the Device Wizard:

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.

NOTICE

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 your storage node is a Microsoft Windows system and you will 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
```

NOTICE

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
```

[“Configure multiple devices for a single volume” on page 178](#) 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 your 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, similar to the example in [Figure 14 on page 177](#) (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 176](#).

d. In **Target Sessions** specify the number of sessions that a nsrmmmd 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 nsrmmmd processes.

If another device is not available, then another nsrmmmd 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 attempted 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. In the **Review the Device Configuration** window, review the settings and click **Configure**.

14. In the **Check results** window, review whether the devices were successfully completed or any messages. Click **Finish**, or to go back, click **Back** or the appropriate wizard step.

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

To create an AFTD on a storage node running Linux and UNIX:

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 NetWorker disk file devices. The data could be overwritten.

2. In the server’s **Administration** interface, 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.

NOTICE

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 you created for the AFTD.

For example:

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 **Name** attribute must be specified with **rd=** in the following format, similar to the example in [Figure 14 on page 177](#) (for Microsoft Windows):

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.

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 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 root-only Client Direct Access

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, 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 176](#) provides additional details for Client Direct configurations.

[“Configure multiple devices for a single volume” on page 178](#) 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.

[“Configure multiple devices for a single volume” on page 178](#) 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 nsrmmmd count** limits the number of nsrmmmd processes that can run on the device. Use this setting to balance the nsrmmmd 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 nsrmmmd count.

At least one nsrmmmd process will always be 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 user name 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 user name, 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 attempt to make this change 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.

Create an AFTD by using the Properties window (Windows)

To configure an AFTD on a storage node running Microsoft Windows:

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 server's **Administration** interface, 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.

NOTICE

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, similar to the example in [Figure 14 on page 177](#):

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 directly access the device, bypassing the storage node.

[Figure 14 on page 177](#) shows this example.
 - For a CIFS-mounted AFTD, specify the complete paths of the directory created by using the Universal Naming Convention (UNC), for example:
 - \\CIFS_host\share-point-name\path
 - To enable UNIX/Linux clients to use this AFTD for Client Direct access, you must provide a NFS path in the host:/path format, for example:
 - NFS_host:/path

where:

- *NFS_host* is the hostname of the NFS file server
- *path* is the NFS-mountable path exported by the file server

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

[“Considerations for Client Direct clients” on page 176](#) provides additional details for Client Direct configurations.

[“Configure multiple devices for a single volume” on page 178](#) 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. 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.

[“Configure multiple devices for a single volume” on page 178](#) 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.

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

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

Device target and max sessions default values and ranges

[Table 21 on page 176](#) lists the default values and ranges for device target and max sessions in the NetWorker Administration interface.

Table 21 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 - 512
AFTD (including Data Domain CIFS/NFS)	4	32	1 - 10	1 - 512
Data Domain (DD Boost)	6	60	1 - 10	1 - 60
Cloud	1	512	Any	1 - 512
NDMP	4	512	1 - 32	1 - 512
FTD (traditional)	4	32	1 - 16	1 - 512
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.

Considerations for Client Direct clients

The Client Direct backup feature enables clients to back up directly to the storage devices, bypassing the storage node. The storage node manages the devices but does not handle the backup data. Device configuration for Client Direct clients depends on what type of storage device you will use and how it is connected to the storage nodes:

- ◆ Client Direct clients require 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 will generally yield better performance. UNIX clients must use a NFS path.
- ◆ If the storage device is to a NAS piece that is not directly attached to any storage node, the device access information (path) would be the same for all storage nodes and Client Direct clients.

- ◆ If the storage device is directly attached to a Windows storage node, then the storage node would use different device access information than the Client Direct clients. The device access information should specify multiple access paths to cover local and remote use cases. [Figure 14 on page 177](#) shows an example for a CIFS AFTD.

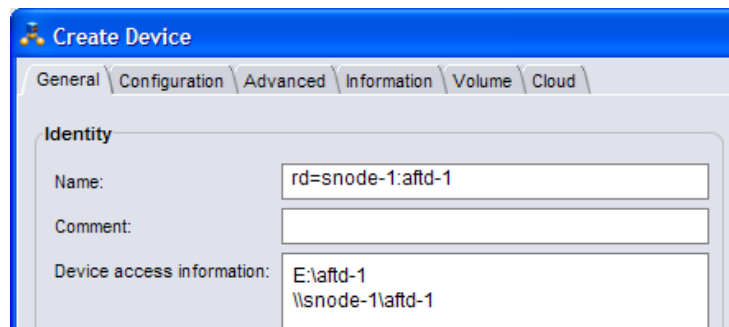


Figure 14 Example name and CIFS access information.

- ◆ As of NetWorker 8.1, non-root and cross-platform Client Direct backups to AFTDs are now supported for NetWorker clients on UNIX/Linux or Microsoft Windows. The AFTD can be managed by a Linux/UNIX or a Windows storage node. The AFTD can be local or mountable on the storage node. To configure such a device, you must:
 - Specify a NFS path in the AFTD's Device access information attribute. Specify the path using the *NFS_host:/path* format. Use this format whether or not 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.
 - You can also specify a CIFS path for Windows Client Direct backups. A CIFS path generally yields better backup performance than a NFS path for Windows Client Direct backups. 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
```
 - Specify the user name and password that is required to access the storage server on which the AFTD resides. Enter the user name and password in the Remote User and Password attributes in the NetWorker AFTD device resource.
 - The NFS server that provides AFTD storage must permit access using the NFSv3 protocol with AUTH_SYS (AUTH_UNIX) authentication.
 - The NFS server that provides AFTD storage must not restrict access to clients by using only privileged ports.

Note: When using a NetWorker version 8.1 server with an earlier version of a NetWorker storage node, it is possible to configure an AFTD using the new NFS path format (*NFS_host:/path*). However, this path will not work correctly because the older-version storage node does not have the required NFS support.

- ◆ Checkpoint restart does not support Client Direct backups 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, checkpoint restart points are not made less than 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.

Configure multiple devices for a single volume

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 167](#) provides more information.

Labeling and mounting an AFTD

To label and mount an AFTD:

1. Right-click the AFTD storage device and 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 dedicated to AFTD backup devices only.

NOTICE

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

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

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

Providing sufficient disk space for an AFTD

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 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. [“Viewing log files” on page 803](#) provides information about viewing log files.

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 invoke custom scripts and other programs to expand the capacity of existing AFTDs. [“Indexes” on page 585](#) and [“Configuring NetWorker SNMP notifications” on page 694](#) provides information on using notifications.

Create 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:

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

Create 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.

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

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. [Table 21 on page 176](#) provides default target and max session values and recommendations for other device types.

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

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 filesystem 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 filesystem.

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**.

Verifying AFTD operations

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.

Change 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.

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.

To set the maximum potential block size for an AFTD device:

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.

Recover savesets by using AFTD concurrent operations

When recovering from an AFTD, save sets are recovered concurrently. Multiple save sets can be simultaneously recovered to multiple clients. AFTD save sets can be cloned to two different volumes simultaneously. [“File type devices” on page 166](#) provides more information.

Limitations with concurrent AFTD recovery operations

AFTD concurrent recovery currently has these 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.
- ◆ NetWorker release 7.2x clients on Windows recover data from AFTD storage nodes sequentially.

Concurrent recoveries must be performed from the command line by using the **recover** command, either by using multiple **-S** options to identify multiple save sets, or executing multiple **recover** commands concurrently.

Deactivating and erasing an AFTD

To deactivate an AFTD device so it does not interfere with normal backup operations, use one of the following options:

- ◆ [“Convert a device to read-only” on page 183](#)
- ◆ [“Disable a device” on page 184](#)
- ◆ [“Delete a device” on page 184](#)

Convert 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.

To convert a device to read-only:

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 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**.

Disable 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.

To disable a device:

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.

Delete 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.

To delete an AFTD:

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**.

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.

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**.

DD Boost devices

DD Boost devices are covered separately in the *EMC NetWorker Data Domain Deduplication Devices Integration Guide*.

Cloud devices

This section describes how to configure the NetWorker Cloud Backup Option (NCBO) to perform backup, staging, cloning, and recovery operations to cloud configurations. Backups to cloud occur over a TCP/IP network and can be compressed and encrypted. NetWorker supports EMC Atmos-based cloud storage. More information on Atmos is available at www.EMC.com.

Cloud backup devices compared to other device types

NetWorker backup, staging, cloning, and recovery to cloud storage devices are similar to those operations that are performed with conventional devices. However, cloud devices also have unique features.

[Table 22 on page 185](#) lists the major similarities with other backup device types as well as the unique features of a cloud storage device.

Table 22 A comparison of cloud devices to other device types

Feature	Cloud backup device	AFTD device	Tape device
Same volume mounted simultaneously on multiple devices	Yes	No	No
Staging source	No	Yes	No
Staging destination	Yes	Yes	Yes
Cloning	Yes	Yes	Yes
Auto mount and unmount	Yes	No	Yes *
Data transformation engine (enables encryption and compression on storage node)	Yes	No **	No **
* When the tape is controlled by a tape library. ** Encryption and compression can still be enabled through NetWorker client side directives.			

Cloud backup requirements and considerations

The following conditions must be met before you can backup to the cloud:

- ◆ The NetWorker Cloud Backup Option must be licensed and enabled. The *EMC NetWorker Licensing Guide* provides information about licence enablers.
- ◆ The NetWorker Cloud Backup Option is supported on Windows and Linux storage nodes only.
- ◆ An Atmos cloud account is set up and you have a username and password to access the cloud account. The *EMC Atmos Installation Guide* and the *EMC Atmos System Management GUI Guide* provides information about setting up and managing an Atmos account.

- ◆ If the Atmos server and the NetWorker server are separated by a firewall, TCP ports 80 and 443 must be open to allow outgoing communication from the NetWorker server to the Atmos server. If a proxy server is configured in the environment, a firewall exception may also need to be created to ensure unrestricted access. If these ports are not open, device operations will fail with the following error:

```
Atmos label operation failed: Failed to write cloud label: Couldn't connect to server.
```

Be aware of the following considerations with respect to cloud support:

- ◆ NetWorker Avamar deduplication storage nodes do not support cloud backups.
- ◆ For NDMP, only a Data Server Agent (DSA) is supported for cloud backups.
- ◆ NetWorker versions prior to 7.6 do not support cloud backups.

Cloud best practices

Consider the topics and recommendations in this section before implementing cloud backups.

Directing NetWorker client backups to a cloud storage device

You direct client backups to a cloud storage device using media pools, in the same way you would direct any other client backup to a device or set of devices. However, be aware of the following recommendations:

- ◆ Set up a media pool for cloud storage devices and give the pool a unique label template.
- ◆ Do not mix cloud backup devices with other types of backup devices in a media pool. [“Sorting Backup Data” on page 303](#) provides information about setting up media pools.

Concurrent backup and recovery operations

A single cloud volume can be mounted on multiple cloud storage devices to support concurrent backup and recovery operations. For example, to optimize performance you could mount cloud volume A on three cloud storage devices: one for backup (device CL1), one for recovery (device CL2), and one for clone operations (device CL3). There is no limit to the number of cloud storage devices that can be mounted on a single cloud volume. Consider such an approach to optimize backup and recovery performance.

Network dependencies

Cloud backups are highly dependent on the network connection that is used to access the cloud service. Any disruption in connectivity or a slowdown in network access speed may adversely affect cloud backups or recoveries.

Creating and labeling a cloud storage device

This section contains the following topics:

- ◆ [“Create the cloud storage device” on page 187](#)
- ◆ [“Label and mount the cloud storage device” on page 189](#)

Create the cloud storage device

To create a cloud storage device:

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. In the **Name** field, replace the default name with a name that uniquely identifies the cloud storage device. If the device is configured on a remote storage node, indicate that the storage node is remote, by including *rd=hostname:* in the name. For example, if the remote storage node is neptune, then the device path might be rd=neptune:cloud1.

Note: A cloud storage device name does not specify a path to the device. You can use any combination of alphanumeric characters for the device name.

4. In the **Comment** field and the **Description** field, add an optional comment and description, respectively.
5. In the **Media Type** field, select Atmos COS as the device type if you are using Atmos as the cloud server.
6. In the **Remote User** field, type the username that is used to access the cloud server. For an Atmos COS device, this is the *token-id*.
7. In the **Password** field, type the password that is used to access the cloud server. For an Atmos device, this is the *shared secret*.
8. Select the **Cloud** tab to specify additional information specific to the cloud backup device.
9. In the **Server** field, type the IP address or fully qualified domain name of the cloud server.
10. Use the Parameter options to adjust network communication attributes:
 - a. In the **Network Write Size** field, specify the amount of backup data, in kilobytes, to cache in memory before sending to the cloud. Larger write sizes typically result in better performance but results vary depending on the underlying network characteristics. Also be aware that larger sizes consume more memory on the storage node for the duration of the backup or recover operation.
 - b. In the **Number of Retries** field, specify the number of times that NetWorker will attempt to send backup or receive recover data in the event of a network failure.
 - c. In the **Send/Receive Timeout** field, specify the number of seconds that NetWorker will wait for confirmation that network send and receive transmissions to the cloud server have occurred successfully. If the timeout period expires, the data transmission is considered to have failed.
 - Set the value of this field in conjunction with the Network Write Size field. Larger Network Write Size values require larger Send/Receive Timeout values to avoid failures. Optimal values for the Send/Receive Timeout field vary depending on the network speed and bandwidth.

- The save group’s Inactivity Timeout value can potentially interact with the Send/Receive Timeout value in unintended ways. To avoid this possibility, ensure that the save group’s Inactivity Timeout value (default is 30 minutes) is greater than the Send/Receive Timeout value (default is 30 seconds).
 - d. In the **Network Failure Retry Interval** field, specify the number of minutes that a backup or recover session must wait before a failed network connection results in an aborted backup or recover session.
11. In the **Compression** field, select a compression level for data that is sent to the cloud. Faster compression speeds result in less data compression but also require less CPU resources. The fastest compression speed, Compression Speed Fast, performs the least amount of data compression and is selected by default.

To choose an optimal compression value, balance the potentially longer backup window of using a slower compression speed against the potential efficiency and cost savings of sending less backup data to the cloud.

NOTICE

If the NetWorker Cloud Back Option determines that backup data cannot be compressed effectively, compression may not occur regardless of the setting in this field.

12. In the **Encryption** field, specify whether to enable or disable encryption of data sent to the cloud. Encryption is standard NetWorker AES 256 bit encryption and is selected by default. If desired a NetWorker datazone pass phrase can be defined that would be used to recover encrypted data.


If this option is selected, encryption will occur regardless of any client-side encryption directives. For more information about encryption including how to specify a new datazone pass phrase, refer to [“Encrypting backup data” on page 108](#).

NOTICE

If encryption is already enabled for the NetWorker client and encryption is enabled in this field, backups will be slower because encryption functions will occur twice.

13. Use the **Cloud network interface** field if the Storage node has multiple network interfaces. If it does, specify the IP address of the network interface that will send backup data to the cloud.

To display the **Cloud network interface** field, select **View>Diagnostic Mode** from the menu bar.

14. Select **Throttling** and then click the Bandwidth icon  to display a dialog box where you can adjust the maximum internet bandwidth that a cloud backup or recovery operation can consume at any given time of the day or week. This option enables you to prevent network congestion by limiting cloud backup and recovery activity during peak internet usage.
- a. Select **New** to add a bandwidth throttling policy.
 - b. From the **Day** field, select the day of week to which the policy applies.

- c. Click the up and down arrows to select a time of the day to which the policy starts and ends. Alternatively, type the times directly into the **Start time** and **End time** fields.
- d. Click the up and down arrows to select the maximum possible network bandwidth, in megabits per second, that a backup or recovery operation can consume when the policy is in effect. Alternatively, type the values directly in the fields.

You can create as many policies per day as required. You can also modify or delete existing throttling policies as necessary.

15. Click **OK** when the configuration is complete.

Label and mount the cloud storage device

To label and mount a cloud storage device:

1. Select the cloud storage device, right-click and select **Label**. The **Create new cloud volume** dialog box appears.
2. In the Pools field, select the media pool to be used for cloud storage devices.

Note: It is recommended that the media pool you select be used for cloud backup devices only.

A label for the cloud storage device is generated and displayed in the **Volume Label** field. The label name is based on the label template that was specified for the cloud media pool.

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

If there are multiple cloud volumes, you will be able to select the volume to associate with the cloud storage device.

Gathering report information on cloud backup

Use cloud backup information to monitor backup costs and help optimize your cloud backups. Cloud backup information can be obtained from the following sources:

- ◆ [“Cloud backup and recover reports” on page 448](#)
- ◆ **mminfo** command
Use the **mminfo -avot** command to get information on how much data is consumed in a cloud backup.

Staging with a cloud storage device

Staging with a cloud storage device works the same way as staging to a tape device. You cannot however, use a cloud storage device as the source for a staging operation.

[Chapter 13, “Staging Backups”](#) provides more information.

Cloning to a cloud storage device

Cloning with a cloud storage device works the same way as cloning with any other advanced file type device. [Chapter 12, “Cloning”](#) provides more information.

CHAPTER 5

Backup to Tape and VTL

This chapter covers these topics:

◆ Overview of tape device storage	192
◆ Stand-alone devices	192
◆ SCSI data block size issues between UNIX and Windows	194
◆ Device parameter settings	195
◆ Common device interface	199
◆ Device ordering	199
◆ Dynamic drive sharing	203
◆ Improvements to deduplication rates for Data Domain VTL multiplexed backups ...	210
◆ Nonrewinding tape device usage (UNIX/Linux only)	212
◆ Support for LTO-4 hardware-based encryption	213
◆ Recycling compared to adding more volumes	213
◆ Display device operations messages	213
◆ Service mode	214

Overview of tape device storage

This chapter contains information on the creation, configuration, and management of tape devices. These 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.

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 132](#) 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

To configure a new stand-alone tape drive, automatically, by using Scan for Devices:

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

To configure (add) a new stand-alone device on a storage node:

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 Microsoft Windows would have a format similar to \\.\Tape0.
 - b. If the device is configured on a remote storage node, however, 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 166](#) provides instructions and restrictions on backing up to a file type device.

4. In the **Identity** area, configure:
 - 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.

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.

Determine the allowable block size

To determine the allowable block size, check the **Properties** window of a mounted volume while in **Diagnostic Mode**:

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**.

Set the block size for a device type

To set the block size for an entire device type:

- ◆ 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.

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 these sections:

- ◆ [“Device settings in the NetWorker Administration interface” on page 195](#)
- ◆ [“Device settings as environment variables” on page 196](#)

Device settings in the NetWorker Administration interface

To locate and change the device parameters in the Administration interface:

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.

- Select the **Advanced** tab. In the **Device Configuration** area, the device settings are the first fields shown. [Table 23 on page 196](#) lists the fields and their corresponding environment variables:

Table 23 Device settings and environment variables

Device setting	Corresponding environment variable
Device Block Size	NSR_DEV_BLOCK_SIZE_ <i>MEDIA_TYPE</i>
Device File Size	NSR_DEV_TAPE_FILE_SIZE_ <i>MEDIA_TYPE</i>
Device Load Time	NSR_DEV_LOAD_TIME_ <i>MEDIA_TYPE</i>
Device Eject Time	None
Device Poll Interval	NSR_DEV_LOAD_POLL_INTERVAL_ <i>MEDIA_TYPE</i>
Device Min Load Tries	NSR_DEV_LOAD_TRY_LIMIT_ <i>MEDIA_TYPE</i>
Device Default Capacity	NSR_DEV_DEFAULT_CAPACITY_ <i>MEDIA_TYPE</i>
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.

Device settings as environment variables

The device-related environment variables are:

- ◆ `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.

Setting environment variables for the NetWorker software differs on Windows and UNIX operating systems:

Windows:

Environment variables on Microsoft Windows are set using the Control Panel System applet on the NetWorker server:

1. Navigate to **Control Panel -> System and Security -> System -> Advanced System Settings**.
2. In the **General** tab click **Environment Variables...**
3. Click the **New** button.
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.

UNIX:

NetWorker 8.0 introduces support for a new NetWorker environment variable file. This file, **/nsr/nsrrc**, will be sourced prior to starting the NetWorker processes.

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.

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 194](#) 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 **nsrmmmd** 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 **nsrmmmd** 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 **nsrmmmd** 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 **nsrmmmd** will attempt to open a drive. The **nsrmmmd** 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
```

Common device interface

The common device interface (CDI) allows the NetWorker server to send commands to tape devices. The CDI feature is not supported within an NDMP environment.

CDI support can be set in the NetWorker Administration interface:

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 the CDI feature can be disabled through selecting the Not Used option, it can be time-consuming to disable a large number of devices.

In this situation, access the `/nsr/debug` directory and create a file named `cdidisable`. Then restart the NetWorker server. This file does not need any content, it just needs to exist. This disables the use of CDI for that server and all storage nodes controlled by that server.

Note: Use of CDI does not change what is written to tape. A tape written with CDI enabled can be read with CDI disabled. Conversely, a tape 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 EMC Customer Support representative. If tape or SCSI issues occur while the CDI feature is enabled, go to the EMC online support.

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 480](#). CDI must be enabled for this functionality. [“Common device interface” on page 199](#) provides more information about enabling CDI.

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).

Once persistently named device files are created and 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.

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 **sjsn** commands, then verify that the device configuration within your NetWorker configuration conforms to this.

To detect device ordering issues:

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 **sjsn** command to determine the current order of the devices:

```
sjsn 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 `sjsn` 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.

Correcting drive ordering changes

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

To correct drive ordering changes by using the NetWorker Console:

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 155](#) 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

To correct drive ordering changes by using the `jbedit` command:

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.

“Using the `jbedit` command to configure a library” on page 147, or the UNIX man page for `jbedit` or the *NetWorker Command Reference Guide* provides more information on 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:

```
"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. To resolve the problem and clear the error message:

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.

The Alert will be removed and the event dismissed.

Reordering tape drive numbers (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.

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 do the following:

- ◆ 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.

Two terms 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.

NOTICE

DDS is currently supported only 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 — A single tape drive can be shared among several storage nodes. In fact, since NetWorker software uses the same open tape format for UNIX, Windows, NetWare and Linux, the same tape can be shared between different platforms (assuming that respective save sets belong to the same pool).
- ◆ Reduces LAN traffic — Clients can be configured as SAN storage nodes that can send save sets over the SAN to shared drives.
- ◆ Provides fault tolerance — Within a SAN environment, hardware can be configured to eliminate a single point of failure.
- ◆ Provides configuration over a greater distance — Allows configuration of a system over a greater distance than with SCSI connections.

DDS configuration overview

Figure 15 on page 204 illustrates the DDS process and potential configurations for sharing drives. This basic configuration consists of a server, two storage nodes, and a library with two tape drives.

In this figure:

- ◆ Storage nodes sn_1 and sn_2 are attached to the library.
- ◆ Each node, on its own, has access to drive_1 and drive_2.
- ◆ With DDS enabled, both nodes have access to both drives and can recognize when a shared drive is in use.
- ◆ Under such a configuration, two DDS licenses are required, one for each drive.




Figure 15 Dynamic Drive Sharing

NOTICE

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 and/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.

[“SCSI data block size issues between UNIX and Windows” on page 194](#) provides information about how to set block sizes for individual drives or tapes on different platforms.

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.

Preventing unintended access to DDS devices

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. [“Common device interface” on page 199](#) provides more information regarding CDI.

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

Note these restrictions for this feature:

- ◆ 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.

Enabling DDS with NDMP

These sections explain the requirements for successfully enabling DDS with NDMP.

DDS on NDMP nodes in a SAN environment

Drives can be shared between NDMP nodes in a SAN environment.

Consider the following:

- ◆ All the components of a SAN configuration must be compatible when DDS is enabled with the NetWorker NDMP feature.
- ◆ The Fibre Channel switches must be compatible with any NDMP hosts within a SAN.
- ◆ NDMP hosts and libraries in the SAN must also be compatible with each other.
- ◆ The NDMP nodes that will share the drives are homogeneous. For example, DDS can be enabled in these configurations:
 - EMC Celerra® to EMC Celerra
 - NetApp to NetApp (any NetApp nodes that Network Appliance supports within a zone)

The current NDMP implementation does not allow the sharing of drives between non-homogeneous NDMP nodes. There is, however, no inherent limitation within DDS that would prevent this.

NetApp zoning requirements for DDS in a SAN environment

In order to configure DDS with NetApp filers, a zoned SAN configuration is required. Zoning is a feature of the Fibre Channel switch.

Consider the following when configuring DDS with NetApp filers:

- ◆ The NetApp zone, which contains only the NetApp filers and tape devices, must be configured on the Fibre Channel switch. This NetApp zone may also include the robotic arm and must also be configured in an arbitrated loop.
- ◆ All non-NetApp servers that are attached to the same Fibre Channel switch must be excluded from the NetApp zone. A separate zone must be configured for the non-NetApp servers, in which an arbitrated loop may or may not be a requirement.
- ◆ The NetApp zone and all other zones can overlap on the tape devices within the SAN, so that the tape devices are visible to both zones.

Figure 16 on page 208 illustrates a basic DDS configuration with NDMP.




Figure 16 DDS with NDMP

DDS attributes in the device properties

Attributes used in DDS are found in the Properties window for a device:

- ◆ Hardware ID
- ◆ Shared Devices

Hardware ID attribute

The Hardware ID attribute tracks the drives that are being shared by multiple hosts. Device instances sharing the same physical drive across multiple hosts have the same hardware ID. The hardware ID is automatically assigned during the device autoconfiguration process, or it can be added when manually configuring a device. It is not editable by users.

The hardware ID can be viewed in the Properties window for a device, on the General tab, in the Device Sharing area.

The hardware ID is generated when a device is scanned or configured. The hardware ID consists of the following:

- ◆ The hardware serial number
- ◆ The device type
- ◆ The worldwide part number (WWPN)
- ◆ The worldwide name (WWN)

Do not try to change a hardware ID once it has been generated. It is read-only.

Shared Devices attribute

The Shared Devices attribute appears on the Operations tab of a device's Properties when in diagnostic mode. It features values that can be used to manipulate all shared instances of a drive at the same time. This attribute enables or disables all devices that share the same hardware ID with a single action. [Table 24 on page 209](#) lists this attribute's allowed values and their descriptions.

Table 24 Shared Devices attributes

Value	Description
Enable All	When selected, enables all devices by using the same hardware ID.
Disable All	When selected, disables all the devices by using the same hardware ID.
Done	This is the default setting. After the server has enabled or disabled all devices with the same hardware ID, the attribute is reset to Done.

The Shared Devices attribute is not reflected in the **jbconfig** program.

Idle Device Timeout attribute and DDS

A tape might remain mounted in its drive after a backup has completed. Other requests for the drive from another device path must wait during this timeout period. The timeout value can be adjusted by changing the Idle Device Timeout attribute.

The Idle Device Timeout is not specifically a DDS attribute, but it can be useful in configuring shared drives. This attribute appears on the device Properties Advanced tab when displayed in diagnostic mode. The default value is 0 (zero) minutes, which means that the device never times out and the tape must be ejected manually.

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's value will take effect only on those devices whose Idle Device Timeout value is 0. The library's Idle Device Timeout value is located on the Timer tab of the library Properties window. [“Automatic unmounting of volumes \(idle device timeout\)” on page 231](#) provides more information.

High availability and DDS

The NetWorker software relocates and restarts operations that were in progress when a failure occurs on a cluster node. Currently, savegroups are the only highly available operations.

The nsrjb program high availability limitations

If the NetWorker server fails over from one node to a new target node, standard library operations (such as performing an inventory, labeling, mounting, or unmounting a volume) do not automatically restart on the new target node.

Example 6 Host crash requires user intervention

This example scenario includes: two physical hosts, A and B, with DDS enabled, sharing the drives on a library.

Physical host A mounts a tape in a shared drive on the library. If physical host A subsequently crashes, the volume is held in that shared drive until the reset command **nsrjb -H** is issued (from host B, in this example) along with a reset from the Library Operations Windows in the NetWorker Console.

This command unloads the drive and makes it available for future backups. The reset command clears the drive by accessing the device through another shared path. In this example, the other shared path would be on host B.

Successfully unloading a volume requires that the NetWorker software be able to access the same path through which the volume was initially loaded.

Improvements to deduplication rates for Data Domain VTL multiplexed backups

New functionality has been introduced to improve tape deduplication ratios of multiplexed backups for remote client save streams. This pertains to non-dedicated storage nodes that use DD5.x VTL devices. With this functionality, significantly higher NetWorker de-duplication ratios are obtained without noticeable impact to the device MB/s throughput.

No special configuration is required when using a NetWorker dedicated storage node (DSN) or NetWorker saves directed to a local device.

To obtain the best deduplication rates, you would generally disable multiplexed deduplication by setting each VTL device's properties **target sessions** and **max sessions** values to **1**. However, for environments with high overall concurrency requirements (for example, a large number of parallel backups required due to a limited backup window duration business requirement), an excessive number of virtual tape drives might not be possible, or might be in use but causing backup runtime stability problems, such as intermittent volume mount operation delays or timeouts. In this case, consider applying the new functionality, with slight increases to the max sessions, target sessions and device block size values and using a fraction of the previous number of VTL devices.

Note: When the new functionality is first enabled with the recommended increases in device block size and max sessions, backups will initially have lower deduplication ratios prior to achieving the higher target deduplication rate. This also occurs without the new functionality enabled, such as when switching from multiplexed to non-multiplexed backups, and reflects initial DDR re-priming or re-analysis overhead. Therefore, extra caution should be observed on DDR systems which are already heavily utilized (for example, 75% or more DDR disk space already used).

To enable the new functionality for network saves globally:

1. Shut down the relevant NetWorker storage node's service (or the entire NetWorker server if possible).
2. After verifying there is no backup activity on the NetWorker DD VTL storage node, change the **max sessions** value for each relevant VTL device to 4 (recommended) up to 32 (maximum) using NMC or nsradmin. **Target sessions** and **device block size** should also be set accordingly.

Recommended values are **max sessions = 4**, **target sessions = 4**, and **device block size = 512KB**. However, the optimal values will depend on your environment (for example, **max sessions = 2** may provide better deduplication while still meeting your backup window without stability problems).

The following example provides a way to set the recommended values using the UNIX/Linux or Windows command line.

Run the following:

```
> nsradmin -i input_file.txt
```

where **input_file.txt** contains the following example lines to be customized for your 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
```

Note: If the NetWorker server was shut down in step 1, run this **nsradmin** command with the **-d resdir** option, which uses the NetWorker resource database resdir instead of opening a network connection.

3. Create a no-intra-block-multi-plexing (**nibmp**) tag file in the NetWorker debug folder on the NetWorker storage node. For example, using the standard NetWorker installation paths:

On Unix/Linux:

```
#touch /nsr/debug/nibmp
```

On Windows:

```
> echo > "⟨NetWorker_install_path⟩\nsr\debug\nibmp"
```

You can limit the patch to specific pools by adding **⟨Pool name⟩** to the end of the tag file. For example, if the pool name was **My Pool**:

On Unix/Linux:

```
#touch "/nsr/debug/nibmp_My Pool"
```

On Windows:

```
> echo > "⟨NetWorker_install_path⟩\nsr\debug\nibmp_My Pool"
```

The pool name can include spaces. Ensure there are quotes around the specified pathname.

4. Restart the NetWorker services to enable the functionality.

Note that deduplication efficiency is still relative to the number of multiplexed save sets, and the ratio drops with higher concurrency. From measured tests, expected efficiency dropped by 4% to 8% for each additional parallel save stream.

For example, using a sufficiently large device block size and 4 parallel streams (where device property **max sessions** is set to **4**), expected de-duplication ratios are 12-24% below the ideal (a non-multiplexed backup where max sessions is set to 1).

Additionally, the following best practices are recommended:

- ◆ Install DD OS version 5.0.2 or later if currently using DD OS 5.0.x. This version of DD OS addresses a large number of DD VTL specific issues, and has higher deduplication ratios in general than DD OS 4.9 versions for DD VTL use cases.

If using DD OS 4.x, install DD OS version 4.9.3.1.

- ◆ The recommended block size is 512 KB. The standard NetWorker default block size (handler default) depends on the device type and varies between 32KB and 128KB. From measured tests, increased block size had an additional positive impact on deduplication efficiency of 15% to 25%.
- ◆ Adjust the number of VTL devices to a lower value to correspond to the new max sessions value. If considering the recommended values in step 2, the number of VTL devices can be inversely reduced: increase each device max sessions from 1 to 4, then decrease the number of required VTL devices to $\frac{1}{4}$ of the current number.
- ◆ Change the device parameters for **max sessions** and **device block size** until a combination is obtained that results in higher deduplication ratios, a smaller backup window and increased system stability. Recommended values are provided in step 2.
- ◆ As a general best practice when performing deduplication, ensure that no client/server side encryption or compression occurs prior to reaching the Data Domain device.

More information is provided in the “NetWorker Improved Deduplication with Multiplexing to Data Domain VTLs” technical note available on the EMC Online Support website at [https://support.emc.com/docu37676_NetWorker-7.6.3-\(and-Later\)-Improved-Deduplication-with-Multiplexing-to-Data-Domain-VTLs.pdf](https://support.emc.com/docu37676_NetWorker-7.6.3-(and-Later)-Improved-Deduplication-with-Multiplexing-to-Data-Domain-VTLs.pdf)

Nonrewinding tape device usage (UNIX/Linux only)

Tape drives used as storage devices must be accessed by nonrewinding 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 nonrewinding, 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 nonrewinding behavior on Solaris, HP-UX, Linux, and HP-Tru64.

On AIX, the number following the decimal selects the BSD and nonrewinding behavior and must be either 1 or 5 for NetWorker software (for example `/dev/rmt2.1`)

Note: Never change a device pathname from nonrewinding (`/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.

Support for LTO-4 hardware-based encryption

The use of LTO-4 hardware-based encryption is supported by NetWorker 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; however, 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. Neither the encryption nor the key management process is managed by the NetWorker application. This includes the ability to turn encryption on or off within NetWorker, and the management of encryption keys.

Recycling compared to adding more volumes

The NetWorker server saves files on volumes marked appen (appendable). If the volumes are marked full, they cannot receive backups.

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 a volume’s mode” on page 222](#) 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.

Display device operations messages

To display device operations messages:

1. In the NetWorker **Administration** window, click **Monitoring**.
2. Select the **Operations** tab.
3. Right-click an operation, and select **Show Details**.
The details window for the selected device appears.
4. Click **Close** to exit the window, or **Save** to save the message.

Service mode

Use the service mode setting to take a device offline temporarily. Service mode differs from the disabled state in that the **nsrmm** 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 **nsrjb** or **nsrmm** command with the **-f** option. For more information, refer to the *EMC 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 236](#) 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.

To put a device in service mode:

1. Open the device’s **Properties** window
2. On the **General** tab, set **Status Enabled** to **Service**.

CHAPTER 6

Media Management

This chapter covers these topics:

- ◆ Storage management operations 216
- ◆ Auto Media Management 218
- ◆ Volume operations 220
- ◆ Media handling errors 236
- ◆ Media management in a silo 237
- ◆ Volume save sets 242

Storage management operations

This section describes the components involved in the operation of storage volumes done through the NetWorker server. The details that describe a particular volume can be viewed, and often changed, by right-clicking the volume and making a selection from the menu. As with other Console functions, users can view and work with only those NetWorker servers for which they have access permission.

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 304](#) 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 it labels a volume, the NetWorker server:

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.

NOTICE

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.

How the NetWorker server selects a volume

When a backup takes place, the NetWorker server searches for a volume from the appropriate pool to accept the data for backup. 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:

1. Mounted volumes from the appropriate pool with the mode appendable are selected. This includes newly labeled volumes. If more than one mounted volume is appendable, the server uses this hierarchy:
 - a. Device availability. The server writes to the volume from the appropriate pool that is mounted on the device with the fewest current sessions.
 - b. Volume label time. The server writes to the volume with the oldest label time if the mounted volumes are appendable and session availability is not an issue.
2. If a library is in use and there is no mounted, appendable volume in the library, the server determines whether there is an unmounted, appendable volume available. This includes newly labeled volumes.
3. If multiple unmounted, appendable volumes are available, the volume with the oldest label time is selected.
4. If no mounted volumes are appendable and Auto Media Management is enabled, a mounted volume with the mode recyclable is selected. The server relabels and mounts the volume.
5. If a stand-alone device is being used and Auto Media Management is not enabled, the server sends a mount request notification.
6. If a library is in use and no unmounted, appendable volumes exist, the server determines whether there is an unmounted, recyclable volume.
7. If Auto Media Management is not enabled, or if there are no appendable or recyclable volumes, the server sends a mount request notification.

[“Volume operations” on page 220](#) provides information about appendable and recyclable volumes.

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 relabel

The NetWorker 8.0 and later software releases provide the ability to automatically relabel recyclable volumes when needed or when scheduled.

Consider the following:

- ◆ If Auto Media Management is enabled and a volume has the mode recyclable, the server automatically relabels the volume. 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 218](#) provides more information on Auto Media Management.
- ◆ A media pool can be configured to automatically relabel recyclable volume at a user defined time and interval. [“Managing volumes in a media pool” on page 313](#) provides more information about configuring the automatic relabel process for recyclable volumes in a media pool.
- ◆ The mode of a volume can also be manually changed to recyclable. [“Changing a volume’s mode” on page 222](#) provides information about changing the mode of a volume.

Auto Media Management

This section describes how NetWorker works with Auto Media Management.

Auto Media Management gives the NetWorker server automatic control over media loaded in the storage device. When Auto Media Management is enabled during device configuration, the NetWorker server automatically:

- ◆ Labels the volume (recognizes EDM labels and does not overwrite them).
- ◆ Mounts the volume.
- ◆ Overwrites volumes it considers 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 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.
- ◆ Because the Auto Media Management feature can relabel 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.
- ◆ When the Auto Media Management feature is not enabled, the NetWorker server ignores unlabeled volumes and does not consider them for backup.

NOTICE

The NetWorker server considers volumes that were labeled by a different application to be valid relabel candidates if Auto Media Management is enabled. Once the NetWorker server relabels the volume, the previously stored data is lost.

Using Auto Media Management

This section describes how to use Auto Media Management.

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.

Enabling 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:
 - a. If yes, the write operation continues.
 - b. 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.

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 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.

To enable auto media management:

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**.

Volume operations

The volume operations sections describe the tasks involved in the operation of storage volumes done through the NetWorker server. Information about storage volumes is available for each device on a NetWorker server. All volume operations are performed in the Media task in the **Administration** window.

If a volume is not mounted when a backup is initiated, 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 more than one volume is needed to recover the files, the server lists all the volumes in the order in 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 stored in the library.

The NetWorker server reports on the status of volumes using values such as:

- ◆ Volume name
- ◆ Written
- ◆ %Used
- ◆ Location
- ◆ Mode

Performing volume operations requires that the user have the correct permissions to use the NetWorker server and its storage nodes.

Viewing volume status information

To find information about a volume and its status:

1. In the server's **Administration** window, click **Media**.
Media-related topics appear in the navigation tree.
2. Select **Volumes**. The Volumes detail table appears. [Table 25 on page 221](#) lists the volume-related categories displayed in the Volumes detail table.

Table 25 Volumes detail

Category	Description
Volume Name	Within the Administration interface, the volume name is the same as the name that appears on the volume label. At the end of the name, these designations might appear: <ol style="list-style-type: none"> 1. (A) indicates an archive volume. 2. (R) indicates a volume that is considered read-only. 3. (W) indicates that the volume is a write once, read many (WORM) device.
Barcode	Barcode label, if one exists.
Used	Indicates the amount of space currently in use on the volume (shown in KB, MB, GB, as appropriate). When Used is equal to full, there is no more space on the volume and the end-of-tape marker has been reached or an error with the volume has occurred.
% Used	An estimate of the percentage used, based on the total capacity of the volume, and on the specified value of the Media Type of the Device resource. When %Used is equal to 100%, it means that the value is equal to, or exceeds, the estimate for this volume. When the word "Full" appears in the % Used column, it is not based on an estimate of the volume's capacity. "Full" literally means that the volume is full. This attribute applies only to tape volumes. File type and advanced file type devices always display 0% Used.
Mode	Choices are appendable, read-only, and recyclable. Table 26 on page 221 lists the NetWorker volume modes and their definitions. " Changing a volume's mode " on page 222 provides information on changing volume modes.
Expiration	Changing the expiration date is only possible from the command prompt. Use the nsrmm command to do this.
Location	Refers to an administrator-defined description of a physical location of the volume within the user's environment, for example, 2nd shelf, Cabinet 2, Room 42.
Pool	Name of the pool to which the volume belongs.

Table 26 Volume modes (1 of 2)

Mode value	Meaning	Description
appen	appendable	This volume contains empty space. Data that meets the acceptance criteria for the pool to which this volume belongs can be appended.
man	manual recycle	This volume is exempt from automatic recycling. The mode can be changed only manually. The manual recycle mode and the option to change it are available from the Volumes menu. The default recycle mode is Auto.

Table 26 Volume modes (2 of 2)

Mode value	Meaning	Description
(R)	read-only	The save sets on this volume are considered read-only. The mode can be changed only manually. “Using the read-only mode” on page 222 provides more information.
recyc	recyclable	The save sets on this volume have exceeded their retention policies.
full	full	The volume is full. There is no more space for data in the volume, and the save sets have not yet passed the time periods specified by the retention policies. This mode can be set only from the command-prompt. Use the nsrjb command with the -o option for libraries, and the nsrmm command with the -o option for stand-alone drives. Refer to the respective UNIX man pages of those commands (or to the <i>EMC NetWorker Command Reference Guide</i>) for more information.

Using the read-only mode

When the mode of a volume is read-only, no new data can be written to the volume. A read-only volume is *not* a write-protected volume. The save sets on the volume are still subject to their browse and retention policies, and the volume is not protected from being overwritten. When all the save sets on the volume have changed their status to recyclable, the mode of the volume changes to recyclable, and the volume becomes eligible for automatic recycling.

[“Changing a volume’s mode” on page 222](#) provides information on changing the volume mode.

Changing a volume’s mode

To change the mode of a volume:

1. In the server’s NetWorker **Administration** window, click **Media**.
2. In the navigation tree, select **Volumes**. The **Volumes detail** table appears, showing all of the server’s volumes.
3. Right-click a volume in the **Volumes detail** table, and select **Change Mode**. The **Change Mode** window appears.
4. Select a mode and click **OK**.

Recycling volumes

A volume's retention policy can be overridden by changing it to manual recycle. One reason to change to manual recycle is when save sets must be kept on a volume longer than its retention policy specifies. A volume marked for manual recycle can be changed back to recycle automatically, so that the volume once again uses its original retention policy.

Change a volume’s recycle policy

To change a volume’s recycle policy:

1. In the server’s **Administration** window, click **Media**.
2. In the navigation tree, select **Volumes**. The **Volumes detail** table appears.

3. Right-click a volume in the **Volumes** detail table, and select **Recycle**. The **Recycle** window appears. It names the selected volume.
4. Select the recycle policy: **Auto** (default) or **Manual**.
5. Click **OK**.

NOTICE

A volume that has been set to manual recycle retains that setting, even after relabeling. It must be explicitly reset to use auto recycle.

Labeling volumes

The NetWorker software labels each storage volume with a unique internal label that corresponds to a pool. During backup and other operations, this label identifies the pool to which a volume belongs. NetWorker software applies a label template to create a unique internal label for each volume. The NetWorker server uses label templates and pool configuration settings to sort, store, and track data on media volumes.

NOTICE

Label templates are created in the Media task, but they are applied to volume labels in the Devices task. Data that exists on a tape is effectively gone after the tape has been relabeled.

Labeling a volume does the following:

- ◆ Writes a label on the volume.
- ◆ Adds the volume label to the media database.
- ◆ Prepares the tape to have data written to it.

During data recovery, the server asks for a specific volume that contains the required data, identifying the required volume by the name with which it was labeled. [Chapter 10, “Sorting Backup Data”](#) provides information about label templates and pools.

Label templates

Several preconfigured label templates are supplied with the NetWorker software. These preconfigured label templates cannot be deleted. [“Naming label templates” on page 321](#) provides more information about label templates and preconfigured label template.

Label or relabel library volumes

Labeling volumes in a library is time-consuming, so consider labeling volumes before it is time to back up or recover files. Library volumes are labeled in the Devices task.

To label a library volume:

1. In the **Administration** window, click **Devices**.
2. Open the **Libraries** folder. The **Libraries** detail table appears.
3. In the navigation pane, right-click the appropriate library and select **Label**. The selected library’s details appear, including divided tables for devices and slots. The **Label Library Media** window also appears.

4. The **Default** pool name appears in the **Target Media Pool** field. To select a different pool, click the field's down arrow for a list of other pool choices. The pool determines which label template is used in labeling the volume.
5. If the volume should not be recycled automatically, click **Allow Manual Recycle**. If **Allow Manual Recycle** is enabled when the volume is labeled, the volume is not automatically marked as recyclable when all of its save sets have expired. Only an administrator can mark the volume recyclable.

NOTICE

A volume that has been set to manual recycle retains that setting, even after relabeling. A Manual Recycle policy cannot be changed back to Auto Recycle simply by unselecting the Manual Recycle checkbox. The volume must be explicitly reset to use auto recycle. [“Recycling volumes” on page 222](#) provides more information.

6. To be prompted before the existing label is overwritten, select **Prompt to overwrite label**.
7. Click **OK**. The **Library Operation** window appears, stating that the library operation has started.
8. Select **Monitoring**, then the **Operations** tab, to track the status of the label operation.
9. If **Prompt to overwrite label** was selected, right-click the label operation in the **Operations Status** window to confirm intent to overwrite the existing volume label with a new label, and select **Supply Input**.

A question window appears displaying this message:

```
Label <labelname> is a valid NetWorker label. Overwrite it with a
new label?
```

10. Click **Yes** to overwrite the label with a new label, or **No** to cancel the label operation.

When a volume is relabeled, that volume is initialized and becomes available for writing again.

Verifying the Label when volume is unloaded

If a SCSI reset is issued during a backup, the volume will rewind and NetWorker may overwrite the volume label. To detect if the label is overwritten in this circumstance, set the **Verify Label On Eject** attribute in the Device resource, or the **Verify Label On Unload** attribute in the Jukebox resource to Yes. When either of these attributes is set to Yes, 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.”

Using barcode labels

The option to label a library volume with a barcode is available during automatic device configuration. [“How to configure devices” on page 48](#) provides more information. 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.

Configure a library to use volumes with barcodes

To select whether barcodes are used or matched after configuration:

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**.

[“Barcode labeling tips” on page 226](#) provides more information.

Use unmatched volume and barcode labels

Note: If unmatched volume and barcode labels are to be used, ensure that labels are attached to the outside of the volumes.

To use unmatched volume and barcode labels:

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.

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.

Mounting and unmounting volumes

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.

Mount or unmount a volume in a library

To mount a volume in a library:

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**.

Mount or unmount a volume in a stand-alone tape drive

To mount a volume in a stand-alone drive:

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 appropriate 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**.

Label and mount a volume in one operation (stand-alone tape drive)

When more than one storage device is connected to the NetWorker server, the device to be used 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.

To label and mount a volume in a single operation in a stand-alone tape drive:

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 appropriate 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 associated with the selected pool.

NOTICE

If the volume is unlabeled, the NetWorker server assigns the next sequential label from the label template associated with the selected pool. If a recyclable volume from the same pool is being relabeled, then the volume label name and sequence number remain the same. Access to the original data on the volume is destroyed, however, and the volume becomes available.

- b. Select a pool on the **Pools** menu. The NetWorker server automatically applies the label template 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 manual recycle, the NetWorker server disregards the assigned browse and retention policies. Therefore, only an administrator can mark the volume recyclable.

NOTICE

A volume that has been set to manual recycle retains that setting, even after relabeling. 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. [“Recycling volumes” on page 222](#) provides more information.

- 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 relabel 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.

[“Configure a library to use volumes with barcodes” on page 225](#) provides information on using barcode labels.

Note: If you are in the process of relabeling a mounted volume and you choose to not overwrite the existing label, the volume will be left in an unmounted state. To use this volume, mount it again.

Label 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.

Mount an uninventoried volume

To mount a volume that is not included in the library inventory, but which is a valid (properly labeled) NetWorker volume:

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. [“Supply user input” on page 472](#) 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 as described in [Chapter 1, “Overview”](#).

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.

Automatic unmounting of volumes (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.

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.

To change the Idle Device Timeout attribute for a library:

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.

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:

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. Right-click the device and select **Properties**.
5. Select the **Advanced** tab.
6. 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.

Using libraries with a 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.

To deposit a volume from the CAP into a library:

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 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

To withdraw a volume from a library slot and place it in the CAP:

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 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 its status.
7. Click **OK**.
8. Select **Monitoring > Log** to see the result. A successful **Eject/Withdraw** operation ends with a **Succeeded** comment in the Log.

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** command with the **-P** option to specify the port or ports from which to withdraw volumes. Refer to the **nsrjb** man page or the *EMC NetWorker Command Reference Guide* for more information.

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 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.

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 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. [“Using barcode labels” on page 225](#) provides more information on using barcode labels.

To inventory volumes in a library:

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.

Working with volumes

This section describes how to work with volumes.

Removing volumes from the media database and 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. A volume entry should be removed from the media database only if the volume has become physically damaged or unusable.

Both the client file index and media database entries can be removed. This action removes all information about the volume from the NetWorker server. Even if the database entries for a volume are removed, as long as the volume is undamaged, the data remains recoverable by using the **scanner** program.

In general, do not remove both the client file index and media database entries at the same time unless the volume is damaged or destroyed.

The presence of a clone of the 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. The volumes's entry in the media database is never actually purged. Because of this, removing volume entries from the media database is not a particularly effective way to reduce index size, although it does reduce the size of the online indexes by purging index entries associated with specific volumes.

Deleting volume data

To delete a volume:

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. The library's drives and mounted volumes are listed, as well as its slots and all volumes, mounted or unmounted.
 - Only unmounted volumes can be deleted.
4. Right-click the volume to be deleted, and select **Delete**. The **Delete** window appears and displays a request to select from where the volume should be removed:
 - a. File and media index entries
 - b. File index entries only

NOTICE

Do not remove the indexes of save sets on bad volumes.

5. Click the appropriate selection.
6. Click **OK**.
7. After a bad volume has been removed, type the **nsrck** command at the command prompt.

The **nsrmm** and **mminfo** UNIX man pages or the *EMC NetWorker Command Reference Guide* provide more information.

Marking a volume as full for offsite storage

When removing a volume from a library to store offsite, mark the volume as “full” so that the NetWorker software will not continue to ask for the volume.

To mark a volume as full, use the **nsrjb** command (for libraries) or the **nsrmm** command (for stand-alone drives) from the command-prompt. Note that the volume must be unmounted before this operation can be completed.

The format is as follows:

- ◆ For libraries:

```
nsrjb -o full valid
```

- ◆ For stand-alone drives:

```
nsrmm -o full valid
```

Where *valid* is the volume identifier of the volume. When a volume is marked as full, it is also marked as read-only.

You can also change the volume's Location attribute to an informational message, such as "Move to offsite storage in September 2009".

To change the volume's Location attribute:

1. In the **Administration** window, click **Media**. Media-related topics appear in the navigation tree.
2. Select **Volumes**. The **Volumes** detail table, which includes information about all of the server's volumes, appears.

3. Right-click a volume in the detail table, and select **Set Location**. The **Set Location** window appears.
4. Type a location description.
This is descriptive information. Filling in this field does not send a volume anywhere.
5. Click **OK**.

Cloning volumes

[“Volume cloning” on page 350](#) provides information about volume clone operations.

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-enable 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.

To reenable a device:

1. Once 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**.

If the disabled drive is part of a library, it might be necessary to reset the device. To do this:

1. From the command prompt, change the path to the directory that contains the NetWorker binaries.
2. 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.

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.

Numbering a silo slot

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 *EMC NetWorker Command Reference Guide* provide more information.

Mounting and unmounting silo volumes

Mount and unmount operations for silos are the same as for 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.

[“Mounting and unmounting volumes” on page 227](#) provides more information.

Labeling a silo volume

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 223](#) and [“Using barcode labels” on page 225](#) 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 a silo with volume import/export capability

NetWorker software supports the use of the import/export feature 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.

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 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 a large number of volumes.

On some silos (for example, StorageTek silos with the import/export feature set to automatic mode), the silo management software inserts volumes automatically. On these silos, 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 **nsrjb** UNIX man page or the *EMC NetWorker Command Reference Guide* provides more information.

Allocating (adding) silo volumes

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.

Add a silo volume

To add a silo volume:

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 its status.

8. Click **OK**. On return to the **Library** detail table, the added volumes will be shown.

Troubleshooting

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, such as DAS:

```
nsrjb -a -T tags -d
```

where *tags* specifies the tags or barcodes of volumes in a remote silo. The **-d** flag performs the manual deposit.

- ◆ On silos where the silo management software deposits volumes automatically, such as StorageTek silos:

```
nsrjb -a -T tags
```

[“NetWorker software interactions with a silo” on page 156](#) provides more information on STLIs.

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.

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.

To remove a silo volume:

1. Unmount the volume from the device. [“Mounting and unmounting volumes” on page 227](#) 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.

[“NetWorker software interactions with a silo” on page 156](#) provides information on STLs.

Inventorying a silo

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 233](#) 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.

Volume save sets

Information about individual save sets on volumes can be displayed from the Volumes detail table. Refer to this information to determine how resources are being used. For example, knowing the size of a save set can help in planning the amount of disk space needed for the online indexes.

Viewing save set details in the Volume Save Sets window

To view save set information in the Volume Save Sets window:

1. In the **Administration** window, click **Media**. Media-related topics appear in the navigation tree.
2. Select **Volumes**. The **Volumes** detail table, which includes information about all of the server's volumes, appears.
3. Right-click a volume in the detail table, and select **Show Save Sets**.

4. The **Volume Save Sets** window appears. [Table 27 on page 243](#) shows the attributes and their descriptions.
5. Click **OK** to close the **Volume Save Sets** window.

Table 27 Volume Save Sets window

Category	Description
Client	Name of the NetWorker client computer that created the save set.
Save Set	Pathname of the file system containing 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.
Time	Date and time when the save set was created.
Level	Level of backup that generated the save set. This refers only to scheduled backups. For manual backups, the level is blank.
Status	Type of save set. Table 47 on page 347 provides a listing of save set values and descriptions.
Size	Save set size, in appropriate units.
Flags	<p>First flag shows which part of the save set is on the volume:</p> <ul style="list-style-type: none"> • c: Completely contained on volume. • h: Spans volumes, and head is on this volume. • m: Spans volumes, and a middle section is on this volume. • t: Spans volumes, and the tail section is on this volume. <p>Second flag shows save set status:</p> <ul style="list-style-type: none"> • b: In the online index and is browsable. • r: Not in the online index and is recoverable. • E: Marked eligible for recycling and may be overwritten at any time. • a: Aborted before completion. Aborted save sets whose targets were AFTD or DD Boost devices are never shown in the Volume Save Sets window nor in mminfo reports because such save set entries are removed from the media database immediately. • i: Still in progress. <p>Optional third flag:</p> <ul style="list-style-type: none"> • N: NDMP save set • R: Raw partition backup (such as for a supported module). • P: Snapshot <p>Optional fourth flag:</p> <ul style="list-style-type: none"> • s: NDMP save set backed up by nsrdsa_save command to a NetWorker storage node.

Changing save set status within the Volume Save Sets window

To change the save set status within the **Volume Save Sets** window:

1. Select a save set.
2. Click **Change Status**. The **Change Save Set Status** window appears.
3. Select either:
 - Normal (default)
 - Suspect
4. Click **OK**, to leave the **Change Save Set Status** window.

5. Click **OK**, again, to leave the **Volume Save Sets** window.

Viewing save set details from the Save Set detail table

To view save set details from the Save Sets detail table:

1. In the **Administration** window, click **Media**. Media-related topics appear in the navigation tree.
2. Select **Save Sets**.

The **Save Sets** detail table appears with two tabs for configuring save set queries and listing save set details:

- The **Query Save Set** tab
 - The **Save Set List** tab
3. Select either tab. [“The Query Save Set tab” on page 244](#) and [“The Save Set List tab” on page 245](#) provide more information.

The Query Save Set tab

The **Query Save Set** tab allows users to search for save sets that meet specific criteria. Click the **Query Save Set** tab to access these query fields:

- ◆ Query Parameters area:
 - Client Name
 - Save Set
 - Save Set ID
 - Volume
 - Pool
 - Copies
 - Save Time (a range)
- ◆ Status area:
 - All
 - Select from:
 - Browsable
 - Scanned-In
 - In-Progress
 - Suspect
 - Recyclable
 - Recoverable
 - Aborted

- ◆ Type area:
 - All
 - Select from:
 - Normal
 - Deduplication
 - NDMP
 - Snapshot
 - Raw
- ◆ Maximum level area:
 - Full
 - 1 through 9
 - All

If no save sets are found that match the query parameters, an error message appears when closing the tab:

`No save sets were found that matched the specified query.`

The Save Set List tab

The **Save Set List** tab lists detailed save set information.

Click the **Save Set List** tab to view this tabular information:

- ◆ Save Set
- ◆ SSID (Save Set ID)
- ◆ Level
- ◆ Status
- ◆ Volume Name
- ◆ Type
- ◆ Client
- ◆ Size
- ◆ Files
- ◆ Pool
- ◆ Time

CHAPTER 7

Backup Groups and Schedules

This chapter covers these topics:

- ◆ Overview of NetWorker scheduling 248
- ◆ Backup groups 248
- ◆ Managing backup groups 257
- ◆ Backing up open files 259
- ◆ Schedules 260
- ◆ Backup levels 267

Overview of NetWorker scheduling

Together, the following two items enable the scheduled backup of client data:

- ◆ Group
- ◆ Schedule

Time-based groups (backup groups) specify either the time of day when a backup occurs, or a probe-based backup that is user defined.

For time scheduled backups, times typically occur after regular work hours. All clients assigned to a group will be backed up at the time specified by the backup group. Schedules enable you to specify the day of the week or month that the backup occurs, as well as the level of backup (full, incremental, synthetic full or level 1-9).

For probe-based backups the probe interval and backup window are used to schedule group probes with clients, and clients with groups. The execution of the probes determines if the backup of the group will proceed.

Note: Each client in a group can have a probe associated with it, but a probe is not required. However, a probe-based backup group must have at least one probe-enabled client associated with it.

Backup groups

Time-based backup groups specify the starting time for a client's scheduled backup. These backup groups enable you to:

- ◆ Schedule the backups to take place in the middle of the night, or some other time when network traffic is low.
- ◆ Balance the backup loads by grouping clients in specific groups and staggering their start times.

NOTICE

Do not place both regular and deduplication clients in the same group.

- ◆ Sort data to specific backup volumes.

To sort data, groups are used in conjunction with backup pools. [Chapter 10, "Sorting Backup Data"](#) provides more information.

The NetWorker server and time-based backup groups

When a Client resource is created, it is assigned to a backup group. The clients in each time-based backup group begin their automatic scheduled backups according to the start time of the group. Backup loads are balanced by taking the client's backup schedule into account when determining which clients to include in a specific group.

Example 7 Using groups to balance client backups

[Figure 17 on page 250](#) illustrates how the NetWorker server uses two time-based backup groups to back up multiple client save sets. In [Figure 17 on page 250](#), three client computers (mars, jupiter, and saturn) are part of a group named Weekly Full. The Weekly Full group starts its automatic scheduled backup at midnight.

- ◆ Client mars runs a full backup of all its save sets every Monday and incremental backups of its save sets on the other days.
- ◆ Client jupiter runs a full backup of all its save sets on Tuesday and incremental backups on the other days.
- ◆ Client saturn runs a full backup of all its save sets on Wednesday and incremental backups on the other days of the week.

Because each client runs its full backup on a different day of the week, the server is not overloaded.

The second group, Accounting, illustrates how you can group clients by department. The Accounting group contains client computers *mercury* and *venus* and starts its backups at 7:00 P.M., when the computers in that department are available for backup. Although the two client computers run full backups on the same day, computer *venus* is scheduled to perform a full backup on only the `/usr/home` save set, whereas all the save sets on computer *mercury* are backed up. By estimating how long a backup takes, you can determine the start time to set for the next group.

The save sets from each group are written to appropriate volumes mounted on storage devices. The NetWorker server uses pools to organize, track, and store save sets. The server uses groups to determine the time clients start their scheduled backups.

Group:	Weekly Full	Group:	Accounting
Daily Start Time:	24:00 (Midnight)	Daily Start Time:	19:00 (7:00 P.M.)
Client:	mars	Client:	mercury
Save Set:	All	Save Set:	All
Schedule:	Full on Monday Incremental, all other days	Schedule:	Full on 1st of month Level 5 on 15th of month Incremental on all other days
Client:	jupiter	Client:	venus
Save Set:	All	Save Set:	/usr/home
Schedule:	Full on Tuesday Incremental on all other days	Schedule:	Full on 1st and 15th of month Incremental on all other days
Client:	saturn		
Save Set:	All		
Schedule:	Full on Wednesday Incremental on all other days		




Figure 17 How NetWorker groups are used to back up multiple clients

Preconfigured groups

The NetWorker product ships with a single preconfigured group named Default. To ensure that all data is backed up, the NetWorker server automatically adds all clients to the Default group. However, you must enable the Default group for the NetWorker server to back it up. You can keep a client in the Default group, or you can put the client in one or more customized groups.

You can also make changes to any Default group attribute, but you cannot delete the group. You can, however, create or delete as many customized groups as required.

Key Group attributes

Table 28 on page 251 displays some of the key attributes for the Group resource.

Table 28 Group attributes (1 of 2)

Attribute	Description
Start Time	The Default group is preconfigured to start its daily backup at 3:33 A.M. This time can be changed.
Autostart	Specifies whether the group is started automatically at a designated start time. You must enable the Autostart attribute for the Default group, and any other group you create, before a scheduled backup can be run.
Schedule	This optional attribute can be used to select a Schedule resource for the group. If this attribute is set, it overrides the schedules selected in the Client resource's Schedule attribute for all clients in this group.
Interval	This attribute dictates how often a group starts a scheduled backup. The default value is 24:00 (once a day), but you can change this value to start backups more often. "Setting the backup group time interval" on page 255 provides instructions on modifying this attribute.
Autorestart	Specifies whether the group will be automatically restarted after an incomplete backup due to a power failure or administrator intervention. If this attribute is enabled, the backup will restart when the NetWorker server is restarted provided that the period of time specified in the Restart Window attribute has not elapsed.
Restart Window	For either auto or manual restarts, this attribute specifies the period of time in which an incomplete backup can be restarted. If the period of time has elapsed, the restart will be treated as a regular backup start operation. The restart period is calculated from the beginning of the start of the last incomplete backup. The default value is 12:00 hours.
Client Retries	When the NetWorker server fails to connect to a client, this attribute specifies the number of times that the server will reattempt the connection before the backup is considered a failure. The first retry will not occur until after an attempt has been made to at least contact each client in the group.
Inactivity Timeout	This attribute specifies the maximum time, in minutes, that a client is given to fail to communicate back to the server. If a client fails to respond beyond the Inactivity Timeout value, the server will consider the client as having failed. If a client fails due to any reason, a retry is initiated immediately. This ensures that no time is lost during the scheduled backup due to any failures. Note: For large save sets, for save sets with large sparse files, and for incremental backups of a large number of small static files, increase the timeout value if the backup consistently aborts due to an inactive job.
Soft runtime limit	This attribute indicates the time in minutes since the start time of a given group after which no new child process will be launched. The Soft runtime limit is measured for each savegroup separately. Index and bootstrap saves are exempt and will be started regardless of this setting. The default value is 0, which indicates that no Soft runtime limit is in effect.
Hard runtime limit	This attribute indicates the time in minutes that any save session still running is ended and the savegrp program stopped. The default value is 0, which indicates that no Hard runtime limit is in effect.
Success Threshold	This attribute sets the criteria for reporting the success of all save sets within a group. The default value is Warning which means if any save set completes with warnings it will be reported as successful. The client will also be reported as successful with warnings in the completion report. If set to "Success", any save sets completed with warnings will be reported as failures. The client will also be reported as failed in the savegroup completion report. Note: Any failures will attempt a retry of the save set if the retry count is not 0.
Probe based backup	If this attribute is set to ON, the probe attributes listed below become enabled. Boolean ON/OFF. It is OFF by default.
Probe interval	This attribute indicates how often probes (in minutes) should be run. Default value is 60 minutes.

Table 28 Group attributes (2 of 2)

Attribute	Description
Probe start time/probe end time	Probe start time and probe end time together define the backup window. Probe end time minus probe start time should be greater than the probe interval. Start time default: 0:00; end time default: 23:59
Probe success criteria	This attribute determines if all probes or only one probe needs to succeed for a backup to proceed. Values are Any or All.
Time since successful backup	If the value is 0, the time since the last successful backup does not matter: the savegrp program always runs probes. If the interval is specified and reached, savegrp runs the backup no matter regardless of the result of all other probes. The probes are run so that the probe state data can be updated. The default value is 0 days.
Time of the Last Successful Backup	Set to the time of the last successful backup by the savegrp program. Used to calculate interval since the last successful backup. GUI read only.
savegrp Parallelism	Maximum number of save sets that can be backed up simultaneously by a NetWorker group. The default value is 0, which means that parallelism is not restricted. “Parallelism” on page 553 provides information about savegrp parallelism.

Probe Group

Probe-based backup groups specify probe interval, and backup window to schedule the group.

Probing occurs continuously throughout the probing window (the hours defined Probe start time, and Probe end time), and only when the Autostart attribute of the save group is enabled. If a save group is started manually, probes run immediately. If autostart is used, then probes only run during the specified probe window.

Clients are associated with probe-based backup groups in the same manner as they are with regular backup groups. However probe-based backup groups must include at least one client which references a probe resource as described in [“Creating a client probe” on page 610](#).

Each client can reference only one probe but, since backup groups can contain many clients, multiple probes can be run with the group.

Instead of a start time, the probe start and end times are used to schedule the group. It is the outcome of the probing which determines if the backup will proceed.

Note: Each client in a group can have a probe associated with it, but a probe is not required. However, a probe-based backup group must have at least one probe-enabled client associated with it.

Configuring a probe-enabled group

To configure a probe-enabled group:

1. In the server **Administration interface**, click **Configuration**.
2. Right-click **Groups**, and select **New**.

- In the **Advanced** tab of the **Create Groups** window, complete the fields in the **Probe** section as described in [Table 29 on page 253](#).

Table 29 Probe group fields

Field	Description
Probe-based group:	Click the checkbox to enable probing.
Probe interval:	Determines the frequency of probing. Can be set from a minimum of 15 to a maximum of 10,000 minutes. Note: A successful backup does not disable probing. Be sure to set the Probe interval to an appropriate value.
Probe start time:	The time at which probing will begin.
Probe end time	The time at which probing will end.
Probe success criteria	<ul style="list-style-type: none"> Any: Any one of the probes associated with the group must succeed for the backup to be performed. All: All of the probes associated with the group must succeed for the backup to be performed.
Time since successful backup	The longest period of time tolerated without a backup.

Aborted backup groups

If the backup of a save set fails, then the NetWorker server marks the save set as ABORTED. In this situation, the automated report from the **savegrp** program does not always show that the backup has completed. For example, if the client is being backed up over a Network File System (NFS) connection and the NFS server crashes and reboots, the NetWorker backup hangs until it times out. The NetWorker server marks the save set ABORTED.

Note: Aborted save sets whose targets are AFTD or DD Boost devices are never shown in the Volume Save Sets window nor in **mminfo** reports because such save set entries are removed from the media database immediately.

How to create a group

[“Task 2: Set up a group for backup clients” on page 61](#) provides information about creating a group.

How to edit a group

NOTICE

You cannot change the name of an existing backup group.

To edit a group:

- From the **Administration** window, click **Configuration**.
- In the expanded left pane, select **Groups**.

3. Select the group to be edited.
4. From the **File** menu, select **Properties**.
5. Edit the attributes of the group and click **OK**.

How to delete a group

NOTICE

You cannot delete the preconfigured Default group nor any group currently applied to a Client resource.

To delete a group:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Groups**.
3. Select the group to be deleted.
4. From the **File** menu, select **Delete**.

How to copy a group

To copy a group:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Groups**.
3. In the right pane, right-click the group to be copied and select **Copy**. The **Create Group** dialog box appears, containing the same information as the group that was copied, except for the Name attribute.
4. In the **Name** attribute, type a name for the new group.
5. Edit the attributes and click **OK**.

Copying a group with clients

The Copy with Client feature allows you to copy an existing group resource including the associated group and all client resources. The Copy with Clients operation enables the following actions:

- ◆ Copy an existing NSR group.
- ◆ Ensure the original client list is preserved in the new group.
- ◆ All NSR client resources are automatically updated.

Note: The Copy with Clients operation is only available to a NSR group resource type. Consequently, this functionality is only available if a NetWorker group is selected in the configuration window in the NetWorker console.

The Copy with Clients option is available in under the Edit menu. The option is also available in a pop-up menu that appears when an individual group is selected in the details pane or navigation tree.

Setting the backup group time interval

The NetWorker server allows you to run an individual scheduled backup group more than once within a 24-hour period. The Interval attribute value of the Group resource determines the frequency (in hours) that an individual group will start a backup.

The default value is 24 hours (24:00), which results in one backup group run per day. If you set the Interval attribute value at 12 hours, then the same group will back up twice a day. For example, a group with the default start time of 3:33 A.M. and an interval of 12:00 would back up twice a day, first at 3:33 A.M., and then again twelve hours later at 3:33 P.M.

To set backup group time intervals:

1. Select the group to edit. For information about editing a group, see [“How to edit a group” on page 253](#).
2. Select the **Advanced** tab.
3. In the **Interval** attribute, type a value in the hh:mm format.

For best results, use time interval values that make it easy to determine the backup group time, such as 24, 12, or 6 hours.

4. Click **OK**.

Note: An increase in the backup group time interval (for example, changing the interval from once every 24 hours to once every 12 hours) can add strain on a network, the NetWorker server, and associated devices.

Limiting full backups when the time interval is less than 24 hours

For groups that have more than one scheduled backup within a 24-hour period, use the Force Incremental attribute to prevent more than one full or level backup per 24-hour period. By default, the Force Incremental attribute is set to Yes. If the Force Incremental attribute is set to Yes, the first backup is performed at the configured level. All subsequent scheduled backups during the next 24 hours after the start of the first backup will be incremental. This means that only changed files will be backed up regardless of the configured level. The Force Incremental attribute applies only to scheduled backups that the NetWorker server runs automatically. If the **savegrp** program is run by other means, such as from the command prompt or a script, this attribute is not used.

If the **Force Incremental** attribute is set to NO, multiple full or level backups are allowed during the 24 hours after the start time of the first backup.

Forcing an incremental backup

To force incremental backups on groups:

1. Select the group to edit. [“How to edit a group” on page 253](#) provides information about editing a group.
2. Select the **Advanced** tab.
3. Select the **Force Incremental** attribute and click **OK**.

[“Setting the backup group time interval” on page 255](#) provides information about how to configure backup groups to occur more than once every 24 hours.

Running a backup group from the command line or a script

Instead of scheduling a backup group through the NetWorker Administration window or through the `nsradmin` program, you can back up a group directly from the command line or from a script by using the `savegrp` program. However, there are some considerations that you need to be aware of when running the `savegrp` program in this way.

When the `savegrp` program is executed directly from the command line or from a script, some attributes that may have been specified for the backup group resource through the NetWorker Administration Group Properties window or the `nsradmin` program will not take effect. You must specify these options directly on the command line or in the script if you want to include these options with your group backup.

For example, suppose that in the NetWorker Administration window you have set up a backup group resource named *Accounting*. And you specified that the Accounting backup group use a schedule named *Full Every Friday*. If you back up the Accounting group from a command line or a script, then you must specify the schedule explicitly by using the `-C` option, for example:

```
savegrp -C "Full Every Week" Accounting
```

[Table 30 on page 256](#) lists the backup group resource attribute and the corresponding `savegrp` option that must be specified explicitly on the command line or in a script, if they are required for your backup group. For more information on these options, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

Table 30 `savegrp` options that are not taken from backup group resource attributes (1 of 2)

For this attribute from the Group Properties window or <code>nsradmin</code> program...	Use this <code>savegrp</code> program option	Description
Level	-l	Backup level
Force incremental	-l incr	Forces an incremental level backup. The Force incremental attribute behaves differently than the <code>-l incr</code> option. The Force incremental attribute only forces an incremental if the group is scheduled to backup more than once in a 24 hour period. “Limiting full backups when the time interval is less than 24 hours” on page 255 provides more information. The <code>-l incr</code> option forces an incremental backup every time.
Printer	-P	Printer to which bootstrap information will be printed
Schedule	-C	Backup schedule
Schedule time	-t	An explicit time can be specified when looking at a schedule to determine what level of save to perform. No value means use the current date to determine the level.
Savegrp parallelism	-N	Parallelism for the backup group. The parallelism attribute will be taken from the backup group resource if no value is specified with <code>-N</code> or if the value specified with <code>-N</code> is 0 (unlimited parallelism for the backup group). “Parallelism and multiplexing” on page 553 provides more information on how various NetWorker parallelism attributes work together.

Table 30 savegrp options that are not taken from backup group resource attributes (2 of 2)

For this attribute from the Group Properties window or nsradmin program...	Use this savegrp program option	Description
Expiration time	-w -y	-w specifies browse time -y specifies retention time
Verbose	-v	savegrp job information detail level
Estimate	-E	Estimate amount of save data before performing a save
No save	-n	Estimate amount of save data but do not perform a save
No monitor	-m	Do not report on status of savegrp operation
Index only	-O	Save only each client index and bootstrap
Preview	-p	Run a probe step on each client but do not save data
Revert to full when synthetic full fails	-F	Automatically perform a full level backup if a synthetic full backup operation fails.
Verify synthetic full	-V	If a synthetic level backup is performed, verify that synthetic full save sets are indexed (available for browsing) after the save sets are created.

Managing backup groups

[Table 31 on page 257](#) lists backup group management tasks and where to find more information about these tasks.

Table 31 Backup group management

Group Management Task	For more information
Start a group immediately	“Starting a group immediately” on page 466
Stop a group immediately	“Stopping a group” on page 466
Restart a group	“Restarting groups” on page 466
Preview a backup group	“Previewing a backup group” on page 257

Previewing a backup group

You can simulate a backup for a specific group. This feature generates an output that includes this information:

- ◆ File System to be backed up
- ◆ Backup level
- ◆ Backup pool

Preview a backup to identify potential problems before the NetWorker server runs a backup group.

To preview a backup group:

1. Select the group to edit. Information about editing a group is available in [“How to edit a group” on page 253](#).

2. Select the **Advanced** tab.
3. In the **Option** attribute, select **Preview** and click **OK**.

To see the simulated results of the backup, open the daemon log file located in the <NetWorker_install_dir>\logs directory. Information about viewing log files is available in [“Viewing log files” on page 803](#).

Moving clients between groups

Multiple clients may be moved from one group to another by selecting the clients and dragging them to another group.

NOTICE

Do not place both regular and deduplication clients in the same group.

Estimating save set sizes of a backup group

You can estimate the size of the save sets that will be produced in a backup group before the backup is run.

To estimate save set sizes of the backup group:

1. Select the group to edit. Information about editing a group is available in [“How to edit a group” on page 253](#).
2. Select the **Advanced** tab.
3. In the **Option** attribute, select **Estimate, No Save, and Verbose**.
4. Click **OK**.

To see the estimated save set sizes, open the daemon log file located in the *NetWorker_install_dir*\logs directory. [“Viewing log files” on page 803](#) provides information about viewing log files.

Note: Selecting the Estimate, No Save, and Verbose options produces an estimate that shows all paths and filenames that will be saved in the backup group. Selecting the Estimate and Verbose options produces a detailed estimate and performs the save in a single operation. Selecting only the Estimate option (without the No Save or Verbose option) results in a save operation without an estimate.

Backing up status reports

When the backup is completed, several types of backup status reports are generated:

- ◆ [Chapter 15, “Enterprise reporting and events monitoring”](#) provides information about these reports.
- ◆ Information about the status of backed-up groups is also written to the savegrp log file, which is located in the *NetWorker_install_dir*\logs directory.
- ◆ [“Viewing log files” on page 803](#) provides information about viewing log files.

Generating and printing bootstrap reports

When the backup group includes the NetWorker server, or if the server is not in an active group, the server generates a special save set called the *bootstrap*, which includes the media database and configuration files. In both of these cases, a bootstrap email (default) or printout is generated whether the scheduled backup is initiated automatically or manually. The bootstrap information is essential for recovery from a disaster. For information on how the bootstrap is used during a disaster recovery operation, refer to the *NetWorker Procedure Generator*.

Note: If the NetWorker server is not a member of an active group, the bootstrap is created when any group backup is run, even if the group has a level of SKIP. However, in this case the bootstrap will be created only once every 24 hours, regardless of how many groups are run during that period. If you would like to create bootstraps for every group backup, you should include the server in the group, with a very small save set (such as */etc/hosts*).

By default, the bootstrap reports are generated and sent as an email to the default email recipient, either the administrator or root. To change the email recipient, open the Bootstrap notification and configure a new email recipient.

Note: User can also choose to get the reports printed through the default printer configured for the NetWorker server. To change the default printer, edit the **Printer** attribute in the Group resource.

If the bootstrap notification is configured for email (default option) and an email recipient is not configured, then the bootstrap reports are lost. However, when an email recipient is later configured, the bootstrap reports are generated the next time as part of the **savegrp** operation and the previous reports are also sent to the email recipient along with the current report.

If the bootstrap notification is configured to the printer (not the default configuration) and bootstrap report fails for any reason, then the contents can be viewed in the *savegrp.log* file which is located in the *<install_dir>/logs* directory, or the savegrp report. [“Viewing log files” on page 803](#) provides information about viewing log files.

Backing up open files

Open files are a problem that all data backup applications must solve. Open files that are not backed up properly represent a potential data loss. They might be skipped, improperly backed up, or locked.

NetWorker can open two different types of files. Those that are owned by the operating system and those that are owned by a specific application.

Opening 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 message:

```
save: file_name cannot open
```

Additionally, a permission denied error may be returned from the operating system.

To back up locked open files, close any open files. However in most cases, this is impractical. 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. [Chapter 2, “Backing Up Data”](#) provides more information. Also use Open File Manager to back up open files.

Opening files owned by a specific application

The NetWorker software by itself 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 Oracle to back up open files in an Oracle database.

Files that change during backup

If a file changes during a backup, the NetWorker software alerts you by displaying the following message in the **Groups** tab of the **Monitoring** option:

```
warning: file_name changed during save
```

To ensure that the changed file is backed up, do either of the following:

- ◆ Restart the backup group.
- ◆ Perform a manual backup. Information is available in [“Manual backups” on page 70](#).

Note: NetWorker Modules can back up these types of files correctly, if they are files related to the database the module is backing up.

Backing up open files with VSS

In NetWorker releases 7.6 and later, the software takes advantage of VSS technology to create snapshot backups of volumes and exact copies of files, including all open files. In this way, files that have changed during the backup process are copied correctly. [Chapter 26, “Volume Shadow Copy Service,”](#) provides more information about VSS.

Schedules

Each Client resource is backed up according to a schedule. A Client resource’s backup schedule tells the NetWorker server what level of backup (for example, full, incremental, or synthetic full) to perform on a given day. For instance, on Fridays it might perform a full backup on a Client resource and the rest of the week perform incremental backups. The time of day the backup begins is determined by the group to which the Client resource is associated.

Schedules can be simple or complex, depending on the needs of the environment. All Client resources can share the same schedule, or each can have a unique schedule.

The type and scope of the backup is determined by the specified backup level. The level can be set to back up a client's entire file system, or only data that has changed since the last full backup.

[“Backup levels” on page 267](#) provides information about backup levels.

Schedules for Avamar deduplication clients

Backups must be scheduled to avoid the Avamar node's read-only periods when such cron jobs as checkpoint and garbage-collection are run. The Avamar server documentation provides more information.

The *NetWorker Avamar Integration Guide* provides backup level and schedule information that is specific to Avamar clients.

Preconfigured NetWorker schedules

The NetWorker software ships with preconfigured schedules. If these schedules meet backup requirements, use them as is. Otherwise, create new schedules to accommodate any site-specific needs.

Preconfigured schedules cannot be deleted. Preconfigured schedules that contain “overrides” (indicated by an asterisk next to a backup level in the schedule's calendar) cannot be modified. All other preconfigured schedules can be modified.

[Table 32 on page 261](#) describes the preconfigured schedules.

Table 32 Preconfigured NetWorker schedules (1 of 2)

Schedule name	NetWorker backup operation
Default	Completes a full backup every Sunday, incremental backups on all other days.
Full Every Friday	Completes a full backup every Friday, incremental backups on all other days.
Full on First Friday of Month	Completes a full backup on the first Friday of the month, incremental backups on all other days. This schedule cannot be modified.
Full on First of Month	Completes a full backup on the first calendar day of the month, incremental backups on all other days.
Quarterly	Completes a full backup on the first day of a quarter. Performs a level 5 backup on the first day of the other months in the quarter. Every seven days, a level 7 backup occurs. Incremental backups are performed on all other days. This schedule cannot be modified.
Incremental and Synthetic Full 1st Friday of Month	Completes an incremental backup prior to performing a Synthetic Full backup on the first Friday of every month. Completes incremental backups on all other days. This schedule cannot be modified.

Table 32 Preconfigured NetWorker schedules (2 of 2)

Schedule name	NetWorker backup operation
Incremental and Synthetic Full Every Friday	Completes an incremental backup prior to performing a Synthetic Full backup on every Friday. Completes incremental backups on all other days. This schedule cannot be modified.
Incremental and Synthetic Full 1st of Month	Completes an incremental backup prior to performing a Synthetic Full backup on the first calendar day of the month. Completes incremental backups on all other days. This schedule cannot be modified.
Incremental and Synthetic Full Quarterly	Completes an incremental backup prior to performing a Synthetic Full backup on the first day of each quarter. Completes incremental backups on all other days. This schedule cannot be modified.

Backup cycles

The period of time from one full backup to the next full backup is called a backup cycle.

The following examples demonstrate how to use schedules for different backup cycles and client backup needs.

Example 8 Weekly backup cycle

Figure 18 on page 262 illustrates a weekly backup cycle. In this example, a full backup is performed on a client each Sunday, and incremental (Inc) backups are performed on the other days of the week.




Figure 18 Weekly backup cycle

Use backup schedules to balance and stagger the load on a NetWorker server. Depending on the size of a network, you could apply the same schedule to all clients. For example, if no one works over the weekend and you want to run full backups during this time, you could apply the Default schedule to all of the clients.

The Default schedule tells the NetWorker server to perform full backups on Sunday, and incremental backups the rest of the week.

Example 9 Default schedule for multiple clients

Figure 19 on page 263 illustrates how the Default schedule works for three clients.




Figure 19 Default schedule for multiple clients

NOTICE

If you have a short backup window period and need to create a full backup, you can use the synthetic full backup. “Synthetic full backups” on page 76 and Example 11 on page 264 provide details.

Since full backups transfer large amounts of data and typically take longer than other backup levels, you may want to stagger them throughout the week. For example, you could apply a schedule that performs a full backup for Client A on Thursday, a second schedule that performs a full backup for Client B on Tuesday, and a third schedule that performs a full backup for Client C on Sunday.

Example 10 Staggered weekly schedules for multiple clients

Figure 20 on page 263 illustrates how to use a staggered backup schedule for multiple clients.




Figure 20 Staggered weekly schedules for multiple clients

By balancing and staggering the load, and using different start times for different groups of clients, you can increase the efficiency of a NetWorker server.

Example 11 Weekly synthetic full backup cycle to reduce backup window

Figure 21 on page 264 illustrates a weekly full synthetic backup cycle. In this example, a synthetic full backup is performed on a client each Sunday, and incremental backups are performed on the other days of the week.




Figure 21 Weekly synthetic full backup cycle to reduce backup window

Scheduling and planning considerations

Deciding which schedules are most appropriate for an environment requires planning.

When you create backup schedules, consider:

- ◆ How much data do you have to back up?
- ◆ How many backup media volumes do you plan to use?
- ◆ How much time do you have to complete a backup?
- ◆ Does it matter how many volumes are required to recover from a disaster, such as a disk crash?

Additionally, determine a policy for recovering files. For example, if users expect to be able to recover any version of a lost file that was backed up during a three-month period (that is, the retention policy is three months), you need to maintain all of the backup volumes for a three-month period. On the other hand, if users expect to be able to recover data from only the last month, you will not need to maintain as many volumes.

The length of time that data is available for recovery by the NetWorker server is determined by the browse and retention policies associated with each client. [Chapter 8, “Browse and Retention Policies”](#) provides information about browse and retention policies.

NOTICE

If you have a short backup window period and need to create a full backup, you can use the synthetic full backup. [“Synthetic full backups” on page 76](#) and [Example 11 on page 264](#) provides details.

Scheduling 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. Consequently, it might not be convenient to perform a scheduled, full backup for client save sets as large as this because of the amount of time required.

Schedule the client's disk volumes for backup at different times by separating them into different backup groups. When you split one client's save sets into multiple backup groups, you back up all the client's files, but not all at once. It is less time-consuming than a full backup of all the local data at one time.

To back up the client's file systems individually, add and configure the same client several times by addressing the different file systems in the Client resource. For example, configure the first Client resource to back up one file system () with a single backup schedule in one group. Then, configure the second Client resource to back up another file system () with a second backup schedule in another group.

NOTICE

When you create separate backup schedules and explicitly list save sets, any files or file systems not included in that list are omitted from backup. This includes any new disk volumes that are added to the system. To avoid this risk, type the value All in the Save Set attribute.

Key components of a schedule

[Table 33 on page 265](#) describes the key components of a Schedule resource.

Table 33 Key components of a schedule (1 of 2)

Attribute	Description
Name	The name assigned to a customized schedule that appears in the Client resource as an attribute, and can be applied to a client/save set. Assign a simple, descriptive name such as Monday Full.
Period	Determines how often a full backup is to run. Set the schedule to apply to either a weekly or a monthly period. When you select Week and set up a schedule, the backup level full is applied to the same day of the week for all weeks in the calendar year. For example, full backups every Sunday. Week is the default setting. When you select Month and set up a schedule, the backup level full is applied to the same day of the month for all months in the calendar year. For example, full backups on the fifteenth of each month.

Table 33 Key components of a schedule (2 of 2)

Attribute	Description
Calendar	<p>Displays the days of the month and the backup level scheduled for each day. In addition to full and incremental backups, you can set intermediate backup levels. You can include one or more of these levels in a backup schedule:</p> <ul style="list-style-type: none"> • Full • Incremental • Level (1 – 9) • Synthetic Full+Incr • Skip <p>The Override Levels option allows you to override an existing backup level for a specific day. For example, you might not want a full backup to run on a holiday. You can override the schedule so 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.</p> <p>If you are overriding 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 man page or the <i>EMC NetWorker Command Reference Guide</i> contains more information about overriding backup levels.</p>

NOTICE

The Force Incremental attribute in the backup group, determines the level used by the NetWorker server when there is more than one backup per day. The default value for this attribute is Yes, which means that an incremental backup will occur if the group is run more than once a day. To perform more than one full or level backup per day, set this attribute to No. [“Limiting full backups when the time interval is less than 24 hours” on page 255](#) provides more information.

Working with schedules

This section provides information on how to edit, delete, and copy schedules.

[“Task 1: Set up a schedule for backups” on page 60](#) provides information about creating a schedule.

Editing a schedule

To edit a schedule:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Schedules**.
3. In the right pane, select the schedule to edit.
4. From the **File** menu, select **Properties**.
5. Edit the attributes and click **OK**.

Deleting a schedule

You cannot delete preconfigured schedules or schedules that are currently selected in a client’s Schedule attribute.

To delete a schedule:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Schedules**.
3. In the right pane, select the schedule to delete.
4. From the **File** menu, select **Delete**.
5. When prompted, click **OK** to confirm the deletion.

Copying a schedule

To copy a schedule:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Schedules**.
3. In the right pane, right-click the schedule to be copied and select **Copy**. The **Create Schedule** dialog box appears with the same information as the schedule that was copied, except for the Name attribute.
4. In the **Name** attribute, type a name for the new schedule.
5. Edit the attributes and click **OK**.

Overriding a client's regular backup schedule

You can set the Schedule and Level attributes in the group to override a client's regular backup schedule. For example, one evening you could run a full backup on all the clients in a group, regardless of the clients' regular backup schedules. The value specified in the group's Level attribute overrides the backup level setting for every client in the group.

Alternatively, you could have a group of clients follow the same backup schedule instead of each client's individual schedule. You could assign a group of clients to follow the default schedule (full every Sunday) regardless of each client's individual schedule. If you leave the group's Level and Schedule attributes blank (the default setting), clients follow the backup schedule assigned in the Client resource.

Disabling or enabling a client backup schedule

By default, the schedule assigned to the backup client is enabled.

To disable scheduled backups for a client:

1. Open the **Client** resource whose scheduled backups are to be disabled. [“Editing a client” on page 606](#) provides more information.
2. Clear the **Scheduled Backup** attribute and click **OK**.

Backup levels

Because it may not be practical or efficient to run full backups every day, you can specify the level of the backup to be performed during scheduled backups. By limiting the frequency of full backup, you help maintain server efficiency, while still ensuring that data

is protected. Different backup levels enable you to trade off the number of volumes and amount of time required to complete a backup with that required to recover from a disk crash.

[Table 34 on page 268](#) describes the five kinds of backup levels:

Table 34 Backup levels

Backup Level	Function
Full	Backs up all files, regardless of whether or not they have changed.
Level [1 – 9]	<p>Backs up files that have changed since the last backup with a lower-numbered backup level. For example:</p> <ul style="list-style-type: none"> • A level 1 backup backs up all files that have changed since the most recent full backup (considered a level zero). • A level 3 backup backs up all files that have changed since the most recent backup at level 2, level 1, or full. For example, if the most recent backup was at level full, then a level 3 backup will back up all files that changed since the full backup. However, if the most recent backup was at level 2, then a level 3 backup will back up only those files changed since the level 2 backup. • A level 9 back up backs up all the files that have changed since the most recent backup of any level except level 9. <p>Note: The NetWorker software ignores any incremental-level backups when determining what files should be backed up.</p>
Incremental	Backs up files that have changed since the last backup, regardless of level.
Synthetic Full	Backs up all data that has changed since last full backup and subsequent incrementals to create a synthetic full backup. “Synthetic full backups” on page 76 provides more information.
Synthetic Full+Incr	<p>Perform an incremental backup and a synthetic full backup on the same day in the same group.</p> <p>Use this level when synthetic full backups fail because the incremental backup chain is broken. This level triggers an incremental to be created before the synthetic full backup. This mends the broken chain of incremental backups.</p> <p>“Synthetic full backups” on page 76. provides more information</p>
Skip	Skips the scheduled backup. For example, you can skip a backup on a holiday if no one will be available to change or add more media volumes.

Note: Information on the special nature of deduplication backups is available in the *NetWorker Avamar Integration Guide* and the *NetWorker Data Domain Deduplication Devices Integration Guide*.

How NetWorker backup levels work

Backup levels work in conjunction with a client’s backup schedule. The way you define the backup levels directly affects how long a recovery from a disk crash takes and how many backup volumes you need.

Planning level backups helps to minimize the number of volumes used. The fewer volumes required to recover from a disk crash, the less time spent restoring the disk.

You can also reduce the size and 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. More information on directives is available in [Chapter 9, “Directives.”](#)

The following three examples illustrate the how backup levels affect the requirements for data recovery.

Example 12 Backup levels (part 1)

As shown in [Figure 22 on page 269](#), a full backup runs on October 1. On October 2, an incremental backup saves everything that changed since the full backup. On October 3, another incremental backup backs up everything that changed since October 2. Then, on October 4, a level 7 backup backs up everything that changed since the full backup on October 1.

To fully recover from a disk crash on October 4, you need the data from the full backup from October 1 and the new level 7 backup. You no longer need the data from October 2 and 3, because the level 7 volume includes that information.

Also, incremental backups on October 5, 6, and 7 back up everything that has changed since the level 7 backup on October 4.




Figure 22 Backups for October 1 through October 7

Example 13 Backup levels (part 2)

[Figure 23 on page 269](#) continues the example illustrated in [Figure 22 on page 269](#) by showing a level 5 backup on October 8, which backs up everything that changed since the full backup on October 1. To fully recover from a disk crash on October 8, you only need the data from October 1 and the new level 5 backup. You no longer need the data from the level 7 backup on October 4 (or the subsequent incremental backups performed on October 5, 6, and 7) because the new level 5 backup includes that data.




Figure 23 Backups for October 1 through October 14

Also, a level 7 backup on October 11 backs up all of the data that changed since the last lower-numbered backup (in this case, the level 5 backup on October 8). To recover from a disk crash on October 11, you need three volumes: the full volume from October 1, the level 5 volume from October 8, and the new level 7 volume.

Example 14 Backup levels (part 3)

Figure 24 on page 270 continues the example by showing a level 5 backup on October 15, which backs up all of the data that changed since the last lower-numbered backup. Because no backup lower than level 5 has been performed since the full backup on October 1, this level 5 backs up all of the data that changed since the full backup. Therefore, to recover from a disk crash on October 15, you only need the data from the full backup on October 1 and the new level 5 backup.




Figure 24 Backups for October 1 through October 15

The NetWorker server and backup levels

A backup schedule defines what level backup to perform on a given day. You can apply one or more backup levels to customize a backup schedule. To use backup levels in a customized schedule, consider these issues to help you decide what best suits your environment:

- ◆ Full backups generally take more time to complete than incremental backups. The exception to this is the full backup of deduplication clients. With deduplication, only the initial full backup takes longer. Thereafter, every full backup captures only the exact bits that have changed. This allows the subsequent full backups to be significantly smaller.
- ◆ 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 level backups.
- ◆ Level 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 do 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. [“Synthetic full backups” on page 76](#) provides more information.

[Table 35 on page 271](#) lists advantages and disadvantages of each backup level.

Table 35 Advantages and disadvantages of backup levels

Backup level	Advantages	Disadvantages
Full	<ul style="list-style-type: none"> • Faster recovery 	<ul style="list-style-type: none"> • Slow backup • High server load • High load on client and network • Uses more volume space
Level	<ul style="list-style-type: none"> • Faster than performing a full backup • Low load on server • Uses least volume space 	<ul style="list-style-type: none"> • Slow recovery • Data can spread across multiple volumes
Incremental	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> •
Synthetic Full	<ul style="list-style-type: none"> • Faster than performing a full backup • Faster recovery • Low load on server, client and network • Requires less volumes for recover 	<ul style="list-style-type: none"> • High load on storage node • Requires at least two volume drives • Uses most volume space

Online indexes and backup levels

The NetWorker server backs up the online indexes (client file index and media database) differently than regular files and other data.

Knowing how the online indexes are backed up is particularly important during disaster recoveries.

The online indexes are backed up in the following way:

- ◆ The client file index for a NetWorker client is backed up every time the client is backed up. When a client's backup level is incremental, the backup of its client file index is at level 9. For a synthetic full backup, the backups of the indexes are level 9. The NetWorker server does not perform a level 1 backup for this data.

The files associated with the client file index of a NetWorker client reside on the NetWorker server. So, when a client is being backed up, its client file index is also backed up on the NetWorker server.

- ◆ The client file index for the NetWorker server client is backed up every time the NetWorker server is backed up. When the server's backup level is incremental, the backup of its client file index is at level 9. For a synthetic full backup, the backups of the index is level 9. The NetWorker server does not perform a level 1 backup for this data.

For example:

- If the NetWorker server is backed up at the level full, the backup levels of the NetWorker server’s client file index, the media database, and the resource database are also full.
- If the NetWorker server’s backup is a level 5, the backups of the server’s client file index is also a level 5.
- ◆ The media database and the resource database is backed up whenever the NetWorker server is backed up, or after every scheduled backup if the server is not in an active group.

[Table 36 on page 272](#) compares the level at which the NetWorker server backs up regular files and the online indexes.

Table 36 Regular file and index backup levels

Regular files	Online indexes
Full	Full
Level 2 – 9	Level 2 – 9
Incremental	Level 9
Manual (using the User program)	Not saved
Synthetic full	Level 9

The Windows Archive attribute

The Windows file Archive attribute is used by the NetWorker software to help determine if the file should be backed up. The NetWorker software backs up a file if the Archive attribute is enabled.

- ◆ After backing up the file, the NetWorker software turns off the Archive attribute.
- ◆ After restoring the file, the NetWorker software turns on the Archive attribute.

To disable NetWorker’s use of the Windows file Archive attribute:

1. Set the environment variable **NSR_AVOID_ARCHIVE** to a value of **Yes**.
 To set this as a system environment variable, use the Control Panel's System applet. If this variable is used, NetWorker determines a file's need to be backed up based on the traditional save criteria of modification time.
2. Log off, reboot, or restart the **NetWorker Remote Exec Service** to make Windows aware of the system environment variable change.

Backup levels for Windows SYSTEM and VSS SYSTEM save sets

This section discusses the required backup levels for the five SYSTEM save sets and six VSS SYSTEM save sets used to back up Windows 2003 Servers system files. More information about these save sets is available in [Appendix A, “SYSTEM, VSS SYSTEM, and WINDOWS ROLES AND FEATURES Save Sets.”](#)

A backup or restore of SYSTEM or VSS SYSTEM save set automatically includes all eligible components of that save set.

Deduplication does not support the backup of SYSTEM save sets.

SYSTEM save sets

These are the required backup levels for the SYSTEM save sets.

SYSTEM STATE

For Windows XP and Windows Server 2003 with no VSS license or VSS disabled, the system state is a single logical entity. To maintain the integrity of the system state and ensure that a recovery will accurately restore the state of the machine to its condition at the time of the backup, Windows requires that the system state always be backed up in this manner:

- ◆ As a single, indivisible unit. Therefore, components of the NetWorker SYSTEM STATE save set (such as COM+ and the registry) cannot be backed up or restored individually as separate entities.
- ◆ At level full. Therefore, individual components of the SYSTEM STATE save set cannot be backed up at any level other than full. Requesting an incremental backup of the SYSTEM STATE save set, for example, always results in a full backup of the save set.

SYSTEM FILES

As with the other SYSTEM save sets, when a backup of the SYSTEM FILES save set occurs, it is always at level full. However, a requested backup of the SYSTEM FILES save set is performed only if one or more of the system-protected files have changed since the specified date and time. If no system-protected files have changed, none will be backed up and no corresponding save set entry is made in the server's media index. However, on an incremental or level 1-9 backup of the SYSTEM FILES save set, *all* system-protected files are backed up if *any* system-protected files have changed since the specified time.

SYSTEM DB

Each component of the SYSTEM DB save set (such as Disk Quota or Removable Storage) is always backed up at level full, even if another backup level is requested. SYSTEM DB save set components cannot be backed up or restored individually.

SHAREPOINT

The SHAREPOINT save set is always backed up at level full, even if another backup level is requested. There are no individual components to back up or restore.

Automated System Recovery

The Automated System Recovery (ASR) save set is always backed up at level full, even if another backup level is requested. There are no individual components to back up or restore.

VSS SYSTEM Save Sets

These are the required backup levels for the VSS SYSTEM save sets.

VSS SYSTEM BOOT

For Windows Server 2003 with VSS licensed and enabled, the system state is a single logical entity. To maintain the integrity of the system state and ensure that a recovery will accurately restore the state of the machine to its condition at the time of the backup, Windows requires that the system state always be backed up in this manner:

- ◆ As a single, indivisible unit.

Components of the NetWorker VSS SYSTEM BOOT save set (such as COM+ and the registry) cannot be backed up or restored individually as separate entities.

- ◆ At level full.

Individual components of the VSS SYSTEM BOOT save set cannot be backed up at any level other than full. Requesting an incremental backup of the VSS SYSTEM BOOT save set, for example, always results in a full backup of the save set.

VSS SYSTEM FILESET

As with the other VSS SYSTEM save sets, when a backup of the VSS SYSTEM FILESET save set occurs, it is always at level full. However, a requested backup of this save set is performed only if one or more of the system-protected files have changed since the specified date and time. If no system-protected files have changed, none will be backed up and no corresponding save set entry is made in the server's media index. However, on an incremental or level 1-9 backup of the VSS SYSTEM FILESET save set, if any system-protected files have changed since the specified time, all system-protected files are backed up. VSS SYSTEM FILESET components cannot be backed up or restored individually.

VSS SYSTEM SERVICES

Each component of the VSS SYSTEM SERVICES save set is always backed up at level full, even if another backup level is requested. VSS SYSTEM SERVICES components cannot be backed up individually, but they can be restored individually.

VSS USER DATA

Each component of the VSS USER DATA save set is always backed up at level full, even if another backup level is requested. VSS USER DATA save set components cannot be backed up individually, but they can be restored individually.

VSS OTHER

Each component of the VSS OTHER save set is always backed up at level full, even if another backup level is requested. VSS OTHER save set components cannot be backed up individually, but they can be restored individually.

VSS ASR DISK

The VSS ASR DISK save set is no longer supported for backup, but is supported to recover existing save sets.

CHAPTER 8

Browse and Retention Policies

This chapter covers these topics:

- ◆ [About browse and retention policies](#) 276
- ◆ [Managing the data lifecycle](#)..... 281
- ◆ [Browse and retention policies for manual backups](#)..... 286
- ◆ [Modifying the browse and retention policy on a save set](#)..... 286

About browse and retention policies

The *browse policy* determines how long files are maintained in the client's file index on the NetWorker server. During the period of the browse policy, users can browse backed-up data from the NetWorker client computer, and select individual files or entire file systems for recovery. After the browse policy for a file is exceeded, the entry for that file is deleted.

The *retention policy* determines how long save set entries are maintained in the NetWorker server's media database. During the period of the retention policy, an entry for a save set cannot be accidentally overwritten.

For at least the period of the retention policy, you can recover a client's backed-up data by save set selection:

- ◆ No save set is considered recyclable until, at a minimum, it has exceeded its retention policy.
- ◆ No storage volume can be relabeled and overwritten until, at a minimum, all save sets on the storage volume (including save sets that depend on them) have exceeded their retention policies.

Entries for a save set can remain in the media database forever, long after the retention policy has expired. Entries are removed from the media database when these occur:

- ◆ Storage volume is relabeled.
- ◆ Entries are manually deleted.

The NetWorker server maintains one file index for each client computer (regardless of the number of Client resources configured for it), and one media database that tracks data from all clients and all save sets.

Browse policies

You can recover a file that has an entry in the client file index by using the NetWorker client computer. Users can browse and mark files, and initiate data recovery. The NetWorker server does not remove the entry for a file until all the save sets that are dependent on the file have also exceeded their browse policies. In general, the entries for a full backup that are older than the browse policy are not removed until one backup cycle has passed. This extra time ensures that you can reconstruct a file to any point in time included in the browse policy period.

The example in this section demonstrates how a browse policy affects data availability in the client file index. [“Schedules” on page 260](#) provides information about schedules and [“Backup levels” on page 267](#) provides information about backup levels

Example 15 One-week browse policy

In [Figure 25 on page 277](#), both the backup cycle and the browse policy are set at one week. A backup cycle is the length of time between full backups. Entries for the first full backup on October 1 remain in the client file index until all the dependent incremental

and level 5 backups exceed the one-week browse policy. The full backup performed on October 1 is not removed until October 15, when the incremental and level 5 backups that depend on the full backup expire.

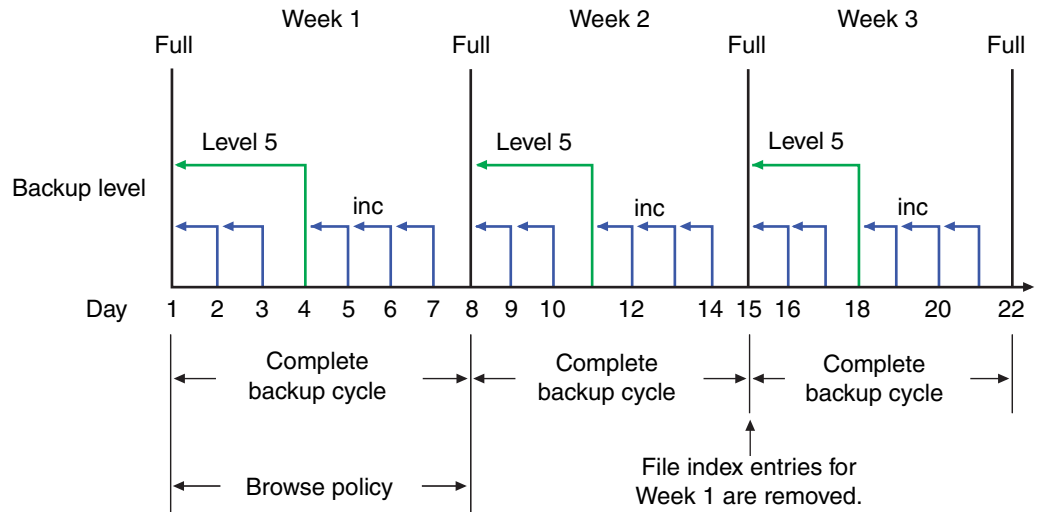


Figure 25 One-week browse policy

To further illustrate, suppose that on October 12, you recover information that is backed up on October 5. The backup performed on October 5 is an incremental backup that is dependent on the October 4 backup, which is a level 5 backup. The October 4 (level 5) backup, in turn, is dependent on the full backup performed on October 1.

The entry for the full backup performed on October 1 must remain in the client file index for a period of time equal to the sum of the following:

- ◆ The browse policy (one week)
- ◆ One complete backup cycle (one additional week)

That is, it must remain in the client file index until the level 5 backup on October 4 and all incremental backups dependent on the full backup pass their browse policy. In [Figure 25 on page 277](#), entries from the Week 1 backup cycle are removed from the client file index on October 15.

Example 16 Two-week browse policy

In [Figure 26 on page 278](#), the browse policy is two weeks, which is twice as long as the backup cycle (one week). In this example, on October 18 a user can still find browsable entries in the client file index from backups created on October 4. The backup performed on October 5 is an incremental backup dependent on the October 4 backup, which is a level 5 backup. The October 4 (level 5) backup, in turn, is dependent on the full backup performed on October 1.

The full backup performed on October 1, and the incremental and level backups that depend on it, must remain in the client file index for a period of time equal to the combination of the following:

- ◆ The browse policy (two weeks)
- ◆ One complete backup cycle (one additional week)

In this example, entries for the Week 1 backup cycle are not removed from the client index until October 22.

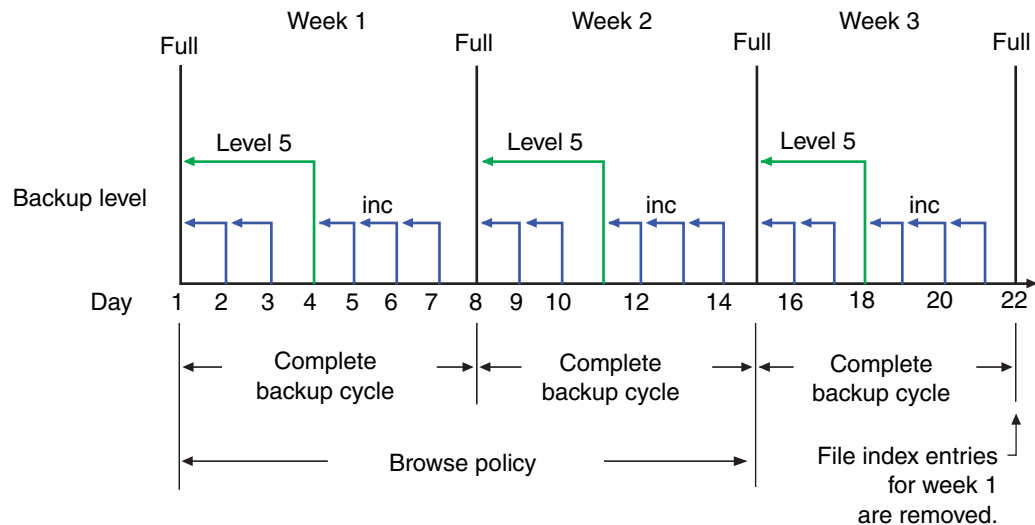


Figure 26 Two-week browse policy

Retention policies

The NetWorker media retention policy specifies a period during which backed-up data is protected from an accidental overwrite. After the retention period is exceeded, the save set is eligible to change its status from recoverable to recyclable. The term recyclable means “eligible for recycling.” The save set’s status, however, does not change to recyclable until it, and all the save sets that depend on it, have passed their retention policy. The NetWorker server keeps track of save set dependencies regardless of whether the dependent save sets are stored on the same or different volumes. The expiration of a save set’s retention policy does not remove the save set’s entries from the media database.

A storage volume becomes recyclable when:

- ◆ The retention policy for every save set on a volume expires.
- ◆ The status for every save set on a volume changes from recoverable to recyclable.

Since a volume can contain save sets from multiple backup sessions, all with different retention policies, the mode of a volume might not change to recyclable for a long time. All the data on the volume remains available for recovery by using either save set **recover** or the **scanner** program. All the entries for recyclable save sets remain in the media database.

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, the volume cannot be made recyclable or relabeled. Furthermore, when deduplication save sets pass their retention time, the NetWorker server will begin the process of deleting the deduplicated data from the deduplication node. Therefore, deduplication data may not be recoverable using the scanner program once the deduplication save set has passed its retention time.

The change in status to recyclable means that the volume can be overwritten if conditions are right. The volume can be relabeled under the following conditions:

- ◆ The volume is placed in an autochanger or mounted in a stand-alone device.
- ◆ The Auto Media Management attribute in the Device resource is enabled.

The existing data is nonrecoverable after the volume is relabeled. [“Auto Media Management” on page 218](#) provides information about the Auto Media Management attribute.

Save set entries are also removed from the media database when they are manually deleted. However, the data on that volume is still available for recovery by using the **scanner** program. The **scanner** program retrieves the information 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’s 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.

[“Recovering a recyclable or recoverable save set entry in the online indexes” on page 399](#) provides more information about re-creating entries in the client file index or the media database.

Note: If only one full browsable saveset backup exists, then its browse policy is equal to its retention policy.

Example 17 Three-week retention policy

Figure 27 on page 280 illustrates how a retention policy works. In this example, the backup cycle is set at one week and the retention policy is set at three weeks.




Figure 27 One-week backup cycle and three-week retention policy

The save set entries for Week 1 have passed their browse policy and retention policy, but they remain available for recovery by using the **scanner** program or via a save set recovery until you relabel the volume. When all the save set entries on a volume change status to recyclable, the volume mode changes from full or appendable to recyclable, and the volume is ready to be relabeled for reuse.

NOTICE

Once a volume is relabeled, data on the volume cannot be recovered.

For more information on these topics, see these sections:

- ◆ [“Viewing volume status information” on page 221](#)
- ◆ [“Schedules” on page 260](#)
- ◆ [“Backup levels” on page 267](#)

Retention policies for client file index save sets

The client file indexes that reside on the NetWorker server are backed up as are any other files. However, the retention policy for these files is calculated differently than for other files. The retention policy for a client file index is based on the retention policy that is specified for the NetWorker client to which the client file index corresponds. For example, if NetWorker client *jupiter* has a retention policy of seven years, then the client file index that corresponds to *jupiter* will also have a retention policy of seven years regardless of any retention policy that may be set up for the NetWorker server. This ensures that if a NetWorker client is recovered, the corresponding client file index can also be recovered.

Retention policies and media pools

A retention policy can also be specified for a media pool. If the retention policy is specified for a media pool as well as the client, the NetWorker software will be the longer of the two policies.

Assigning a retention policy to a clone pool has special implications. When a retention policy is specified in a clone pool, all save sets that are directed to the clone pool take on the retention policy of the clone pool regardless of the retention policy of the save set client. [“Specifying browse and retention policies for clone data” on page 349](#) provides more information.

When browse and retention policies are specified with a command line program, such as **save -w -y**, the browse and retention policies are taken from that program. [“Browse and retention policies for manual backups” on page 286](#) provides more information.

Managing the data lifecycle

Browse and retention policies control the growth of the client file index and the media database, and how long data remains available for recovery.

[Figure 28 on page 282](#) traces the data lifecycle through the client file index and the media database. In the example, the entries for the September 1 through September 7 backup cycle remain in the client index for one month (the browse policy), plus the length of a full backup cycle (one week), to ensure that all dependent entries pass their browse policies. In this case, the file index entries for the September 1 through September 7 backup cycle are removed on October 13. Since the entries exist in the client file index, you can browse and recover the data through the NetWorker client computer. As long as the save set’s file entries remain in the client file index, the status of the source save set is browsable. After the save set status changes from browsable to recoverable, you must know the full path to the file in order to recover it directly.

The status for each save set backed up during the September 1 through September 7 cycle remains recoverable until their retention policies expire and all the dependent save sets exceed their retention policies. In this case, the entries from the September 1 through September 7 backup cycle change from recoverable to recyclable on December 8. When all of the save set entries on a volume change status to recyclable, the mode of the volume changes to recyclable and the volume can be overwritten.

While a save set is either recoverable or recyclable, you can recover any save set by using either the save set recovery procedure or the **scanner** program. Alternatively, you can use the **scanner** program to re-create a save set's entries in the client file index, which enables file recovery directly from the NetWorker client computer. "Recovering expired save sets" on page 398 provide more information.




Figure 28 Data lifecycle in the client index and the media database

On October 13, all data entries from September 1 to September 7 are removed from the client file index. On December 8, the save set entries from September 1 to September 7 in the media database change status from recoverable to recyclable. After all save sets on a volume change status from recoverable to recyclable, the volume mode changes to recyclable. After the volume is relabeled, all existing data on the volume is unavailable for recovery.

NOTICE

When you relabel a volume for reuse within the same *pool*, the volume identification (the volume name as it appears on the volume label) remains unchanged. Even though the volume has the same label, information required by the NetWorker server to locate and restore data on the volume is destroyed. All existing data is inaccessible and will be overwritten.

Assigning multiple policies to a single client

Identical versions of a client and save set combination can have a different set of browse and retention policies assigned for each different backup group to which it belongs. If you create an identical Client resource with the same name and save set values, but assign it to a different

backup group, you can designate a different set of browse and retention policies. The NetWorker server employs the Browse Policy and Retention Policy attributes that correspond to the unique combination of the Client resource's Name, Save Set, and Group attributes.

Example 18 Assigning different policies for an identical client

You have a Client resource for the host saturn. The Client resource has a save set value of All and is assigned to backup group general. The browse policy is weekly and the retention policy is monthly. You create a copy of the Client resource for saturn, but assign it to the backup group special. This version of saturn has a browse policy of Weekly and a retention policy of yearly.

- ◆ If the group special is backed up, then the weekly and yearly policies are applied.
- ◆ If the group general is backed up, then the policies weekly and monthly are used.

Preconfigured time policies

NetWorker software contains these preconfigured browse policies:

- ◆ Day
- ◆ Week
- ◆ Month
- ◆ Quarter
- ◆ Year
- ◆ Decade

You can use these preconfigured policies, or customize policies to best suit data storage needs. Create the customized policy before you configure the client. Otherwise, the policy name does not appear in the Client dialog box as an option.

[“Task 3: Set up policies for quick access and long term storage” on page 62](#) provides information about creating a policy.

Editing a time policy

You *cannot* change the name of a time policy. To rename a time policy, delete the current time policy and create a new one.

To edit a time policy:

1. In the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Time Policies**.
3. In the right pane, select the time policy to edit.
4. From the **File** menu, select **Properties**.
5. Make any necessary changes in the **Properties** dialog box and click **OK**.

Delete a time policy

Note: Preconfigured time policies cannot be deleted. For more information, see [“Preconfigured time policies” on page 283](#).

To delete a time policy:

1. In the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Time Policies**.
3. In the right pane, select the time policy to delete.
4. From the **File** menu, select **Delete**.
5. When prompted, click **Yes** to confirm the deletion.

Snapshot policies

A snapshot policy is required to perform backups with the NetWorker Snapshot Management (NSM) feature. This policy determines how many snapshots are created, retained, and backed up to permanent storage. You can specify a preconfigured policy or create a custom snapshot policy.

The snapshot policy works in conjunction with the Interval attribute of the Group resource. The value for the Interval attribute must be set low enough that the specified number of snapshots can be created in the 24-hour period. For example, to create four snapshots, the Interval value must be set to six hours or less.

For more information on the NSM feature and creating a snapshot policy, refer to the *EMC NetWorker Snapshot Management Integration Guide*. The *EMC NetWorker Module for Microsoft Applications Administration Guide* provides information about creating a snapshot policy for the VSS Client.

Working with snapshot policies

This section provides information about preconfigured snapshot policies, as well as instructions for creating, editing, and deleting snapshot policies.

Preconfigured snapshot policies

If a new customized snapshot policy is *not* manually created, the NetWorker software provides two preconfigured policies that can be used with the snapshot management feature:

- ◆ Rollover-only

With the Rollover-only snapshot policy, a single snapshot is taken per day. The data is then backed up to conventional storage media and the snapshot is deleted.

- ◆ Daily

With the Daily snapshot policy, eight snapshots are taken in a single day. The data in the first snapshot is backed up to tape. Each snapshot has an expiration policy of 24 hours.

Note: Neither preconfigured snapshot policy may be deleted.

Creating a snapshot policy

To create a snapshot policy:

1. In the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Snapshot Policies**.
3. From the **File** menu, select **New**.
4. In the **Create Snapshot Policy** dialog box, type a name for the snapshot policy in the **Name** attribute and complete other attributes as appropriate.

Note: For information about how to complete the attributes for a snapshot policy, refer to the *EMC NetWorker Snapshot Management Integration Guide*.

5. Click **OK**.

Edit a snapshot policy

To edit a policy:

1. In the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Snapshot Policies**.
3. In the right pane, select the snapshot policy to edit.
4. From the **File** menu, select **Properties**.
5. Make any necessary changes in the **Properties** dialog box and click **OK**.

Copying a snapshot policy

To copy a snapshot policy resource:

1. In the **Administration** window, click **Configuration**.
2. In the left pane, select **Snapshot Policies**.
3. In the right pane, select the **Snapshot Policy** resource to copy.
4. From the **Edit** menu, select **Copy**. The **Create Snapshot Policy** dialog box appears.
5. Type the name for the new **Snapshot Policy** resource in the **Name** attribute, edit any other attributes as appropriate, and click **OK**.

Delete a snapshot policy

To delete a snapshot policy:

1. In the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Snapshot Policies**.
3. In the right pane, select the snapshot policy to delete.
4. From the **File** menu, select **Delete**.
5. When prompted, click **Yes** to confirm the deletion.

Browse and retention policies for manual backups

If a browse or retention policy is specified with a manual backup from the command prompt, the browse or retention policy takes effect for all of the save sets included in the manual backup. Specify browse and retention policies with a manual backup from the command prompt by using the **save -w -y** command. Both the browse and the retention policies must be entered in time and date formats accepted by the **nsr_getdate** program. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information about **save** and **nsr_getdate**.

If a browse or retention policy is *not* specified for a manual backup, the policies are determined as follows:

◆ Browse policy

The save sets included in a manual backup adopt the browse policy of the Client resource. If there are multiple Client resources for the NetWorker host, the Client resource with the longest browse time is adopted. For example, if there are three Client resources for the NetWorker client mars, each with one of these browse periods:

- One week
- One month
- One year

A manual backup of mars adopts a browse policy of one year.

◆ Retention Policy

The save sets included in a manual backup adopt the retention policy of the Client resource according to the same rules that were described previously for browse policies. However, if a retention policy is set up for the media pool to which the backup is directed, the retention policy will be the longer of either:

- The Client resource retention policy
- The media pool retention policy

There are special considerations for retention policies and clone data. [“Specifying browse and retention policies for clone data” on page 349](#) provides more information.

Modifying the browse and retention policy on a save set

Use the **nsrmm** program to modify the browse and retention policy of a save set after the backup has occurred. Use **nsrmm** with these options:

- ◆ **-e *retention_time*** – updates retention time
- ◆ **-w *browse_time*** – updates browse time

Use the **-e** and **-w** options with the **nsrmm** option **-S** (to specify a save set ID).

Note: The retention time must be later than the browse time, and the browse time must be later than the *insertion time*. The insertion time is the time that the save set record was most recently introduced into the save set database.

When the **-e** and **-w** options are used with **nsrmm**, these must be true:

- ◆ Retention-time is greater than the browse-time.
- ◆ Browse-time is greater than the insertion-time.

Both the browse and the retention policies must be entered in time and date formats accepted by the **nsr_getdate** program. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information about **nsrmm** and **nsr_getdate**.

Example 19 Changing browse and retention policies with **nsrmm**

The examples in this section use **nsrmm** to change browse and retention policies:

- ◆ Change the retention time to midnight, January 1, 2016. Change the browse time to midnight, January 1, 2012:

```
nsrmm -S 3315861249 -e "01/01/09 23:59:59" -w "01/01/08 23:59:59"
```

- ◆ Change the browse time to six months from the current date and time:

```
nsrmm -S 5315861249 -w "6 months"
```

- ◆ Change the retention time to two years from the current date and time:

```
nsrmm -S 3315861249 -e "2 years"
```

Reports on browse and retention policies for save sets

The **mminfo** command can be used with the **-p** option to display a report on the browse and retention times for save sets. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information about **mminfo**.

CHAPTER 9

Directives

This chapter covers these topics:

◆ Directives overview	290
◆ Types of local and global directives	290
◆ Creating a global directive resource.....	290
◆ Editing a global directive resource.....	291
◆ Deleting a global directive resource.....	291
◆ Copying a global directive resource	291
◆ Example directives	292
◆ Local directives within the NetWorker User program	293
◆ Preconfigured global directive resources	294
◆ Format of directive statements	297
◆ Application Specific Modules (ASMs)	299

Directives overview

Directives are resources that contains special instructions that control how the NetWorker server processes files and directories during backup and recovery. NetWorker administrators can create directives to customize the NetWorker software to your specific needs, maximize the efficiency of backups, and apply special handling to individual files or directories.

NOTICE

Do not leave blank lines in directive scripts.

Types of local and global directives

There are three types of directives.

NOTICE

If you are using the Windows BMR feature, employ user defined directives with caution. Using such directives in directories where system state files reside will lead to an incomplete BMR backup image and potentially render your BMR backup image unusable. If you create user defined directives, test your BMR backup image to ensure that you can recover your Windows system state correctly. [“Perform a NetWorker Bare Metal Recovery wizard test before recovery” on page 754](#) provides more information about testing your BMR backup image.

Global directives

Administrators can create global directives by using the NetWorker Administration window. These 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.

NetWorker User local directives (Windows only)

On clients that run Microsoft Windows, 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, and are applied throughout the client's file systems during scheduled backups (or **save** operations that do not include the **-i** option).

Local directive files

Users can create local directive files named nsr.dir (Windows) or .nsr (UNIX) anywhere on a client file system that they have permission to create files. These directives apply only to the immediate data within the path where the directive file is located.

Creating a global directive resource

[“Format of directive statements” on page 297](#) provides instructions on the syntax to use when creating directives.

To create a global directive:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Directives**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a name for the new directive.
5. In the **Comment** attribute, type a description of the directive.
6. In the **Directive** attribute, type one or more directives.
7. Click **OK**.

The directive can now be applied to a NetWorker Client resource. [Example 20, “Applying a global directive”](#) provides more information.

Editing a global directive resource

A directive name cannot be changed, the directive must be deleted and a new one created with a new name.

To edit a global directive:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Directives**.
3. In the right pane, select the directive.
4. From the **File** menu, select **Properties**.
5. In the **Directive** attribute, modify the directive as necessary and then click **OK**.

Deleting a global directive resource

You cannot delete preconfigured directives or any directives currently applied to a Client resource.

To delete a global directive:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Directives**.
3. In the right pane, select the directive to delete.
4. From the **File** menu, select **Delete**.
5. Click **Yes** to confirm the deletion.

Copying a global directive resource

To copy a global directive:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Directives**.

3. In the right pane, select the directive to copy.
4. From the **Edit** menu, select **Copy**.
A copy of the directive is created.
5. In the **Name** attribute specify a name for the directive.
6. In **Directive** attribute, modify the directive as necessary and then click **OK**.

Example directives

This section contains some basic examples of global and local directives, and describes how to apply them to NetWorker clients.

Example 20 Applying a global directive

This example shows how to use a global directive to skip all *.tmp files in a particular directory for a particular NetWorker client.

1. Create a global directive by using the appropriate format:
 - On Windows, skip all *.tmp files in the C:\mydir directory:


```
<<"C:\MYDIR">>
skip: *.tmp
```
 - On UNIX, skip all *.tmp files in the /mydir directory:


```
<</mydir>>
skip: *.tmp
```
2. Apply the directive to the appropriate NetWorker Client resource:
 - a. From the **Administration** window, click **Configuration**.
 - b. In the left pane, click **Clients**.
 - c. In the right pane, select a client.
 - d. From the **File** menu, select **Properties**.
 - e. From the **Directives** attribute list, select a directive and then click **OK**.

When a scheduled backup is performed on the NetWorker client, all files that match the *.tmp pattern in the specified directories will be skipped.

Example 21 Applying a NetWorker User program local directive (Windows only)

This example shows how to use a Windows local directive to skip all *.tmp files in the C:\mydir directory on a particular NetWorker client.

Using the NetWorker User program, create a local directive. The directive is saved in the networkr.cfg file in this format:

```
<<"C:\mydir">>
skip: *.tmp
```

When a scheduled backup is performed on the NetWorker client, all files that match the *.tmp pattern in the C:\mydir directory will be skipped.

[“Local directives within the NetWorker User program” on page 293](#) provides information on creating local directives.

Example 22 Applying a local file directive

This example shows how to use a local file directive to skip all *.tmp files in a particular directory for a particular NetWorker client. When a scheduled backup is performed on the NetWorker client, all files that match the *.tmp pattern in the specified directory will be skipped.

On Windows, skip all *.tmp files in the C:\mydir directory:

1. Use a text editor to create a file named nsr.dir and type this directive in the file:

```
skip: *.tmp
```

2. Place the nsr.dir file in the C:\mydir directory on the NetWorker client.

Note: To create directive files on a client that is running Microsoft Windows, an authenticated user must have the appropriate permissions to create files either within the root of a volume, or in a folder within the volume.

On UNIX, skip all *.tmp files in the /mydir directory:

1. Use a text editor to create a file named .nsr and type this directive in the file:

```
skip: *.tmp
```

2. Place the .nsr file in the /mydir directory on the NetWorker client.

Order of precedence of global and local directives

If there is a conflict between directives, global directives are enforced over local directives. And on Windows hosts, NetWorker User program local directives are enforced over local directive files (nsr.dir files).

Local directives within the NetWorker User program

On Windows, users can create local directives with the NetWorker User program. These directives are saved in a file named networkr.cfg.

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.

Set up a NetWorker User program local directive

To set up a User program local directive:

1. Log in to the client computer as a member of either the local **Windows Administrators** or **Backup Operators Windows** 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 in order 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. [“Encrypting backup data” on page 108](#) provides information about setting a password.

5. From the **File** menu, select **Save Backup Directives** to save changes.

Depending on user privileges and OS version, the User program creates networkr.cfg in one of these locations:

- If you are logged on 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 on 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.

Preconfigured global directive resources

The NetWorker software comes with a number of preconfigured global Directive resources. All preconfigured Directive resources can be modified, but they cannot be deleted.

[Table 37 on page 294](#) lists the preconfigured directives and their descriptions.

Table 37 Preconfigured directives (1 of 3)

Directive resource	Description
AES	Encrypts backup data with the aes ASM, which provides 256-bit data encryption. For more information about encrypting backup data, see “Encrypting backup data” on page 108 .
DOS standard	Legacy resource that is used to back up Microsoft Windows 95 and Windows 98 clients. By default, this resource has no directives.
DOS with compression	Legacy resource that is used to back up and compress Microsoft Windows 95 and Windows 98 clients. Applies the compressasm ASM to all files.
Mac OS with compression	Contains the same set of directives as the Mac OS standard directive, along with applying the compressasm ASM to specific directories.

Table 37 Preconfigured directives (2 of 3)

Directive resource	Description
Mac OS standard	<p>Contains a set of directives used to back up standard Mac OS clients. Applies these ASMs:</p> <ul style="list-style-type: none"> • The skip ASM is applied to these files and directories: <ul style="list-style-type: none"> /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 /nsr directory to ensure that local directives in /nsr and subsequent subdirectories are applied. • The logasm ASM is applied to the /nsr/logs and /var directories. • The swpasm ASM is applied to the /private/var/vm
NetWare standard	Is used to back up NetWare clients. By default, this resource has no directives.
NetWare with compression	Is used to back up and compress NetWare clients. Applies the compressasm ASM to all files.
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.

Table 37 Preconfigured directives (3 of 3)

Directive resource	Description
UNIX standard	<p>Contains a set of directives used to back up standard UNIX clients. Applies these ASMs:</p> <ul style="list-style-type: none"> • The skip ASM is applied to the tmp_mnt directory. • The skip ASM is applied to core files on the file system. • The allow save environment keyword is applied to the /nsr directory to ensure that local directives in /nsr and subsequent subdirectories are applied. • The skip ASM is applied to the /tmp directory. • The swapas ASM is applied to the /export/swap directory. If swap files are located in a different directory, modify this directive to use the appropriate directory. • The logasm ASM is applied to the /nsr/logs, /var, /usr/adm, and /usr/spool directories. You can apply this ASM to other directories as well. • The mailasm ASM is applied to the /usr/spool/mail and /usr/mail directories. If email files are located in different directories, modify these directives to use the appropriate locations.
UNIX with compression	<p>Contains the same set of directives as the UNIX standard directive, along with applying the compressasm ASM to all files.</p> <p>Note: This directive is only applied to save sets that contain directories. If the save set is defined by using a filename, this directive will not be applied.</p>
VCB directives	<p>VCB directives are valid for backing up virtual machines using the VCB methodology. This directive is supported in the following scenarios:</p> <ul style="list-style-type: none"> • 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. <p>The vcb directive skips the following files and folders:</p> <ul style="list-style-type: none"> • pagefile.sys • hiberfil.sys (Hibernation file) • WINDOWS\system folder • WINDOWS\System32 folder

Format of directive statements

The format of a directive uses three primary types of statements:

- ◆ “Directory specifications” on page 297
- ◆ “ASM specifications” on page 298
- ◆ “Save environment keywords” on page 298

Any text after a hash (#) character is treated as a comment. More information about directive formats can also be found in the `nsr` and `nsr_directive` pages of the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

NOTICE

Do not leave blank lines in directive scripts.

Directory specifications

Directory specifications indicate the highest-level directory to which these ASMs apply. Be aware of the following considerations when adding directory specifications:

- ◆ You cannot use wildcards in directory specifications.
- ◆ When multiple directory specifications are used, directives that follow a directory specification apply to that directory until the next directory specification.
- ◆ Mount points, including nested mount points, must have their own directory specification.
- ◆ For directives applied to clients on Windows systems, pathnames are not case-sensitive. If there is a colon (:) in the pathname, the entire path must be enclosed in quotation marks.

The format for a directory specification is:

`<<directory>>`

- ◆ On Windows:

```
<<"C:\BIN">>
asm
<<"C:\TEMP">>
asm
```

- ◆ On UNIX:

```
<<"/etc">>
asm
<<"/tmp">>
asm
```

ASM specifications

ASMs specify the action to take on one or more files. The syntax for an ASM specification is:

[+]asm: argument

where:

- ◆ The optional plus sign (+) indicates that the directive applies to both the current directory and all subdirectories.
- ◆ *asm* is the name of the ASM to be executed
- ◆ *argument* is a list of names (files or directories) that are acted upon by the ASM. The argument can include multiple names, separated by spaces, and can also specify wildcards. The argument can contain names that are in the current directory only. Subdirectories cannot be specified in the argument.

If an ASM or argument name includes a space, enclose the name or argument in double quotes. “[Save environment keywords](#)” on page 298 provides a description of available ASMs and examples.

Note: For directives applied to clients on Microsoft Windows systems, filenames are case-sensitive.

Using wildcards in ASM specifications

Wildcards can be used in ASM specifications to replace a single character or string of characters. Standard shell command interpreter filematching patterns are supported.

On UNIX systems, when applying a directive to all files, including hidden files, use * . ?* (insert a space after the first asterisk).

Save environment keywords

Save environment keywords are used to control how the current ASM and subsequent ASMs that apply to the current directory and subdirectories will be applied.

Table 38 on page 299 lists the three save environment keywords.

Table 38 Save environment keywords

Keyword	Instruction	Example
forget	Instructs the NetWorker server to no longer apply inherited directives (those directives that begin with a +).	To skip all *.o files in the directory G:\SRC, except those in the G:\SRC\SYS directory, type: <pre><<G:\SRC>> +skip: *.o <<G:\SRC\SYS>> forget</pre> <p>This uses the skip ASM to instruct the NetWorker server to skip all files 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.</p> <p>The forget keyword works only if the corresponding directories are also explicitly specified in the NetWorker client resource's Save Set attribute. In the previous example, both the G:\SRC and the G:\SRC\SYS directories must be explicitly specified on separate lines in the client's Save Set attribute.</p>
ignore	Instructs the NetWorker server to ignore all directives applied to the subdirectories below the current directory.	To override any local directives set in users' home directories, type: <pre><<HOME>> ignore</pre>
allow	Used in subdirectories that currently have the ignore keyword applied to them, and overrides the ignore.	Building on the preceding example for the ignore keyword, to allow directives in the HOMEDOC directory to be applied, type: <pre><<HOME>> ignore <<HOMEDOC>> allow</pre>

Application Specific Modules (ASMs)

Directives use Application Specific Modules (ASMs) to process files and directories. ASMs are programs that operate within the NetWorker environment to perform various tasks during backup and recovery. For example, the **compressasm** program instructs the NetWorker server to compress data during backup.

ASMs are specified in a directive and are then executed during the backup of client data. Directives can contain one or more ASMs. [Table 39 on page 300](#) describes the NetWorker ASMs.

Table 39 NetWorker ASMs (1 of 2)

ASM name	Description	Example
aes	Encrypts backup data when included in a global directive that is associated with a Client resource. “Encrypting backup data” on page 108 has information about using the AES ASM.	+aes: *
always	Always backs up the specified files, independent of the change time of the file, and ignores the scheduled level. This ASM can be used to ensure that important client files that change constantly are always fully backed up.	always: *.c
compressasm	<p>Compresses files so that they use less network bandwidth and fewer volumes. This ASM does not compress directories. The amount of compression achieved is data-dependent. This ASM can use considerable CPU resources, so its benefits could be limited on low-powered systems.</p> <p>Some storage devices such as cloud devices and deduplication devices have their own encryption capabilities. If such a device is already set up to compress data, then using the compressasm will likely yield no added benefit. In fact, the additional compression may result in slightly more data being written to your device.</p> <p>Three types of compression are supported:</p> <ul style="list-style-type: none"> • Default NetWorker compression • gzip compression • bzip2 compression 	<p>For default NetWorker compression, specify one of the following:</p> <pre>+compressasm: .</pre> <pre>+compressasm -default: .</pre> <p>For gzip compression, specify the -gzip argument with a compression level from 0 to 9. A value of 0 provides the least compression and 9 provides the most compression. If no level is specified, the default value is 6:</p> <pre>+compressasm -gzip 3: .</pre> <p>For bzip2 compression, specify the -bzip2 argument with a compression level from 0 to 250. A value of 0 provides the least compression and 250 provides the most compression. If no level is specified, the default value is 0:</p> <pre>+compressasm -bzip2 250: .</pre> <p>Added compression uses more CPU resources and therefore, could increase backup times.</p> <p>Both gzip and bzip2 compression cannot be used with the aes encryption ASM.</p>
holey	Handles holes or blocks of zeros when backing up files and preserves these holes during recovery. This ASM is normally applied automatically and does not need to be specified.	+holey: *
logasm	Instructs the NetWorker server to not generate errors when the files specified by this ASM are in use. This ASM is useful for files involved in logging, and other similar files that might change during a backup operation.	+logasm: *.log
mailasm	Uses mail-style file locking and maintains the access time of a file, preserving "new mail has arrived" flag on most mail handlers.	+mailasm: *.mbx

Table 39 NetWorker ASMs (2 of 2)

ASM name	Description	Example
mtimeasm	Backs up files by using the modification time, rather than the inode change time, to determine which files should be backed up. The modification time is the last time the file's contents were modified, while the inode change time is the last time the file's mode, owner, or link count was changed.	mtimeasm: *.log
nsrindexasm	Used to recover from NetWorker file index backups performed by NetWorker servers prior to release 6.0. During recovery from these older index backups, nsrindexasm is invoked automatically by nsrck and mmrecov .	Not applicable
nsrmmdbasm	Used to process the media database. Normally, nsrmmdbasm is invoked automatically by savegrp and mmrecov , and should not be used in NetWorker directives.	Not applicable
null	Does not back up the specified files, but does back up the directory containing the files so entries for the files get added to the online indexes. The NetWorker server uses this ASM to back up the online indexes during a scheduled backup.	+null: *.tmp
nullasm	Another name for the null ASM, used for backward compatibility.	See null .
posixcrcasm	Calculates a 32-bit cyclic redundancy check (CRC) for a file during backup. This CRC is stored along with the file and is verified when the file is restored. No verification occurs during the backup itself. With this ASM, it is possible to validate a file at restore time, but it does not provide a way to correct any detected errors.	posixcrcasm: *.*?
rawasm	Specifies the back up of UNIX raw disk partitions. The /dev entries (block and character-special files) and their associated raw disk partition data is backed up. On some systems, /dev entries are symbolic links to device specific names. Unlike other ASMs, this ASM follows symlinks, allowing the shorter /dev name to be configured. “Precautions when using rawasm to back up UNIX raw partitions” on page 302 provides more information.	rawasm: /dev/oracle1
skip	Omits files and directories from the backup, and does not place the directory or filename in the online index. In the example given, all files and directories with the name temp will be omitted from the backup.	+skip: temp
swapasm	Does not back up actual file data, but re-creates a zero-filled file of the correct size on recovery. This ASM is used on systems where the swapping device is a swap file that must be recovered with the correct size, but the contents of the swap file are not important and do not need to be backed up or restored.	swapasm: compression.doc
xlateasm	Translates file data so that data backed up is not immediately recognizable	xlateasm: *.*

Precautions when using rawasm to back up UNIX raw partitions

One can specify the **rawasm** directive to back up raw disk partitions on UNIX. However, if the raw partition contains data managed by an active database management system (DBMS), ensure that the partition is offline and the database manager is shutdown. For greater flexibility when backing up partitions that contain data managed by a DBMS, use a NetWorker Module application.

Similarly, if **rawasm** is used to save a partition containing a UNIX file system, the file system must be unmounted or mounted read-only to obtain a consistent backup.

Note: Do not specify the **rawasm** directive to backup or recover raw partitions on Windows. [“Backing up raw partitions on Windows” on page 128](#) provides more information.

Using rawasm to recover a UNIX raw partition

When recovering data, **rawasm** requires that the file system node for the raw device exist prior to the recovery. This protects against the recovery of a /dev entry and the overwriting of data on a reconfigured disk. You can create the /dev entry, having it refer to a different raw partition, and force an overwrite if needed. If you create the /dev entry as a symbolic link, the data is recovered to the target of the symbolic link.

Recovery of a raw partition must occur on a system configured with the same disk environment and same size partitions as the system that performed the backup:

- ◆ If the new partition is *smaller* than the original partition, the recovery will not complete successfully.
- ◆ If the new partition is *larger* than the original partition the estimated size reported upon recovery is not accurate.

File matching with multiple ASMs in a directive

When a file matches multiple ASMs in a directive, the action taken on the file depends on the order of the ASMs in the directive. For example, if these ASMs are listed in a directive:

```
+always: master.mdf master.ldf
+skip *.mdf *.ldf
```

Then the master.mdf and the master.ldf files will be backed up because the always ASM is processed first. All other files with a .mdf or .ldf extension will not be backed up.

However, if the order of the ASMs is reversed:

```
+skip *.mdf *.ldf
+always: master.mdf master.ldf
```

The master.mdf and the master.ldf files will not be backed up because the **skip** ASM is processed first.

Note: To simplify directives that include multiple potential matches for the same file, consider using save environment keywords. [“Save environment keywords” on page 298](#) provides more information.

CHAPTER 10

Sorting Backup Data

This chapter covers these topics:

- ◆ [Media pools](#)..... 304
- ◆ [Label templates](#) 318

Media pools

Backup data is sorted onto backup media volumes by using media pools and volume labels. A media “pool” is a specific collection of volumes to which the NetWorker server writes data. The server uses media pools to sort and store data. A volume is identified with a unique label based on configurable label templates.

Media pools act as filters that tell the server which backup volumes should receive specific data. The NetWorker server uses media pools in conjunction with label templates (if the Match Bar Code Labels attribute is not used for the library resource) to keep track of what data is on which specific volume. [“Using label templates” on page 318](#) provides more information on label templates.

Note: Media pools do not apply when deduplication is used.

Using media pools

Each media pool configuration includes criteria that must be met in order for the data to be written to associated volumes.

When a scheduled backup occurs, the NetWorker server tries to match the save stream to a media pool configuration. If the save stream matches the criteria of a media pool configuration, it is directed to a labeled volume in the media pool. The 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 requests that such a volume be mounted and waits until an operator or an autochanger mounts an appropriate volume.

Default media pool

If the save stream does not match the criteria for any custom (user-created) media pools, the NetWorker server directs the save stream to the Default media pool. Even if you use customized media pools, ensure that appropriate storage devices and media are available for the Default media pool for cases when the NetWorker server directs save streams there. If the media pools are not properly configured for backup, an error message similar to this may appear in the Monitoring Alerts tab in the Administration window:

```
media waiting (or critical): backup to pool 'Default' waiting for 1
writable backup tape or disk
```

NetWorker media pool types

NetWorker software contains preconfigured media pool types to keep different types of data separate. NetWorker software does not mix these types of data within a media pool:

- ◆ Backup
- ◆ Backup clone
- ◆ Archive
- ◆ Archive clone

- ◆ Migration
- ◆ Migration clone
- ◆ WORM (Write Once-Read Many)
- ◆ DLTWORM

Without any special sorting, all backup data is routed to the Default media pool and all archive data is routed to the Default Archive media pool. Likewise, clone data is routed to the appropriate Default Clone media pool. [“Creating a media pool” on page 311](#) provides information on customizing media pools.

Sorting data with media pools

When you configure the NetWorker server, you can create additional media pools and sort data by media pool type and any combination of the following:

- ◆ Group (backup group)
- ◆ NetWorker client
- ◆ Save sets (file or file systems)
- ◆ Backup levels (full, levels 1 – 9, incremental, manual)

When you select a group, the media pool accepts only data associated with the named group. If a second group name is added, the media pool accepts data associated with either group, but no others. If you enter configuration criteria in both the Group attribute and Save Set attribute, only data that meets both the group criteria *and* the save set criteria is written to volumes from the specified media pool. [Chapter 7, “Backup Groups and Schedules”](#) provides information about groups and backup levels.

Example 23 Using media pool configurations to sort data

[Figure 29 on page 306](#) illustrates how the NetWorker server uses media pool configurations to sort data. The save stream contains data from a full backup that was performed on client and save sets in a group called Accounting. The NetWorker server looks for a media pool configuration that matches the group named Accounting and the level full. When the NetWorker server finds the matching media pool configuration, it writes the data to a volume with a label from the corresponding Accounting Full media pool of volumes mounted on one of the storage devices.




Figure 29 Using media pool configurations to sort data

Directing client file indexes and bootstrap to a separate media pool

You can use regular expression matching to direct the client file indexes and bootstrap to a media pool other than where you send the backup data.

Example 24 Sending bootstrap information and all the client file indexes to the same media pool

To send the NetWorker server’s bootstrap and client file indexes to the same media pool, create a media pool (in the Media Pool resource) with the values listed for the attributes in [Table 40 on page 306](#).

Table 40 Attributes to direct client indexes and bootstrap to a separate media pool

Attribute	Value
Name	Index
Pool Type	Backup
Save Sets	Bootstrap index

When the group’s scheduled backup runs, the client save sets are written to a volume labeled for the appropriate save set media pools, while the NetWorker server’s bootstrap and index save sets are written to a separate volume labeled for the Index media pool.

Directing consolidated backup data to a specific media pool

By default, save sets from a consolidated backup are written to whatever media is mounted for the group most recently backed up.

To direct consolidated save sets to a specific set of media:

1. Configure a **Group** resource for consolidated backups. [“How to create a group” on page 253](#) provides information.
2. Configure a **Media Pool** resource for consolidated backups. [“Creating a media pool” on page 311](#) provides instructions.
3. In the **Create Media Pool** dialog box, select the name of the **Group** resource created in [step 1](#) for the **Groups** attribute.
4. Add each client that will receive consolidated backups to the group created for those backups.

Meeting the criteria for more than one media pool configuration

Depending on the media pool that is created, there may be data that matches the criteria for more than one media pool configuration. For example, if one media pool is configured to accept data from a group called Accounting and another media pool is configured to accept data from all full backups, it is not immediately clear which pool of volumes will be used for a full backup for the Accounting group. The NetWorker server uses this media pool selection criteria:

1. Group (highest precedence)
2. Client
3. Save set
4. Level (lowest precedence)

When data matches the attributes for two media pools, for example, Group and Level, the data is written to the media pool specified in the Group attribute. For example, in the case where the data from the group matched the criteria for two different media pools, the data is routed to the media pool that accepts data from the Accounting group.

[Table 41 on page 307](#) details the hierarchy that the NetWorker server uses to determine media pool selection when a conflict arises. For example, the media pool criteria for Group takes precedence over the media pool criteria for client, save set, and level. Data that meets the criteria for both media pools is written to the media pool associated with the group. If data does not meet the criteria for any customized group, it is written to the Default media pool.

Table 41 NetWorker hierarchy for resolving media pool conflicts (1 of 2)

Precedence	Group	Client	Save Set	Level
Highest	x	x	x	x
	x	x	x	
	x	x		x
	x	x		

Table 41 NetWorker hierarchy for resolving media pool conflicts (2 of 2)

Precedence	Group	Client	Save Set	Level
	x		x	x
	x		x	
	x			x
	x			
		x	x	x
		x	x	
		x		x
		x		
			x	x
			x	
Lowest				x

When no customized media pool criteria is met

When you sort data by using a customized media pool, you might inadvertently omit a client or save set. During a scheduled backup, if data does not meet the criteria for any customized media pool, the NetWorker server automatically sends the data to the Default media pool. By using the Default media pool, the server ensures that all data is backed up to a volume.

When the NetWorker server sends data to the Default media pool, it looks for a labeled volume from the Default media pool mounted on a storage device. If no Default media pool volume is mounted on a storage device, the server requests the appropriate volume and waits until an operator mounts the volume. If the NetWorker server asks for a Default media pool volume in the middle of a scheduled backup, the backup pauses until it has been mounted. If an operator is available to monitor backups, keep a Default media pool volume on hand should this situation arise.

To plan for unattended backups, run a test of the backup after making any configuration changes. This ensures that all data is written to the appropriate media pools of volumes, and avoids requests for a Default media pool volume.

Configuring media pools

This section provides information specific to the configuration of these media pool types:

- ◆ Incremental backups
- ◆ Manual backups
- ◆ Clone data
- ◆ Archive data

Note: You can create and edit media pools while a backup session is running. For each change, a message is written to the daemon log file, located in the <NetWorker_install_dir>\logs directory. [“Viewing log files” on page 803](#) provides information about viewing log files. You cannot delete a media pool that has labeled volumes in the media database.

Incremental backups

When creating a separate media pool for incremental backups, the NetWorker software’s hierarchy of precedence affects the way the data is stored. When a pool has been configured with a level incremental restriction and an incremental server initiated backup is performed:

- ◆ Incremental data will be routed to the media pool.
- ◆ The client file index will not go to the incremental pool. In an incremental backup, the associated index will backup at a level 9 to speed the recovery operation, if needed.

If the client indexes do not meet the criteria for the media pool associated with the incremental backups (that is, level 9 is not allowed), the indexes are matched to another media pool (usually the Default media pool) and an appropriately labeled volume is searched for. To recover the data, a large number of volumes may be required. To speed the recovery, define the level value of the appropriate pool to accept both level 9 and incremental data, rather than only incremental.

By using the preconfigured NonFull media pool, you ensure that the client file indexes belong to the same media pool as their incremental backups. By keeping the indexes in the same media pool as their incremental backups, you reduce the number of volumes required for a recovery.

Manual backups

You can create a customized media pool to receive data from a manual backup by specifying manual in the Level attribute. The NetWorker server, however, sorts data from a manual backup differently than it sorts data from a regularly scheduled backup. Since a manual backup is not performed as part of a scheduled backup group, by default the data is not associated with any group name. When you perform a manual backup in which only data from a single client or save set is saved, the group normally associated with that client or save set is not considered in any sorting. As a consequence, data from a manual backup may be sent to a media pool other than the one in which data from this client or save set is stored during a regularly scheduled backup. If you do not create a customized media pool to receive data from manual backups, the NetWorker server uses the Default media pool and looks for a mounted volume from the Default media pool on which to write data.

Note: Manual backups only back up file system data. Unlike scheduled backups, they do not back up the client file index at the end of the backup. The next server initiated backup of the client will backup the index. Manual backups can still be browsed at recovery time unless an index recover is performed before the index containing the save information has been backed up.

Clone data

To clone data, use a specific media pool to receive the clone data and a minimum of two devices:

- ◆ One to read the source volume
- ◆ One to write the clone

If data to be cloned is not associated with a customized Clone media pool, the Default Clone media pool is used. [Chapter 12, “Cloning”](#) provides more information.

Archive data

To archive data, use a specific media pool to receive the archived data. You can then store these volumes offsite. If data to be archived is not associated with a customized Archive media pool, the preconfigured Archive media pool is used. [Chapter 11, “Archiving”](#) provides more information about archive feature.

Using storage devices and media pool configuration to sort data

Data can be sorted by configuring media pools, in conjunction with storage devices, to either use specific media to receive data, or designate a specific storage device to receive data from a designated media pool.

Using different media

You can write data across several volumes of different media types (for example, magnetic disk and tapes) as long as the volumes mounted on the storage devices have the appropriate label associated with the media pool.

Using a specific device for backup data

You can associate a media pool with a specific storage device. For example, full backups may be written to optical disk for offsite storage. There are two ways to ensure that data goes to one specific storage device:

- ◆ Always keep a labeled volume associated with the appropriate media pool mounted on the specific storage device.
- ◆ Associate a specific media pool with the storage device in the Volume Pool attribute in the Device resource. Then, for the Media Pool resource, select that device for the Devices attribute. All data will be written only to that device.

Note: The only time you can assign a device to a media pool is when you label it. If you later want to assign the device to a different volume pool, you must relabel it.

Creating a media pool

Resource dialog box attributes vary depending on the server. Use the steps presented here as a general guideline. For additional information about each attribute, click **Field Help**.

Before creating a media pool, do 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. [“Creating a label template” on page 322](#) provides more information.
- ◆ Determine a preconfigured label template to use for the media pool.

If you do not select a label template when creating a media pool, the NetWorker server notifies you that it will create a label template for the media pool.

To have the NetWorker server create the label template, click **OK**. [“Using label templates” on page 318](#) provides more information on label templates.

To create a media pool:

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. Select the **Enabled** attribute.
7. For the **Pool Type** attribute, select the appropriate media pool type.
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. Select the **Selection Criteria** tab.
11. To further restrict which data can back up to this media pool, complete any of these attributes: **Client**, **Save Sets**, or **Level**. For guidelines about typing save set paths, see [“Expression matching of save sets to media pools” on page 312](#).
12. Select the **Configuration** tab.

13. In the **Auto Media Verify** attribute, select whether automated media verification will be performed while data is written to a volume associated with this media pool. For more information, see [“Auto media verification” on page 312](#).
14. Complete the other attributes as necessary, and 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.

Expression matching of save sets to media pools

If you enter a save set path, the Save Set attributes in the Media Pool resource are matched by using regular expression matching. This affects how the save set filenames are written, and how the path is written in this field for a Microsoft Windows system. Backslashes and periods must be preceded with a backslash:

- ◆ A save set path of C:\SomeDir\MyFiles should be written C:\\SomeDir\\MyFiles.
- ◆ A filename of MyFile.txt should be written MyFile\\.txt.

When using the **NetWorker Administrator** program, each save set in the Save Set attribute must be on a separate line. The following is an example of properly written save set entries:

```
/
/usr
C:\\Program Files\\bin
*\\.doc
```

The **nsr_regexp** and **nsr_pool** entries in the *EMC NetWorker Command Reference Guide* or the UNIX man pages provide information about regular expression matching.

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.

Managing volumes in a media pool

In NetWorker 8.0 and later, new pool resource attributes exist to provide support for scheduling the automatic relabeling of eligible volumes in a pool. In prior releases, volumes would be relabelled at the time of backup or clone and only when the selection criteria was met. [“How the NetWorker server selects a volume” on page 216](#) provides details on volume selection criteria during a backup or restore.

Consider the following:

- ◆ Automatically relabeling a recyclable volume provides the following benefits:
 - Volumes can be relabeled outside of the backup window without the need for a scripted solution.
 - Appendable volumes are available at the time of a backup or clone, resulting 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 nsrmmmd process (for example, during a restore operation)
 - In read-only mode
 - Busy
- ◆ The **daemon.raw** file on the NetWorker server is updated with the following message when volumes are automatically relabeled:

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

To configure automatic volume relabeling for a pool:

1. Modify an existing pool or create a new pool resource. [“Creating a media pool” on page 311](#) provides details on how to create a new pool resource.
2. Under the **Miscellaneous** tab, configure the automatic relabeling attributes as required:
 - **Recycle start:** (Format HH:MM) defines the time to start the automatic relabel process each day. By default this attribute is empty and the automatic relabelling of recyclable volumes is not done.
 - **Recycle interval:** (Format HH:MM) defines the interval between two starts of automatic relabel processes. The default value is 24:00

- **Recycle start now:** invokes the automatic relabel process of recyclable volumes for this pool immediately. The default value is **No**.
- **Max volumes to recycle:** defines the maximum number of recyclable volumes that can be relabeled during each automatic relabel process. The default value is 200.

Supporting 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.

Table 42 on page 314 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* on the EMC Online Support web site.

Table 42 WORM supported devices

Device	Description
HP LTO Ultrium 3 and higher	Unique to HP Ultrium-3 and higher <ul style="list-style-type: none"> • 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.) <ul style="list-style-type: none"> • Inquiry data VPD page 0xc0, byte 2, bit 0 to indicate WORM capable • Read attribute # 0x0408 bit 7 to indicate WORM media present
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: <ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> - 00 indicates no overwriting allowed - 01 indicates some labels can be overwritten • Byte 5, bits 0,1: filemark overwrite restrictions <ul style="list-style-type: none"> - 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.

Identifying 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 [Figure 30 on page 315](#), a (W) is appended to the volume names displayed in the NetWorker Administrator window. If a volume is both read-only and WORM, an (R) is appended to the volume name.

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
NECHVAR.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
sait_via_polarbear.001		1091 MB	full		9/12/06		Default

Figure 30 Identifying WORM tapes in the NetWorker Console

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.

Configuring WORM and DLTWORM support

To determine if a device is WORM or DLTWORM capable:

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).

To configure pools to accept WORM or DLTWORM devices:

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.

Table 43 WORM/DLTWORM attributes (1 of 2)

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 (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: <ol style="list-style-type: none"> 1. You do not want to segregate WORM tapes within WORM pools. 2. 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.

Table 43 WORM/DLTWORM attributes (2 of 2)

Attribute	Description
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	<p>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.</p> <hr/> <p>Note: Refer to the Quantum web site for information on which tapes can be converted to DLTWORM tapes.</p> <p>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.</p> <hr/>

Working with media pools

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

Editing a media pool

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

To edit a media pool:

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 **Properties**.
5. In the **Properties** dialog box, make the necessary changes and click **OK**.

Copying a media pool

To copy a media pool:

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**.

For details about the **Media Pool** attributes, click **Field Help** in the **Properties** dialog box.

Deleting a media pool

Note: You can delete a media pool only if there are no active volumes assigned to the media pool. Preconfigured media pools cannot be deleted.

To delete a media pool:

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 deletion.

Creating an archive media pool

To archive data, the NetWorker server requires an archive media pool to receive the archive data. If data to be archived is not associated with a custom archive media pool, the server automatically uses a preconfigured archive media pool. An appropriately labeled volume must be mounted on a storage device for the archive process to proceed.

To create an archive media pool:

1. Create a new media pool. [“Creating a media pool” on page 311](#) provides instructions.
2. From the **Pool Type** attribute, select **Archive** from the list.
3. Select the **Configuration** tab.
4. Clear the **Store index entries** attribute.

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.

Figure 31 on page 319 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 31 Labeling a volume by using a label template

Using preconfigured label templates

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

- ◆ Default
- ◆ Default clone
- ◆ Archive
- ◆ Archive 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
```

Completing Label Template attributes

There are certain guidelines to keep in mind when completing the attributes for a Label Template resource. [Table 44 on page 320](#) describes how to complete the key attributes for this resource. [“Creating a label template” on page 322](#) provides more information on creating a label template.

Table 44 Key label template attributes

Attribute	Guidelines
Name	<p>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 named "AcctFull" would identify volumes that belong to a media pool called "Accounting Full."</p> <p>Do not use these characters in label template names: / \ * ? [] () \$! ^ ; ' " ' ~ < > & { } : - . _</p>
Fields	<p>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.</p> <p>You can use four types of components:</p> <ul style="list-style-type: none"> • 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) <p>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).</p> <p>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. Table 45 on page 321 illustrates how label templates use components to create a number sequence for volume labels.</p>
Separator	<p>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.</p>
Next	<p>Choose the next sequence number to write on the label that the NetWorker server places on a volume (according to the template).</p> <ul style="list-style-type: none"> • 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. <p>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.</p>

Table 45 on page 321 lists examples of number sequences for volume labels.

Table 45 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 Range of numbers	SalesFull 001-100	SalesFull.001,...SalesFull.100	100
Range of lowercase letters Range of numbers	aa-zz 00-99	aa.00,...aa.99, ab.00,...ab.99, ac.00,...ac.99, : az.00...az.99, ba.00,...ba.99 : zz.00,...zz.99	67,600 (26 ² times 10 ²)

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.

[Table 46 on page 322](#) illustrates several ways you can use components to organize labels.

Table 46 Using label template components

Type of organizational structure	Components	Separator	Resulting volume labels
Sequential	AcctFull '001-100	period	AcctFull.001 (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)

Labeling tips

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 described in [Table 44 on page 320](#).

To create a label template:

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.

To edit a label template:

1. In the **Administration** window, click **Media**.
2. In the expanded left pane, select **Label Templates**.
3. In the right pane, select the template to edit.
4. From the **File** menu, select **Properties**.
5. In the **Properties** dialog box, make any necessary changes and click **OK**.

Copying a label template

To copy a label template:

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.

To delete a label template:

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.

CHAPTER 11

Archiving

This chapter covers these topics:

- ◆ Overview of archiving 326
- ◆ Permissions for archiving 328
- ◆ About archive pools 329
- ◆ Creating custom Archive pools 329
- ◆ Archiving data procedures..... 330
- ◆ Retrieving archived data..... 334
- ◆ Archive request management..... 337

Overview of archiving

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.

Archive save sets are similar to backup save sets. The main difference is that archive save sets have no expiration date. By default, the archive backup level is always set to full. Archive data must be written to separate pools. Browse and retention policies do *not* apply to archive data.

Note: The archive feature must be purchased and licensed separately from other NetWorker software components. The NetWorker *Licensing Guide* provides more information on licensing procedures.

Benefits of using the NetWorker archive feature include:

- ◆ Files that have been archived can be:
 - Deleted from the primary disk storage to make space for newer files.
 - Retained in archive volumes for quick retrieval.
- ◆ Archived data is never subject to automatic recycling, so it *cannot* be overwritten accidentally.
- ◆ Files on archived volumes that use the Indexed Archive pool and the PC Archive pool can be browsed indefinitely.

NetWorker software does not support archiving of SYSTEM or VSS SYSTEM save sets. The NetWorker Client Direct feature does not support archiving.

Archive requirements

Before NetWorker archive feature, ensure that you have:

- ◆ A device, either stand-alone or in an autochanger or silo, connected to a NetWorker server or storage node. If you are cloning archives, you must have at least two devices available.
- ◆ A temporary or permanent enabler code to license the product after any evaluation period is over. The NetWorker *Licensing Guide* provides more information.

How the NetWorker server archives data

The NetWorker software provides three preconfigured pools to receive archived data:

- ◆ Indexed Archive pool
- ◆ PC Archive pool
- ◆ Archive pool

You can also create custom archive pools. During the archive operation, the NetWorker server writes data to storage volumes that belong to an Archive pool. Information about archive data is tracked in the NetWorker server’s media database.

If you use the preconfigured Indexed Archive pool or PC Archive pool, or if you create a custom Archive pool that has the Store Index Entries attribute in the Pool resource set to Yes, information about individual files in the archive save set are 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.

Note: Index entries are not generated when the Store Index Entries attribute in the Pool resource is set to No.

The NetWorker server tracks the volumes used for archives separately from those used for backups. You *cannot* archive files to a backup volume, nor can you back up files to an archive volume. An archive volume must be loaded and mounted in the server device to complete an archive.

Whether you initiate the archive on the client or the server, the archive is created by the client's archive program (**nsrarchive**), which is initiated by the client's **nsrexecd** service. You can schedule archives from the server or client by using the Archive Requests resource in the NetWorker Administrator program.

Figure 32 on page 327 illustrates how the NetWorker software archives data.




Figure 32 Overview of archive operation

Indexed and nonindexed archiving

The NetWorker server supports two styles of archiving:

- ◆ Indexed archiving for Indexed Archive pools or PC Archive pools
- ◆ Nonindexed archiving for Archive pools

Indexed archiving (Indexed Archive pool, PC Archive pool)

One can browse and select indexed archives for individual file recovery.

To use indexed archiving, do one of the following:

- ◆ Use the preconfigured Indexed Archive pool or the PC Archive pool.

- ◆ Enable the Store Index Entries attribute of the Pool resource associated with the archive volume.

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.

Nonindexed archiving (Archive pool)

When data is archived by using nonindexed archiving, entries are not added to the client file index. When this data is retrieved, the entire save set is retrieved and you cannot browse or recover individual files.

Permissions for archiving

This section describes the permissions required to use the archive feature.

Enabling archive services for the client

After the archive service is licensed and the enabler code has been entered into the NetWorker server, all clients of that server are enabled for the NetWorker archive feature by default. To disable the archive feature for a specific client, set the Archive Services attribute in the Client resource to Disabled.

To archive data that resides on the NetWorker server, ensure that the Archive Services attribute is enabled in the Client resource for the server. [“Enabling archive services for a NetWorker client” on page 330](#) provides instructions.

When you enable the Archive Services attribute for a Client resource, enable the Archive Services attribute for all other clients of the same name on that server. For example, if the NetWorker Module for a database application and the NetWorker client software are installed on the same computer and both back up to the same NetWorker server, both Client resources have the same name. Ensure that the Archive Services attribute is enabled for both Client resources.

Enabling or restricting archive access

The Archive Users User Group specifies the users who are allowed to archive data. [“NetWorker User Groups” on page 559](#) provides more information.

Users can only retrieve data that they own. If other users need to retrieve data they do *not* own, then enable public archives access.

Enabling public archive access

To allow users listed in the Archive User Group to retrieve any archived files from a client:

1. In the **Administration** window, click **Configuration**.
2. In the left pane of the **Configuration** window, select the NetWorker server.
3. From the **File** menu, select **Properties**.
4. Select the **Public Archives** attribute and then click **OK**.

If, during recovery, the operating system allows you to change the ownership of archived data to that of the original owners, then the retrieved files display the original ownerships. Otherwise, the user who retrieves the files becomes the owner of the files.

About archive pools

The NetWorker software provides these three preconfigured pools to receive archived data:

- ◆ Preconfigured Indexed Archive pool
- ◆ PC Archive pool
- ◆ Preconfigured Archive pool

You cannot change the settings for these preconfigured pools, although you can create custom pools for archiving data. Custom pools can use either indexed or nonindexed archiving. [“Creating custom Archive pools” on page 329](#) provides information on creating custom Archive pools.

If you do not specify a pool to store archived data, the NetWorker software uses the Indexed Archive pool by default.

Preconfigured Indexed Archive pool and PC Archive pool

The preconfigured Indexed Archive pool and the PC Archive pool store entries for individual files in the client file index.

Note: Use of the Indexed Archive pool or the PC Archive pool may create a large client file index that never expires.

Preconfigured archive pool

The preconfigured Archive pool does not have a browsable client file index associated with it. You cannot retrieve individual files from the archive save set. Instead, you must retrieve the entire save set.

Creating custom Archive pools

Two attributes in the Pool resource distinguish Archive pools from other pools:

- ◆ Pool Type - This attribute must be set to Archive, which tells the NetWorker server that volumes belonging to this pool are used for archiving.
- ◆ Store Index Entries - This attribute determines whether the archive is an indexed or nonindexed archive:
 - If this attribute is set to No, entries are *not* written to the client file index (nonindexed archiving).
 - If this attribute is set to Yes, entries are written to the client file index (indexed archiving).

[“Media pools” on page 304](#) provides details and procedures about creating pools.

Archiving data procedures

You can request manual archives from the client, or you can schedule archives from the server.

Enabling archive services for a NetWorker client

To enable archive services for a NetWorker client:

1. In the **Administration** window, click **Configuration**.
2. In the left pane of the **Configuration** window, select **Clients**:
 - a. If you are creating a new client, select **New** from the **File** menu.
 - b. If you are editing an existing client, select the client and then select **Properties** from the **File** menu.
3. On the **Globals (2 of 2)** tab, enable the **Archive Services** attribute. When you enable archive services for one Client resource, archive services are enabled for all Client resources with the same hostname.
4. Make the remaining configuration choices as appropriate. The computer is now an enabled archive client. However, an archive will not occur until it is requested. [“Scheduling data archives” on page 331](#) provides instructions.

Note: If the NetWorker client is set up for encryption with the **aes** ASM, archive data will also be encrypted. [“Encrypting backup data” on page 108](#) provides information about setting up encryption for a NetWorker client.

Manually archiving data

You can manually archive data at any time. Manually archiving data is similar to performing a manual backup.

Perform a manual archive from a NetWorker client on Windows

Note: Manual archives that are performed 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. Scheduled archives, enforce all directives. For more information about scheduled archives, see [“Scheduling data archives” on page 331](#).

To perform a manual archive for a Windows client:

1. In the NetWorker **User** program, click **Archive** to open the **Archive Options** dialog box.
2. Type a comment in the **Annotation** attribute. This annotation is used to uniquely identify each archive save set during retrieval.

Note: Consider adopting a consistent naming convention so that one can easily identify archives based on the annotation name.

3. From the **Archive Pool** list, select the appropriate archive pool.

Note: Only pools with their Pool Type attribute set to Archive are listed.

4. Select the appropriate settings for these criteria:
 - To write a copy of each archive save set to a volume from an archive clone pool, select **Clone**.
 - If you enable cloning, type or select an archive clone pool for the **Archive Clone Pool** attribute.
 - To instruct the NetWorker server to check the integrity of the data on the storage volume, select **Verify**.
 - To instruct the NetWorker server to remove the archived files from the disk, select **Grooming**.
5. Click **OK**. The **Archive browse** window appears.
6. From the **File** menu, select **Mark** to select each file or directory for archiving. When you select an item for archiving, a check mark appears next to that item.

Note: To clear an item currently marked for backup, select **Unmark** from the **File** menu.

7. From the **File** menu, select **Start Archive**.
8. Click **OK** and the **Archive** browse window appears. The NetWorker server appears in the **Archive Status** window, which monitors the progress of the archive. When the NetWorker server is finished archiving, a message similar to this appears in the **Archive Status** window:

```
Archive completion time: 2-21-09 5:18p
```

9. If **Groom** was selected in [step 4](#) , the **Remove Archived File** dialog box prompts for confirmation before NetWorker software deletes archived files from the local disk.

Perform a manual archive from a NetWorker client on UNIX

To perform a manual archive from a UNIX client, use the **nsrchive** command. For information about this command, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

Scheduling data archives

Unlike scheduled backups, scheduled archives run only once. The advantage of a scheduled archive is that the archive can be run when network traffic and computer use is low.

Scheduling an archive

Before you can schedule an archive request, enable the Archive Services attribute in the Client resource. [“Enabling archive services for a NetWorker client” on page 330](#) provides more information.

To schedule an archive:

1. In the **Administration** window, click **Configuration**.
2. In the left pane, select **Archive Requests**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a name for the archive request.
5. In the **Comment** attribute, type a description of the archive request.
6. In the **Client** attribute, type the archive client's hostname.
7. In the **Save Set** attribute, type the pathnames of the files and/or directories to be archived.

Note: If you archive all of the client's save sets, set the **Grooming** attribute (on the **Running** tab) to **None**. If this attribute is set to **Remove**, all of the archived save sets will be deleted from the client computer.

8. Type a comment in the **Annotation** attribute. This annotation is used to uniquely identify each archive save set during retrieval.

Note: Consider adopting a consistent naming convention so that one can easily identify archives based on the annotation name.

9. In the **Directive** attribute, select a directive if special processing is to occur during the archive process. [Chapter 9, "Directives"](#) provides more information about directives.
10. From the **Archive Pool** attribute, select the appropriate pool from the list:
 - To store the entire save set, select the preconfigured **Archive** pool. This pool does *not* store the client file index.
 - To store the client file index in addition to the entire save set, select the preconfigured **Indexed Archive** or the **PC Archive** pool.
11. Select the **Running** tab.
12. For the **Status** attribute, indicate a start time for the archive:
 - To begin the archive immediately, select **Start Now**.
 - To begin the archive at a specified time, select **Start Later** and indicate a time in 24-hour format in the **Start Time** attribute.
13. For the **Archive Completion** attribute, type a notification for the NetWorker server to use after completing the archive. ["Indexes" on page 585](#) provides details.

14. Select the appropriate response for these options:

- To instruct the NetWorker server to remove the archive files from the disk, select **Remove** from the **Grooming** list.
- To instruct the NetWorker server to check the integrity of the data on the storage volume, select the **Verify** attribute.
- To write a copy of each archive save set to a volume in an **Archive Clone pool**, select **Yes** for the **Clone** attribute and select an archive clone pool from the **Archive Clone Pool** list.

15. Click **OK**.

To view information about the status of an archive request, open the **Archive Request Details** window. [“Viewing details of a scheduled archive” on page 337](#) provides more information.

Copying an Archive Request resource

To copy an Archive Request resource:

1. In the **Administration** window, click **Configuration**.
2. In the left pane, select **Archive Requests**.
3. In the right pane, select the archive request to copy.
4. From the **Edit** menu, select **Copy**. The **Create Archive Request** dialog box displays the same information as the archive request that was copied, except for the Name attribute.
5. Type the name for the new archive request in the **Name** attribute, edit any other attributes as appropriate, and click **OK**.

For details about the **Archive Request** attributes, click **Field Help** in the **Properties** dialog box.

Changing the archive time

To change the archive time:

1. In the **Administration** window, click **Configuration**.
2. In the left pane, select **Archive Requests**.
3. In the right pane, select an archive request.
4. From the **File** menu, select **Properties**.
5. Click the **Running** tab.
6. In the **Start Time** attribute, type a new time in this format:

HH:MM [a,p]

7. Click **OK**.

You can also schedule an existing archive by using the **Schedule Archive** operation in the **Activities Monitor**. [“Archive request management” on page 337](#) provides more information.

Editing an archive request

To edit an archive request:

1. In the **Administration** window, click **Configuration**.
2. In the left pane, select **Archive Requests**.
3. In the right pane, select an archive request.

4. From the **File** menu, select **Properties**.
5. Edit the attributes of the archive request and click **OK**.

Deleting an archive request

Note: You cannot delete an archive request that is currently in use.

To delete an archive request:

1. In the **Administration** window, click **Configuration**.
2. In the left pane, select **Archive Requests**.
3. In the right pane, select an archive request.
4. From the **File** menu, select **Delete**.

Retrieving archived data

This section describes how to retrieve archived data.

Retrieval permissions

The following restrictions apply when retrieving or recovering archived data:

- ◆ You must have read permissions to the archived data.
- ◆ If the Server resource's Public Archives attribute is enabled, all users listed in the Archive Users User Group can retrieve data (as long as they have read permissions to the data).
- ◆ If the Server resource's Public Archives attribute is disabled, only the users in the Archive Users User Group that own the file can retrieve the data.

[“Permissions for archiving” on page 328](#) provides more information.

Note: If, during retrieval, the operating system allows you to change the ownership of archived data to that of the original owners, then the retrieved files display the original ownership. Otherwise, the user who retrieves the files becomes the owner of the files.

Retrieving archives from a client on UNIX

To retrieve archive data for a UNIX client, use the **nsrretrieve** program.

Note: You cannot browse the archive data on a UNIX client.

To retrieve archived data:

1. Mount the archive volume in the appropriate storage device.
2. Start the **nsrretrieve** program. Type:

```
nsrretrieve -s NetWorker_servername -A annotation -S ssid/cloneid
-i_recover_option path
```

- ◆ *NetWorker_servername*—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 ssid.
- ◆ *-S ssid/cloneid*—Specifies the archive save set to recover. To recover a cloned archive save set, specify the ssid and cloneid. You must specify at least one annotation or ssid.
- ◆ *-i_recover_option*—Specifies how the NetWorker server should handle a naming conflict between a recovered file and an existing file.

For example:

- 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 **nsrretrieve** man pages and the *EMC NetWorker Command Reference Guide* provides more information about additional options available with the **nsrretrieve** command.

Example 25 Using the nsrretrieve command

In this example, suppose that archive A is annotated with *Accounting_Fed* and archive B is annotated with *Accounting_Local*.

- ◆ If you type this command:

```
nsrretrieve -A Accounting
```

No match is found and no archive is retrieved.

- ◆ If you entered this command:

```
nsrretrieve -A ting_L
```

Archive B is located.

Retrieving nonindexed archives from a client on Windows

Data that was archived with a nonindexed archive pool (such as the Archive pool) must be retrieved by the entire save set rather than by file selection.

To retrieve nonindexed archived data:

1. Mount the archive volume in the appropriate storage device.
2. Start the NetWorker **User** program.
3. From the **Operation** menu, select **Archive Retrieve** to open the **Source Client** window.
4. Select a client to retrieve the archived data from and click **OK**. The **Archive Retrieve** window opens.

5. For the **Annotation String** attribute, type all or part of the annotation assigned to the save set when it was archived.

Note: If no annotation is entered, all archived save sets for the client appear.

6. Click **OK**. The **Save Sets** dialog box opens.
7. From the **Save Sets** dialog box, perform either of these functions, if necessary:
 - To view a list of volumes required to retrieve the data from this archived save set, click **Required Volumes**.
 - 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 retrieve and click **OK**. The **Retrieve Status** window appears.

Note: You can also recover archived save sets by using save set recovery. [“Save set recover by using NetWorker User” on page 385](#) provides more information.

Recovering indexed archive data from a client on Windows

Data archived by using the Indexed Archive pool and the PC Archive pool maintain information in the client file index about the individual files in the save set.

The archived files are recovered the same way as nonarchived files. To recover indexed archived files, the archive must have been saved by using the Indexed Archive pool, PC Archive pool, or be a custom archive pool with the Store Index Entries attribute in the Pool resource enabled.

To recover indexed archived data:

1. Start the NetWorker **User** program.
2. Click **Recover** to open the **Source Client** dialog box.
3. Select the source client whose data is to be recovered, and click **OK**. The local client is the default selection.
4. Select the destination client for the recovered data, and click **OK**. The local client is the default selection.
5. Select the files to be recovered and click the **Mark** button.
6. Click **Start**.

Archive request management

This section describes how to work with scheduled archive requests.

Starting a scheduled archive at any time

You can start a scheduled archive immediately rather than wait for the scheduled start time.

To start a scheduled archive:

1. In the **Administration** window, click **Monitoring**.
2. Select the **Archive Requests** tab.
3. Right-click the archive request and select **Start**.

Stopping a scheduled archive while in progress

To stop an archive request in progress:

1. In the **Administration** window, click **Monitoring**.
2. Right-click the archive request and select **Stop**.

Disabling a scheduled archive

To disable an archive request:

1. In the **Administration** window, click **Monitoring**.
2. Right-click the archive request and select **Disable Archive**.

Viewing details of a scheduled archive

To open the **Archive Request Details** window:

1. In the **Administration** window, click **Monitoring**.
2. Right-click the archive request and select **Show Details**.

The **Archive Request Details** window provides information about the completion of an archive request:

- ◆ The **Completion Time** displays the time the archive finished. This is the difference between the completion and start times of the archive.
- ◆ The success of the archive request is either completed, failed, or partial.

CHAPTER 12

Cloning

This chapter covers these topics:

◆ Overview of cloning.....	340
◆ Save set cloning.....	341
◆ Specifying browse and retention policies for clone data	349
◆ Volume cloning	350
◆ Recovering cloned data	351
◆ Cloning archived data	353
◆ Directing clones to a special storage node.....	354
◆ Using file type devices for clone operations.....	357
◆ Backup-to-tape for Avamar deduplication clients	358
◆ Cloning with Data Domain devices	358
◆ Using the nsrclone command	358

Overview of cloning

Cloning allows for secure offsite storage, transfer of data from one location to another and verification of backups. Cloning can be performed on volumes and on save sets.

Information about the volumes, status, and history of cloning operations can be viewed and monitored from the Administration window. Clone-related messages are also logged to the NetWorker message file and the savegrp log file, which are located in the <NetWorker_install_dir>\logs directory. [“Viewing log files” on page 803](#) provides information about viewing log files.

There are two main methods of cloning:

- ◆ **Save set cloning:**
Save sets can be cloned based on a schedule or on-demand by manual selection. [“Save set cloning” on page 341](#) provides more information.
- ◆ **Volume cloning:**
Backup volumes can be cloned on demand by manual selection. [“Volume cloning” on page 350](#) provides more information.

Note: Cloning works differently for deduplication devices. The *NetWorker Avamar Integration Guide* and the *NetWorker Data Domain Depulication Devices Integration Guide* provides more information.

Cloning requirements

The following requirements apply when performing cloning operations:

- ◆ A minimum of two storage devices must be enabled: one to read the existing data and one to write the cloned data:
 - If libraries with multiple devices are used, the NetWorker server mounts the volumes required for cloning automatically.
 - If stand-alone devices are used, mount the volumes manually. A message displays in the Alert tab of the Monitoring option that indicates which volumes to mount.
- ◆ The destination volume must be a different volume from the source volume, and must belong to a clone pool.
- ◆ You must be a member of the NetWorker Administrators group. [“NetWorker User Groups” on page 559](#) provides information.
- ◆ Only one clone of a particular save set can reside on a single volume. Therefore, if three clones of the same save set are specified, each clone is written to a separate volume.

Save set cloning

NetWorker provides the ability to schedule clone operations by using a Clone user interface in NMC.

The following topics are included in this section:

- ◆ [“Considerations for scheduled clone jobs” on page 341](#)
- ◆ [“Setting up a schedule clone job” on page 342](#)
- ◆ [“Starting a scheduled clone job manually” on page 344](#)
- ◆ [“Monitoring scheduled clone jobs” on page 345](#)
- ◆ [“Setting up automatic cloning from a backup group” on page 345](#)
- ◆ [“Viewing the clone status of a save set” on page 346](#)
- ◆ [“Cloning a save set manually” on page 346](#)
- ◆ [“Additional manual clone operations” on page 349](#)

Considerations for scheduled clone jobs

Be aware of the following considerations when setting up scheduled clone jobs:

Scheduling multiple clone jobs to start at the same time

Do not schedule more than 30 clone jobs to start at the same time. Scheduling 30 or more clone jobs to occur at the same time may result in some clone jobs timing out and not completing.

Mixing save sets from different source devices

Clone operations that mix save sets from different source devices, such as Data Domain devices, AFTD devices, or NDMP devices, may be written to different target volumes. Although this behavior is by design, you may prefer to write all save sets in the clone operation to the same clone volume.

If the clone operation includes save sets from different devices, and you want all save sets to be written to the same volume, include only one volume in the clone target pool.

Unmounted clone source volumes on remote storage nodes

If the clone source volume is on a remote storage node and is unmounted, attempting to start a regular volume clone operation will not complete successfully, even if the source volume is mounted after the clone operation attempts to start. The clone program **nsrclone** will hang with the following message:

```
Server <server_name> busy, wait 30 second and retry
```

This issue does not occur in the following situations:

- ◆ If the storage node is on the NetWorker server, that is, when the storage node is not remote.
- ◆ If performing a clone controlled replication (optimized clone) operation.

Snapshot with scheduled cloning

To set up a scheduled clone for a snapshot backup, you must specify the backup group to which the snapshot rollover belongs. When setting up a scheduled clone job, specify the backup group in the **Save Set Filters** tab.

Clone resources that are created with the nsradmin program

Clone resources (known as NSR clone resources) that are created with the **nsradmin** command line program cannot be edited as scheduled clone resources in the NetWorker Administration graphical user interface.

To avoid this issue, perform one of the following:

- ◆ Create scheduled clone resources in the Administration interface. [“Setting up a schedule clone job” on page 342](#) provides more information.
- ◆ If you must create a NSR clone resource with the **nsradmin** program, create a corresponding NSR task resource with the **nsradmin** program. Together, these resources will enable you to edit the clone item as a scheduled clone resource in the GUI. The corresponding NSR task resource must have its **name** and **action** attributes specified as follows:

- **name:** “clone.nsrclone_resource_name”
- **action:** “NSR clone:nsrclone_resource_name”

For example, if the NSR clone resource was named *TestClone1*, the **name** and **action** attributes of the NSR task resource would be:

- **name:** clone.TestClone1
- **action:** NSR clone: TestClone1

These entries are case-sensitive.

Setting up a schedule clone job

To set up a scheduled clone operation:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Clones**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a unique name to identify the scheduled clone resource. Type additional information in the **Comment** attribute if necessary.
5. If you wish to override the saveset’s original browse and retention policies, select new policies in the **Browse** and **Retention** policies attributes.
6. To specify the storage node that will write data during the clone operation, select a storage node from the **Storage node to WRITE save sets attribute**. If a selection is made in this attribute, it will override any selection criteria that is described in [“Storage node selection criteria and settings for writing a clone” on page 355](#).

This attribute is used primarily in conjunction with the **Storage node to READ save sets** attribute to balance access to storage media across different storage nodes.

7. To specify the storage node that will read data during the clone operation, select a storage node value from the **Storage node to READ save sets** attribute. This attribute is visible only when Diagnostic mode is selected. Additionally, the selected storage node must be included in at least one of the following:
- The **Recover storage nodes** or **Storage nodes** attribute of the NetWorker server's client resource.
 - The storage node listed in the **Read Hostname** attribute for the library resource, if a library is being used.
 - A storage node on which any device in the library is configured, if a library is being used.

This attribute is used primarily in conjunction with the **Storage node to WRITE save sets** attribute to balance access to storage media across different storage nodes and is not intended for use with standalone devices such as AFTD's, file type devices, Data Domain devices, and so on.

8. To specify the clone media pool to write data to during a clone operation, select a clone type media pool from the **Write clone data to pool** attribute. If no selection is made, clones will be written to the default clone pool.

Pools are used to direct backups to specific media volumes. This attribute is particularly useful when you want to ensure that only certain media types are used to hold clone data. For example, to ensure that this clone job only replicates to a certain type of disk, such as a Data Domain type disk, select a clone pool that uses only Data Domain type disks. Likewise, to ensure that this clone job only replicates to tape (tape out), select a clone pool that uses only tape devices.

9. Select **Continue on save set error** to force NetWorker to skip invalid save sets and to continue the clone operation. If this option is not selected (default setting), an error message will be generated and the clone operation will not continue if an invalid save set or invalid volume identifier is encountered.
10. To restrict the number of clone instances that can be created for any save set that is included in this particular scheduled clone operation, type a value in the **Limit number of save set clones** attribute. A value of zero (0) means that an unlimited number of clones may be created for this scheduled clone operation.

Consider limiting the number of save set clones in cases where the clone operation has not completed and is being retried. For example, if you type a value of **1** in this attribute and then retry a partially completed clone operation, only the save sets that were not successfully cloned the first time will be eligible for cloning. In this way, unnecessary clone instances will not be created.

Regardless of the value in this attribute, NetWorker always limits the number of save set clone instances to one per volume. A clone pool can have multiple volumes. This attribute limits the number of save set clone instances that can be created for a clone pool in a particular scheduled clone operation.

11. Select **Enable** to allow the clone job to run at its scheduled time.
12. In the **Start Time** attribute, click the up and down arrows to select the time to start the clone job. Alternatively, type the time directly into the attribute fields.

13. From the **Schedule Period** attribute, select **Weekly by day** or **Monthly by day** depending on how you want to schedule the clone job and then select the day(s) of the week or month on which the scheduled clone is to occur.
14. To repeat the clone job within a day, specify an **Interval** time in hours. For example, if the start time is **6 AM**, and the interval is **6** hours, then the clone job will run at 6 AM, 12 PM, and 6 PM.

If the **Limit the number of save set clones** value is set, then the repeat clone job will fail after the limit is reached.

15. Select the **Save Set Filters** tab to specify the save sets to be included in this scheduled clone job.
16. Select **Clone save sets that match selections** to limit save sets by various filter criteria or select **Clone specific save sets** to explicitly identify the save sets to be cloned.

To clone save sets that match selection criteria:

- Specify selection criteria to limit the save sets that will be included in this scheduled job. You can select the following criteria:
 - Groups (save groups)
Required for snapshot rollover clones.
 - Clients (client resources)
 - Pools (backup pools)
 - Filter save sets by level (backup level)
 - Filter save sets by name (save set name as specified in the client resource)
 - Include save sets from the previous (save sets from the past number of days, weeks, months, or years)

To display a list of the save sets that will be cloned based on the filter criteria that you specified, select **Preview Save Set Selection**.

To clone specific save sets:

- Type the specific save set ID/ clone ID (ssid/clonid) identifiers in the **Clone specific save sets** list box. Type each ssid/clonid value on a separate line.

You can query save set IDs / clone IDs through the Administration > Media user interface or by using the **mminfo** command. [“Querying the media database” on page 588](#) provides more information.

17. Select **OK** to save the scheduled clone job.

Starting a scheduled clone job manually

You can start a scheduled clone job at any time without affecting the regularly scheduled start time.

To start a scheduled clone job manually:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Clones**.

3. Right-click on a clone resource in the right-pane and select **Start**.

You can also start a scheduled clone from the NetWorker Monitoring feature, which is described in [“Clones window” on page 467](#).

Monitoring scheduled clone jobs

You can view the status of scheduled clone jobs in the Monitoring window. You can view the scheduled clone’s last start and end time and you can view the completion status of each save set that is included in the scheduled clone. [“Clones window” on page 467](#) provides more information.

Note: If you change the clone times for an existing scheduled clone job, the Monitoring window will show the time of the old scheduled clone time until the updated schedule is executed.

Setting up automatic cloning from a backup group

You can also set up automatic clone operations for a backup group. The clone operation can be set to start immediately after each save set in the group is backed up (immediate cloning) or the clone operation can be set to start only after all save sets in the group are backed up.

Immediate cloning operations can complete sooner because they can run in parallel instead of sequentially. Performance gains with immediate cloning are most noticeable when there are many savesets in the backup queue or when there are many savesets of different sizes. Immediate cloning is only supported with clone-controlled replication (CCR) using DD boost devices.

[“Setting up a schedule clone job” on page 342](#) provides alternate clone methods that provide more flexibility than what is available with the automatic group clone method described in this section.

Note: All of the save sets that are associated with the group are backed up the first time the automatic clone operation is run regardless of whether the previous backup was full or incremental. Subsequent automatic clone operations for the same group will clone only those save sets that have changed since the previous backup.

To automatically clone the save sets that belong to a group resource:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Groups**.
3. Select the group in the right-pane.
4. From the **File** menu, select **Properties**.
5. Select the **Setup** tab.
6. Select the **Clones** attribute.

7. Select a value from the **Clone mode** attribute.
 - Select **Start on save set completion** to start a save set clone operation after the save set is backed up. If this option is selected, the NetWorker server parallelism attribute must be set to a value of 2 or higher. To access the server parallelism attribute, right-click the NetWorker server name in the left pane of the **Administration** window, select **Properties** and then select the **Setup** tab.

This option is supported only when performing clone-controlled replication with DD Boost devices. If this option is selected for non-DD Boost devices, it will fall back to the **Start on Group Completion** option.
 - Select **Start on group completion** to start clone operations only after all savesets in the group are backed up.
8. Select a clone pool from the **Clone pool** attribute.

Viewing the clone status of a save set

To determine whether a save set on a volume has been cloned, or is itself a clone, check the Query Save Set tab window. [“Cloning a save set manually” on page 346](#) provides more information.

Cloning a save set manually

To manually clone a save set, first query the database, select the save set, and begin the cloning operation.

To manually clone a save set:

1. From the **Administration** window, click **Media**.
2. In the expanded left pane, select **Save Sets**.
3. In the right pane, select the **Query Save Set** tab. Use the **Query Save Set** tab to specify options to limit the range of save sets displayed. All query options are optional except for the date. A date range must be selected.

Note: The text boxes in the **Query Save Set** tab are case-sensitive.

4. Type values in any of these attributes to limit the search:
 - Client Name
 - Save Set
 - Save Set ID
 - Volume
 - Pool
5. Use the **Copies** attribute to limit the search to only those save sets that have already been cloned:
 - a. Select a boolean value of greater than (>), equal to (=), or less than (<) from the list.
 - b. Type the number of clones to complete the search criteria for the **Copies** attribute.

For example, to search for only those save sets that have been cloned at least twice, select greater than (>) and then type **1** as the number of copies.

- Use the **Save Time** attribute to limit the search to a period of time in which the save set was created.

By default, yesterday is used for the start date, and today is used for the end date. This means that save sets backed up between yesterday at 12:01 A.M. and the current time will be displayed.

For the **From** and **To** date fields, any of these formats are acceptable:

- Written out completely (for example, November 1, 2009)
- Numerically as mm/dd/yy (for example, 11/01/09)
- Date and time selection from the list.

A long date range may result in too many selected save sets. This can increase response time or even require that you close and reopen the browser connection to the NetWorker Console.

- Use the **Status** attribute to limit the search to save sets that have a particular status. [Table 47 on page 347](#) lists the values that can be selected.

Table 47 Save set status settings

Status	Description
All	Select all options listed under Select from in the Status area.
Select from	Select one or more of the following options: <ul style="list-style-type: none"> • Browsable: Select if the save set still has an entry in the client file index. • Recyclable: Select if all save sets have passed both the browse and retention policy time periods; the volume may now be available for automatic relabeling and overwriting (provided all save sets on the volume are recyclable). • Recoverable: Select if the entry for the save set has been removed from the client file index, but is still available for recovery from the media (that is, the volume has not passed its retention policy). • In-progress: Select if the save set is currently in the process of being backed up. In-progress save sets cannot be cloned. • Aborted: Select if the save set was either aborted manually by the administrator during a backup, or because the computer crashed. Aborted save sets cannot be cloned. • Suspect: Select if a previous attempt to recover the save set failed.

- Use the **Maximum Level** attribute to limit the search to save sets of a particular backup level.

The level All is specified by default. All the levels up to and including the selected level are displayed. For example:

- If you select level 5, save sets backed up at levels full, 1, 2, 3, 4, and 5 are displayed.
- If you select level Full, only those save sets backed up at level full are displayed.
- If you select All, save sets for all levels are displayed.

- Click the **Save Set List** tab.

The save sets that fit the criteria appear in the **Save Sets** list.

10. Select the save sets to clone from the **Save Set** list.
11. From the **Media** menu, select **Clone**.
12. From the **Target Clone Media Pool** list, select a clone pool.
13. Click **OK**, then click **Yes** on the confirmation screen.

[“Viewing manual clone history” on page 349](#) provides information on viewing the status of a manual clone operation or to cancel a clone operation that is in progress.

Example 26 Manual cloning of a save set

In this example, a user has requested that the NetWorker administrator manually clone several save sets that are not included in an automatic cloning schedule.

The user must send the data to another company located out of country. The administrator must clone the most recent full backup, and any incrementals since the last full backup, to make sure that the most current data is sent.

To clone the save set, the NetWorker administrator must have this information:

- ◆ NetWorker client name
- ◆ Name of the save set
- ◆ Date the data was backed up

To manually clone the save sets:

1. From the **Administration** window, click **Media**.
2. In the expanded left pane, select **Save Sets**.
3. In the right pane, select the **Query Save Set** tab.
4. In the **Client Name** attribute, type the client name.
5. In the **Save Set** attribute, type the save set name.
6. For the **Start Date** and **End Date**, type the dates used for the search.

Note: In this case, the administrator does not need to change or select any status choices other than the defaults.

7. Click the **Save Set List** tab. All save sets that meet the selection criteria appear in the **Save Sets** list.
8. Select the save sets to clone.
9. From the **Save Set List**, determine the size of the data and the original volume that contains the data to be cloned.
10. Mount the original volume.
11. From the **Media** menu, select **Clone**.
12. From the **Target Clone Media Pool** list, select a clone pool.
13. Click **OK** and then click **Yes** on the confirmation screen.

Additional manual clone operations

This section covers operations that can be performed on both volumes and save sets that have been manually cloned. [“Volume cloning” on page 350](#) or [“Save set cloning” on page 341](#) provide information about manually cloning a volume or save set.

Viewing manual clone history

To view history information about manual clone operations:

1. From the **Administration** window, click **Monitoring**.
2. From the **Monitoring** menu, select **Show Manual Clone History**.

A dialog box appears that shows manual clone history information.

Stopping a manual clone operation

To stop a manual clone operation that is in progress:

1. From the **Administration** window, click **Monitoring**.
2. From the **Monitoring** menu, select **Show Manual Clone History**.
3. Select the clone operation to be stopped.
4. Click **Stop Selected Operation**.

The manual clone operation is stopped.

Specifying browse and retention policies for clone data

The browse and retention policy for clone data can be specified independently of the original save set. In this way, clone data can be browsed and recovered after the policies of the original save set have expired.

To specify the browse or retention policy for clone data, perform one of the following:

- ◆ Specify a browse and retention policy in a scheduled clone job
- ◆ Specify a retention policy in the Clone pool.
- ◆ Specify a retention policy from the command prompt.

Specify a browse and retention policy in a scheduled clone job

To specify a browse and retention policy in a scheduled clone job:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Clones**.
3. From the **File** menu, select **Properties**.
4. Select new policies in the **Browse** and **Retention** policies attributes. When the scheduled clone job is next run, the cloned save sets will be given the new browse and retention policies.

Specify a browse and retention policy from the command prompt

To specify a retention policy from the command prompt, perform one of the following:

- ◆ Use the **nsrclone** command with the **-y** option when creating a clone save set.
- ◆ Specify a retention policy for an existing clone save set by using the **nsrmmm -e** command.

To specify a browse policy from the command prompt:

- ◆ Use the **nsrclone** command with the **-w** option when creating a clone save set. However, be aware that this will also change the browse policy of the original save set instance if the original save set's browse time has not passed and is earlier than the new browse time for the clone.

Specify a retention policy for a Clone pool

You can only specify a retention policy for cloned data in a pool resource.

To specify a retention policy for a Clone pool:

1. In the clone pool to which clone backups will be directed, select the **Configuration** tab.
2. From the **Retention** policy list, select a retention policy, then click **OK**.

[“Configuring media pools” on page 308](#) provides information about editing or creating a pool.

Note: Retention policies that are specified in a scheduled clone job or from the command prompt, override the retention policy specified in a clone pool.

Volume cloning

Volume cloning is the process of reproducing complete save sets from a storage volume to a clone volume. You can clone save set data from backup or archive volumes.

Creating a clone volume

To create a clone volume:

1. From the **Administration** window, click **Media**.
2. In the expanded left pane, select **Volumes**.
3. In the right pane, select the volume to be cloned.
4. From the **Media** menu, select **Clone**.
5. From the **Target Clone Media** Pool list, select a clone pool.
6. Click **Ok**, then click **Yes** on the confirmation screen.

Viewing clone volume details

You can view the details of a clone volume, such as the amount of space used, mode, expiration date, pool, and save sets.

To view clone volume details:

1. From the **Administration** window, click **Media**.
2. In the expanded left pane, select **Volumes**.
3. In the right pane, details for each volume are displayed in the table.
4. To view save set information for a volume:
 - a. In the right pane, select a volume.
 - b. From the **Media** menu, select **Show Save Sets**.

Recovering cloned data

The NMC Recovery Wizard, the save set recovery option in NetWorker User, and the **recover** command provides you with the ability choose which volume (original or clone) to use when recovering data.

NetWorker decides which volume to use when you:

- ◆ Perform a browsable recovery in NetWorker User.
- ◆ Perform a browsable recovery with the **recover** command and do not specify the clone pool.
- ◆ Perform a save set recovery with the **recover** command and do not specify the cloneid of the clone save set.
- ◆ Allow the NMC Recovery Wizard to select the required volumes for recovery.

NetWorker bases the volume selection on the following criteria:

1. The highest priority is given 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.
[“Changing save set status within the Volume Save Sets window” on page 243](#) provides information about changing the status of a save set.
2. If the volumes still have equal priority, then priority is given 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:
 - a. Advanced file type device
 - b. File type device
 - c. 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:
 - a. Volumes in a library.
 - b. Volumes in an AlphaStor or SmartMedia controlled library.
 - c. Volumes that are not in a library but are onsite (**offsite** flag is not set).

- d. Volumes that are offsite (**offsite** flag is set).

Use the **nsrmm** command to specify that a volume is offsite. For example:

```
nsrmm -o offsite -V volume_id
```

The volumes required for recovery appear in the Required Volumes window of the NMC Recovery Wizard and the NetWorker User (Windows) programs.

You can also run the **scanner** program on a clone volume to rebuild entries in the client file index, the media database, or both. After you re-create the entries, normal recovery is available. [“Recovering a recyclable or recoverable save set entry in the online indexes” on page 399](#) provides more information.

Recovering a clone save set from the command prompt

When you use the **recover** command to perform a browsable or save set recover, you can recover from the cloned save set.

Example 27 Performing a save set recover from a cloned save set

To perform a save set recover from a clone, use **mminfo** to determine the cloneid of the save set, then specify the cloneid in the **recover** command.

1. Use **mminfo** to determine the cloneid of the save set you want to recover:

```
mminfo -av -r volume, savetime, client, ssid, cloneid, name
```

volume	date	client	ssid	clone id	name
backup.001	05/03/2013	bu_iddnwserver	3644194209	1362492833	C:\ddlib
clone.001	05/03/2013	bu_iddnwserver	3644194209	1362493448	C:\ddlib

2. Specify the SSID and cloneid in the **recover** command.

For example:

```
recover -S 3644194209/1362493448
```

Example 28 Performing a browsable recover from a clone save set

To perform a browsable recover from a clone, use **mminfo** to determine the clone pool, then specify the pool name in the **recover** command.

1. Use **mminfo** to determine the clone pool that contains the clone save set you want to recover:

```
mminfo -r volume, savetime, client, ssid, cloneid, name
```

volume	date	client	pool	name
backup.001	05/03/2013	bu_iddnwserver	Default	C:\ddlib
clone.001	05/03/2013	bu_iddnwserver	Default Clone	C:\ddlib

- Specify the clone date and clone pool in the **recover** command.

For example:

```
recover -t 05/03/2013 -b "Default Clone"
```

- Recover the data. [“Using the recover command” on page 379](#) describes how to use the recover command.

Recovering a save set when all cloned instances have expired

When all cloned instances of a save set have passed their retention period, the following procedure must be used to mark a save set as eligible for recovery:

- Use the **nsrmm** command with the **-e time** option to change the retention time for the clone save set:

```
nsrmm -e time -S ssid/cloneid
```

If the *cloneid* is not identified with the **-S** option, the following error message is displayed:

```
Save set ssid cannot be marked as notrecyclable. Please specify the ssid/cloneid of the particular clone instance.
```

- Use the **nsrmm** command with the **-o notrecyclable** option to instruct the media database that the save set is no longer expired:

```
nsrmm -o notrecyclable -S ssid/cloneid
```

If the **-o notrecyclable** option is used with **nsrmm** prior to changing the retention time as described in [step 1](#), the following error message is displayed:

```
nsrmm: Save set ssid:ssid cloneid:cloneid eligibility cannot be cleared, retention time must be adjusted first
```

After this procedure has been completed, the save set is recoverable.

Cloning archived data

You can schedule a clone job to clone archive data or clone archive data on-demand.

To set up a scheduled clone job for archive data:

- ◆ Follow the steps in [“Setting up a schedule clone job” on page 342](#) and select an archive pool as one of your save set filter criteria.

To clone an archive volume on-demand:

- From the **Administration** window, click **Media**.
- In the expanded left pane, select **Save Sets**.
- In the right pane, select the **Query Save Set** tab.
- In the **Pool** attribute, select an archive pool from the list. Make other selections, as appropriate, to limit the save set search criteria. [“Cloning a save set manually” on page 346](#) provides more information.
- Click the **Save Set List** tab.

6. Select the archive save sets to clone from the **Save Set** list.
7. From the **Media** menu, select **Clone**.
8. From the **Target Clone Media Pool** list, select an archive clone pool.
9. Click **OK** and then click **Yes** on the confirmation screen.

Directing clones to a special storage node

You can direct clone operations to a specific storage node. This section describes the criteria used to determine the storage node from which the clone data is read (read source) and the storage node to which the clone data is written (write source).

Storage node selection criteria for reading the clone data

The following logic is used to determine the storage node from which the clone data will be read (read source):

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.
 - If the `FORCE_REC_AFFINITY` environment variable is set to **Yes**, the selection criteria in [step 1](#) is ignored and the selection criteria behaves as though the volume is not mounted as described in [step 2](#).
 - When cloning is used in a Virtual Tape Library (VTL) environment such as a CLARiiON Disk Library (CDL), the NetWorker software behaves as if the `FORCE_REC_AFFINITY` environment variable is set to **Yes**.
2. If the volume is not mounted or if the `FORCE_REC_AFFINITY` environment variable is set to **Yes**, a list of eligible storage nodes is created. The list is based on the storage nodes that meet the criteria in both [step a](#) and [step b](#) that follow:
 - a. The storage nodes listed in the **Recover Storage Nodes** attribute of the NetWorker server's client resource. If this attribute is empty, the NetWorker server's **Storage Nodes** attribute is used.
 - b. If the requested volume is in a media library, the storage nodes on which the volume can be mounted are determined in the following manner:
 - The storage node 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 is added to the list of eligible storage nodes.
 - If the volume is not in a media library, then the list of storage nodes is based on [step a](#) only.

Example 29 Selecting a storage node read source

In this example, the volume resides in a media library and is not mounted. The **Recover Storage Nodes** attribute in the NetWorker server's 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, however, 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. Therefore the list of eligible storage nodes is as follows:

- ◆ Storage node E
- ◆ Storage node D

The order in which the storage node is selected is based on the order of 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.

In this example, if no matching storage nodes were found in the intersecting list, an error would be written to the daemon log file that indicates no matching devices are available for the operation. To correct the problem, make adjustments so that at least one matching storage node can be found in both lists. [“Viewing log files” on page 803](#) provides information about viewing log files.

Storage node selection criteria and settings for writing a clone

The following priorities determine which storage node will store the cloned backup data. The storage node where the backup data resides is called the “read source.” The storage node that stores the cloned data is called the “write source”:

1. The read source storage node specifies the write source in its **Clone Storage Nodes** attribute.
2. If this attribute is blank, then the NetWorker server's storage node specifies the write source in its **Clone Storage Nodes** attribute.
3. If this attribute is also blank, then the NetWorker server's client resource specifies the write source in its **Storage Nodes** attribute.

Wherever the cloned data is written, the client file index and media database entries for the cloned save sets will reside on the NetWorker server.

Specifying a clone from a volume shared by multiple devices

In backup-to-disk environments it is possible for a single backup volume to be shared by multiple storage devices on different storage nodes. The Clone Storage Nodes attribute on each of these storage nodes can specify a different clone write source. Thus the write source for data cloned from the backup volume can be ambiguous depending on which device reads the volume.

To ensure unambiguous clone write sources in this situation, configure the Clone Storage Nodes attribute of all the storage nodes that have access to the backup volume to specify the same storage node write source.

Cloning from one storage node to another

To clone backup data from one storage node to another storage node:

1. In NMC, connect to the NetWorker server.
2. In the **Devices** view, select **Storage Nodes** in the navigation tree.
3. Right-mouse click the storage node where the backup data resides (read source storage node) and select **Properties**.
4. On the **Configuration** tab, in the **Clone Storage Nodes** attribute, type the hostname of the storage node that will store the cloned backup data.

Cloning from many storage nodes to one storage node

To clone backup data from many storage nodes to a single storage node:

1. In NMC, connect to the NetWorker server.
2. In the **Devices** view, select **Storage Nodes** in the navigation tree.
3. Right-mouse click the NetWorker server's Storage Node resource and select **Properties**.
4. On the **Configuration** tab, in the **Clone Storage Nodes** attribute, type the hostname of the storage node that will store all the cloned backup data
5. Configure each read source Storage Node resource to ensure that the **Clone Storage Nodes** attribute is blank.

Storage node selection criteria for recovering cloned data

The following logic is used to determine the storage node from which the clone data will be recovered:

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.

If the `FORCE_REC_AFFINITY` environment variable is set to **Yes**, the selection criteria in [step 1](#) is ignored and the selection criteria behaves as though the volume is not mounted as described in [step 2](#).

When cloning is used in a Virtual Tape Library (VTL) environment such as a CLARiiON Disk Library (CDL), the NetWorker software behaves as if the `FORCE_REC_AFFINITY` environment variable is set to **Yes**.

2. If the volume is not mounted or if the `FORCE_REC_AFFINITY` environment variable is set to **Yes**, a list of eligible storage nodes is created. The list is based on the storage nodes that meet the criteria in both [step a](#) and [step b](#) that follow:
 - a. The storage nodes listed in the **Recover Storage Nodes** attribute of the NetWorker client resource that is being recovered. If this attribute is empty, the NetWorker client's **Storage Nodes** attribute is used.
 - b. If the requested volume is in a media library, the storage nodes on which the volume can be mounted are determined in the following manner:
 - The storage node 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 is added to the list of eligible storage nodes.
 - If the volume is not in a media library, then the list of storage nodes is based on [step a](#) only.

Using file type devices for clone operations

This section discusses issues related to cloning with file type and advanced file type devices.

Differences in the cloning process

There are differences in the cloning process for the two types of devices:

- ◆ For file type devices, automatic and manual cloning begins only after all the save sets in a savegroup have been backed up.
- ◆ For advanced file type devices, automatic cloning begins only after all the save sets in a savegroup have been backed up. However, you can begin manually cloning a save set as soon as it has finished its backup.
- ◆ As of NetWorker 8.1, automatic cloning can begin after each save set in a savegroup is backed up when using clone-controlled replication (CCR) with a DD Boost device. [“Setting up automatic cloning from a backup group” on page 345](#) provides details.

Manual cloning with advanced file type device

In a situation where there are three save sets:

- ◆ Save set A has a size of 10 KB
- ◆ Save set B has a size of 10 MB
- ◆ Save set C has a size of 10 GB

When save set A has completed its backup, you can begin the manual cloning process while the other two larger save sets are still being backed up.

As each save set is backed up, you can launch the cloning process for that save set.

You can only clone one save set at a time.

Backup-to-tape for Avamar deduplication clients

The *NetWorker Avamar Integration Guide* provides more information.

Cloning with Data Domain devices

Data Domain devices were introduced in NetWorker 7.6 Service Pack 1 and enable one to perform clone controlled replication (optimized cloning) from one Data Domain device to another. You can also clone to tape or to any other device type.

Clone operations with Data Domain devices are set up in basically the same way as any other scheduled clone operation, which is described in [“Save set cloning” on page 341](#). However, there are some special considerations to be aware of when setting up Data Domain devices. These are described in the *EMC NetWorker Data Domain Devices Integration Guide*.

Using the nsrclone command

As of NetWorker 7.5, the **nsrclone** command has been enhanced to provide greater flexibility in selecting save sets for cloning by clients, groups, save set names, save set levels, and by number of valid copies or clones not yet created in the target pool. Also, be aware that as of NetWorker 7.6 Service Pack 1, most of the functionality provided in the **nsrclone** command is now provided in the Clone resource user interface. [“Setting up a schedule clone job” on page 342](#) provides more information.

[Table 48 on page 358](#) provides the descriptions of the new options, in NetWorker 7.5, that can be used with the **nsrclone** command.

Table 48 List of nsrclone options and their descriptions

Options	Description
-C <i>less than copies in target pool</i>	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 for cloning. This option is useful when retrying aborted clone operations. Because the target is a clone pool, each save set's original copy or clone is never considered when counting the number of copies of the save set. Likewise, any AFTD read-only mirror clone is not considered because its read or write master clone is counted and there is only one physical clone copy between the related clone pair. Recyclable, aborted, incomplete and unusable save set or clones are excluded in the counting. This option can only be used with the -t or -e option.
-l <i>level or range</i>	Specifies the level or n1-n2 integer range from 0 to 9 for save sets that are considered for cloning. Manual for ad-hoc or client-initiated save sets, full for level full save sets, incr for level incremental save sets, and integers 0 through 9, where save set0 also means full, can be used. More than one level can be specified by using multiple -l options and the -l n1-n2 range format. This option can only be used with the -t or -e option.

Table 48 List of nsrclone options and their descriptions

Options	Description
-N <i>save set name</i>	Specifies the save set name for save sets that are considered for cloning. More than one save set name can be specified by using multiple -N options. This option can only be used with the -t or -e option.
-c <i>client name</i>	Specifies the save sets in the particular client. More than one client name can be specified by using multiple -c options. This option can only be used with the -t or -e option.
-g <i>group name</i>	Specifies the save sets in the particular group. More than one group name can be specified by using multiple -g options. This option can only be used with the -t or -e option.

Examples

The following examples show how various options can be used with the **nsrclone** command:

Copy all save sets created in the last twenty-four hours for clients *mars* and *jupiter* with save set names */data1* and */data2* for only backup level full:

```
nsrclone -S -e now -c mars -c jupiter -N /data1 -N /data2 -l full
```

Copy all save sets that were not copied to the default clone pool in a prior partially aborted **nsrclone** session:

```
nsrclone -S -e now -C 1
```

Copy all save sets that were not copied to the default clone pool in a previous partially aborted nsrclone session and with extended retention and browse periods:

```
nsrclone -S -e now -C 1 -y 12/12/2010 -w 12/12/2009
```


CHAPTER 13

Staging Backups

This chapter covers these topics:

- ◆ [Save set staging.....](#) 362
- ◆ [Working with staging policies.....](#) 362

Save set staging

Save set staging is the process of transferring data from one storage medium to another medium, and then removing the data from its original location. For example, the initial backup data can be directed to a high performance file type or advanced file type device. In this way, the backup time is reduced by taking advantage of a file or advanced file type device. At a later time, outside of the regular backup period, the data can be moved to a less expensive but more permanent storage medium, such as magnetic tape. After the backup data is moved, the initial backup data can be deleted from the file or advanced file type device so that sufficient disk space is available for the next backup.

A save set can be staged from one disk to another as many times as required. For example a save set could staged from disk 1, to disk 2, to disk 3, and finally to a remote tape device or cloud device. Once the save set is staged to a tape or cloud device, it cannot be staged again. However, you could still clone the tape or cloud volume.

Staging can be driven by any of the following:

- ◆ Calendar-based process, such as keeping the save set for 30 days on the staging device before moving the data to the next device.
- ◆ Event-based process, such as when available space in the staging pool drops below a set threshold. When this happens, the oldest save sets are moved until available space reaches a preset upper threshold.
- ◆ Administrator-based process, such as allowing the administrator to either reset the threshold or manually select save sets to stage.

Staging does not affect the retention policy of backup data. Therefore, staged data is still available for recovery.

When the **nsrstage** process encounters an error after successfully cloning specified save sets, it deletes only those successful save sets from the source volume before the program is aborted. This ensures that only a single set of save sets exist in either of the source or clone volumes after staging.

Working with staging policies

This section describes how to work with staging policies.

[Chapter 4 “Backup to Disk and Cloud”](#) provides information on disk based device configuration.

Creating a staging policy

Before creating a staging policy configure all appropriate devices. Otherwise, no devices will be listed in the Devices attribute.

To prevent an advanced file type device from becoming full during backup, the staging policy must be set up so that save sets are automatically moved to another medium to make disk space available in the advanced file type device.

To create a staging policy:

1. In the **Administration** window, click **Configuration**.
2. In the left pane, select **Staging**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type a name for the staging policy.
5. In the **Comment** attribute, type a description of the staging policy.
6. To enable staging to begin immediately or to be invoked automatically at a later time, set the **Enabled** attribute to **Yes**.

You can enable or disable staging at any time.

7. In the **Devices** attribute, select the file type and adv_file type devices as the source device for staging. You can assign multiple devices to the staging policy, but a given device cannot be controlled by more than one staging policy.
8. For the **Destination Pool** attribute, select the destination pool for the staged data.

If the Clone pool you have selected is restricted to storage node devices, you will also need to modify **Clone Storage Nodes** attribute of the **Storage Node** resource for the NetWorker server to include the storage node name.

[“Storage node selection criteria and settings for writing a clone” on page 355](#) provides details.

9. In the **High-Water Mark (%)** attribute, type or select a number.

This value is the point at which save sets should be staged, measured as the percentage of available space used on the file system partition that the file device is on. Staging continues until the low-water mark is reached (see [step 10](#)).

The high-water mark must be greater than the low-water mark.

10. In the **Low-Water Mark (%)** attribute, type or select a number. This is the point at which the staging process will stop, measured as the percentage of available space on the file system partition that the file device is on.
11. From the **Save Set Selection** attribute, select from the list to determine the save set selection criteria for staging.
12. In the **Max Storage Period** attribute, type the number of hours or days for a save set to be in a volume before it is staged to a different storage medium.

The Max Storage Period attribute is used in conjunction with the File System Check Interval attribute. Once the Max Storage Period value is reached, staging does not begin until the next file system check.
13. In the **Max Storage Period Unit** attribute, select **Hours** or **Days**.
14. In the **Recover Space Interval** attribute, type the number of minutes or hours between recover space operations for save sets with no entries in the media database from file or advanced file type devices.
15. In the **Recover Space Interval Unit** attribute, select **Minutes** or **Hours**.

16. In the **File System Check Interval** attribute, type the number of minutes or hours between file system checks.

At every **File System Check** interval, if either the **High-Water Mark** or **Max Storage Period** has been reached, a staging operation is initiated.

17. In the **File System Check Interval Unit** attribute, select **Minutes** or **Hours**.
18. To invoke the staging policy immediately, complete this step. Otherwise, skip this step.

- a. Select the **Operations** tab.
- b. In the **Start Now** attribute, select one of these operations:
 - **Recover space** — Recovers space for save sets that have no entries in the media database and deletes all recycled save sets.
 - **Check file system** — Checks file system and stage data, if necessary.
 - **Stage all save sets** — Stages all save sets to the destination pool.

The selected operation applies to all devices associated with this policy.

The choice you make takes effect immediately after clicking **OK**. After the staging operation is complete, this attribute returns to the default setting (blank).

19. When all the staging attributes are configured, click **OK**.

Errors with device usage statistics when staging and backup operations are concurrent

When disk devices such as AFTDs perform staging and backup operations concurrently, NetWorker does not accurately display the disk volume's usage total. The inaccurate data can be seen in the Written column when using the **mminfo -mv** report command or in the Used column when viewing volume information in the Media window of the NetWorker Administration application.

Editing a staging policy

To edit a staging policy:

1. In the **Administration** window, click **Configuration**.
2. In the left pane, select **Staging**.
3. In the right pane, select the **Staging** policy to edit.

You cannot edit the name of an existing staging policy.

4. From the **File** menu, select **Properties**.
5. Make any necessary changes and click **OK**.

Copying a staging resource

To copy an Staging resource:

1. In the **Administration** window, click **Configuration**.
2. In the left pane, select **Staging**.
3. In the right pane, select the **Staging** resource to copy.

4. From the **Edit** menu, select **Copy**. The **Create Staging** dialog box appears, containing the same information as the Staging resource that was copied, except for Name attribute.
5. Type the name for the new **Staging** resource in the Name attribute, edit any other attributes as appropriate, and click **OK**.

Deleting a staging policy

The Default staging policy cannot be deleted.

To delete a staging policy:

1. In the **Administration** window, click **Configuration**.
2. In the left pane, select **Staging**.
3. Remove all devices from the **Staging** policy.
 - a. In the right pane, select the **Staging** policy to be deleted.
 - b. From the **File** menu, select **Properties**.
 - c. In the **Devices** attribute, ensure that all listed devices are unselected.
 - d. Click **OK**.
4. In the right pane, select the **Staging** policy to be deleted.
5. From the **File** menu, select **Delete**.
6. When prompted, click **Yes** to confirm the deletion.

Consideration for staging a bootstrap backup

Bootstrap backups can be directed to a disk device such as an AFTD or FTD device. However, if a bootstrap backup is staged to another device, the staging operation will complete and will be reported as complete even though the “recover space” operation will not be executed. This means that the staged bootstrap will remain on the original disk from which it was staged. Therefore, the original disk can be used to scan in the bootstrap data if the staged bootstrap is accidentally deleted. Also be aware that if the bootstrap data is not staged from the original disk, the data on the original disk will be subject to the same browse and retention policies as any other save set backup and will, therefore, be subject to deletion after the retention policy has expired.

Staging and cloning from the command prompt

Staging a save set from the command prompt works differently than staging a save set from the NetWorker Console. When staging from the NetWorker Console, you select save sets that belong to a single device. When staging from the command prompt, specify the save set IDs to be staged.

When a save set is cloned, the cloned save sets are given the same save set ID as the original save set with a new clone ID. When staging a save set from the command prompt, the NetWorker software stages all the save sets with the specified save set ID and then removes those save sets. That means that any cloned versions of the save set are removed when the original is removed.

To ensure that all clones are not removed, specify a clone ID with the save set ID to indicate the source volume of the staging. For example:

```
nsrstage -m -S ssid/cloneid
```

To find the clone ID of a save set, use the **mminfo** command. For example:

```
mminfo -avot -r "volume,ssid,cloneid,name"
```

For information about **nsrstage** or **mminfo**, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

CHAPTER 14

Recovering Filesystem Data

This chapter covers these topics:

- ◆ NetWorker recovery overview 368
- ◆ Overview of NetWorker recovery methods..... 372
- ◆ Recovering the data 373
- ◆ Recovering deduplication data 389
- ◆ Recovering with BMR..... 389
- ◆ Recovering encrypted data 389
- ◆ Recovering the Windows system configuration 390
- ◆ Recovering Windows volume mount points 396
- ◆ Recovering special Windows databases 397
- ◆ Recovering expired save sets 398
- ◆ Recovering client files from an old NetWorker server..... 404
- ◆ Recovering critical NetWorker server databases..... 405
- ◆ Recovering the NMC server database 417

NetWorker recovery overview

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.

The **NetWorker User** and the **NMC Recovery Wizard** programs recover data sequentially. You can use multiple **recover** commands to recover files in parallel.

Three types of NetWorker hosts are involved in a recovery operation:

- ◆ Administering host—The NetWorker host that starts the recovery.
- ◆ Source host—The NetWorker host from which the backup was run.
- ◆ Destination host—The NetWorker host that receives the recover data.

Perform a recovery operation in one of two ways:

- ◆ Local recover—A single NetWorker host is the administering, source and, destination host.
[“Local recoveries” on page 368](#) provides more information.
- ◆ 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 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. For example, if the AUTOEXEC.BAT or .profile file on a client is appropriate for a new client, recover the file to the new client.

[“Directed recoveries” on page 369](#) provides more information.

Local recoveries

In a local recovery, the administering host is also the source and destination host. Local recoveries are the simplest way to recover NetWorker data.

To perform a local recovery, the user on the local host, *user@localhost* must:

- ◆ Belong to a NetWorker User Group that has the Recover Local Data privilege.
[“NetWorker User Groups” on page 559](#) provides more information.
- ◆ Have operating system ownership of the recovered files. The root user on UNIX and the Windows Administrator have this privilege.
- ◆ Have write privileges to the local destination directories. The root user on UNIX and the Windows Administrator have this privilege.

Directed recoveries

A directed recovery enables a NetWorker administrator 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 required to perform the operation.

A user with directed recovery privileges can:

- ◆ Browse the backup data of all NetWorker clients.
- ◆ Recover the data to any NetWorker client.

Figure 33 on page 369 provides an example of a directed recovery. A user on client Saturn performs a directed recovery of data from a remote client to the destination client Mars.

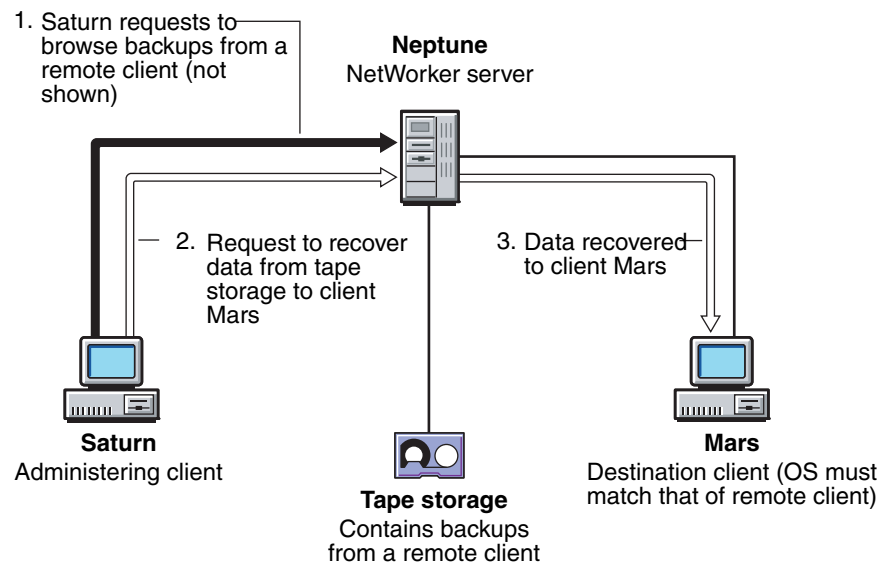


Figure 33 A directed recovery from a remote client

Review the requirements for a directed recovery:

- ◆ “Directed recover requirements” on page 370
- ◆ “Windows requirements” on page 371
- ◆ “UNIX specific requirements” on page 372

Directed recover requirements

[Table 49 on page 370](#) summarizes recover requirements for each host in a recover session.

Table 49 General recover requirements (1 of 2)

Host	Requirements
Destination	<p>Is the same platform as the source host, for example, UNIX to UNIX or Windows to Windows.</p> <p>Uses the same file system as the source host, for example, UXFS to UXFS or NTFS to NTFS.</p> <p>The root user or the Administrator user on the destination host must be one of the following:</p> <ul style="list-style-type: none"> • A member of a NetWorker User Group with Remote Access All Clients privileges. • Added to the Remote Access attribute of the source host. <p>For example: The source client is mars. The destination client venus, is a Windows host. The Remote Access attribute for the client mars contains: <code>Administrator@venus</code></p> <p>The servers file contains an entry for the administering host. “Editing the servers file” on page 625 describes how to modify the servers file.</p> <p>The Disable Directed Recover attribute is set to the default value No, in the NSRLA database. A value of No means that the client accepts directed recoveries from remote hosts. “Editing a client NSRLA database” on page 608 describes how to edit the NSRLA database.</p>
Source	<p>The same platform as the destination host, for example, UNIX to UNIX or Windows to Windows.</p> <p>Uses the same file system as the destination host, for example, UXFS to UXFS or NTFS to NTFS.</p> <p>The Remote Access attribute contains the user account 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: <code>Administrator@venus</code></p>

Table 49 General recover requirements (2 of 2)

Host	Requirements
Administering	<p>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.</p> <p>Use the local root or Administrator account to start the recover. Ensure the user account is a member of one of the following:</p> <ul style="list-style-type: none"> • The Operators, the Application Administrators, the Database Administrators, or the Database Operators User Group. <hr/> <p>Notice: You must have operator privileges in the Operators user group to perform a selective files restore from a Microsoft Windows deduplication backup. Microsoft provides complete documentation for working with the Windows deduplication functionality.</p> <hr/> <ul style="list-style-type: none"> • A customized User Group with the following privileges on the NetWorker server: <ul style="list-style-type: none"> Remote Access All Clients Operate NetWorker Monitor NetWorker Operate Devices and Jukeboxes Backup Local Data Recover Local Data Recover Remote Data <p>“Managing server access” on page 558 provides more information about access requirements</p>
Virtual cluster client	<p>A Client resource for the virtual cluster client must exist on the NetWorker server. The Remote Access attribute for the virtual cluster client must contain an entry for the root or Administrator user for each physical cluster node.</p>

Windows requirements

Review this section before you perform a directed recovery of a Windows client:

- ◆ You cannot recover data from a backup that you performed on a NetWorker 8.1 or higher client to a pre-NetWorker 8.1 client.
- ◆ You cannot perform a directed recovery of the SYSTEM or VSS SYSTEM save sets.
- ◆ You cannot perform a directed recovery to a CIFS share.
- ◆ For Microsoft Windows XP Professional, the NetWorker server and clients must be in a Microsoft Windows Domain, regardless of the username and password.
- ◆ Enable the Windows File and Print Sharing option on a Windows destination client.
- ◆ If you use the **recover** command and the NetWorker server and clients are Microsoft Windows hosts, then:
 - When the NetWorker server and the destination host are in the same domain, start the NetWorker Backup and Recovery 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 NetWorker Backup and Recovery service with a local user that exists on both hosts. The local user must have the same password on both hosts and be a member of the local Administrators group on the NetWorker server.

UNIX specific requirements

When relocating non-ASCII directories 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, based on the locale of the administering client.

Overview of NetWorker recovery methods

Determine the recovery method that you will use:

- ◆ [“Browsable recovery” on page 372](#)
- ◆ [“Save set recovery” on page 372](#)
- ◆ [“Scanner recovery” on page 373](#)
- ◆ [“VSS File Level Recovery \(FLR\)” on page 373](#)

Browsable recovery

The file selection recovery method, or browsable recovery enables you to browse for and select the files and directories to recover. The browse policy applied at the time of backup determines the earliest versions of files and file systems that are available for recovery. [“About browse and retention policies” on page 276](#) provides more information about browse policies.

Use a file selection recovery when:

- ◆ You do not know the exact name of the file or directory to recover.
- ◆ You want to recover a small number of files or directories. When you select many files and directories, the recover process can take some time to complete, particularly from the **NetWorker User** program.
- ◆ You want to perform a directed recovery.
- ◆ You want to recover only the files that you select, not extra files.

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 recover, recover the last full backup first, then recover levels 1 to 9 and incremental backups in the chronological backup order. [“Backup levels” on page 267](#) provides information about the relationship between full backups, backup levels 1 to 9, and incremental backups.

Use a save set recovery when:

- ◆ You want to recover many files or all the data in a save set, for example, in the event of a total disk failure. When you perform a save set recovery, you do not select individual files or directories for recovery.
- ◆ You want to recover data from a recoverable or recyclable save set. [“About browse and retention policies” on page 276](#) provides more information about browse and retention policies. [“Recovering a recyclable or recoverable save set entry in the online indexes” on page 399](#) describes how to repopulate the client file index entries for recoverable and recyclable (expired) save sets.
- ◆ Memory resources on the recovery host is scarce. A save set recovery requires less memory than a browseable recovery.

Scanner recovery

The **scanner** program enables you to recover data directly from a volume

Use the **scanner** program to recover data when:

- ◆ You want to perform a by file selection recovery but the save set information is not in the client file index.
- ◆ You want to recover data directly from a tape.
- ◆ You want to recover data from an incomplete save set.

VSS File Level Recovery (FLR)

VSS File Level Recovery is a new feature in NetWorker 8.1 and 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 un-supported writers (application writers) continue to work and are excluded from file system backups.

Recovering the data

NetWorker provides four applications to recover data. The application that you can use to recover data depends on the administrative host used to recover the data.

To perform recoveries:

- ◆ From the NMC server, use the **NMC Recovery Wizard**. [“Using the Recovery Wizard” on page 374](#) provides more information.
- ◆ From a command prompt, use the **recover** command. [“Using the recover command” on page 379](#) provides more information.

- ◆ From a Windows administrative host, use the **NetWorker User** application. [“Using the NetWorker User program” on page 382](#) provides more information.
- ◆ From the NetWorker server when save set information is not in the media database, use the **scanner** command. [“Using the scanner program” on page 386](#) provides more information.
- ◆ [“Using VSS file level recovery \(FLR\)” on page 388](#) provides information on recovering data by using VSS FLR.

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).
- ◆ 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 501](#) 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 Recovery Wizard” on page 376](#) 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.

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 you use to connect to the Console server has Configure NetWorker privileges. “[NetWorker User Groups](#)” on page 559 provides more information.
- ◆ The appropriate configuration is in place if you will perform a directed recover. “[Directed recoveries](#)” on page 369 provides more information.

Creating a new recover configuration

The Recovery Wizard allows you to create and save a configuration that you can reuse or modify later.

1. Connect to the Console server from a Console client.
2. Connect to the NetWorker server.
3. Click **Configuration** from the left navigation pane, then select **Clients**.
4. Right-click the client from which you want to recover the data, then select **Recover**. The Recovery Wizard appears.
5. Navigate through the Recovery Wizard screens and define the configuration for the recover job. Online help describes how to use the Recovery Wizard.

Modifying a saved recover configuration

The Recovery Wizard allows you to save partial recover configurations and complete the configuration at a later time.

To modify saved recover configurations:

1. Connect to the Console server from a Console client.
2. Connect to the NetWorker server.
3. Click **Recover** on the Administration window toolbar. The **Recover** window appears. “[Recover window](#)” on page 475 provides more information about the **Recover** window.
4. 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.

To copy a configuration for reuse:

1. Connect to the Console server from a Console client. Ensure that the account you use to connect to the console server has Configure NetWorker privileges. [“NetWorker User Groups” on page 559](#) provides more information.
2. Connect to the NetWorker server.
3. Click **Recover** on the Administration window toolbar. The **Recover** window appears. [“Recover window” on page 475](#) provides more information about the **Recover** window.
4. In the **Configured recovers** window, right-click the saved recover configuration, select **Recover Again**.
5. Make changes as required and save the configuration with a new name.

Troubleshooting 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:

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** screen.
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** screen.
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.

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.emc.com -c mnd.emc.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.emc.com is the name of the NetWorker server.
- mnd.emc.com is the name of the source NetWorker client.

3. To 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. Quit the **nsradmin** application.
 - c. Confirm that the **nsrtask** process starts.
 - d. If the **nsrtask** process does not start, the 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.emc.com -c mnd_emc.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 hang, 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 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.

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.

Table 50 on page 380 provides information to review before performing a command line recovery.

Table 50 Command line recovery considerations

When	Consider
Using recover.exe on Windows	<p>To avoid using the Windows version of recover.exe on Windows operating systems, perform one of the following:</p> <ul style="list-style-type: none"> • Include NetWorker_install_path\bin\recover.exe at the command prompt • Ensure that NetWorker_install_path\bin appears before %SystemRoot%\System32 in the \$PATH environment variable. <p>To recover files or directories that begin with a dash (-) such as -Accounting, try one of the following options:</p> <ul style="list-style-type: none"> • Run the recover command and type add ./-Accounting to recover the -Accounting file or directory and its contents. • Run the recover command and use the cd command to change directories to -Accounting. Type add . to add the current directory and its contents for recovery. • When the current directory is /temp and -Accounting resides in the /temp directory, run the recover command and type add /temp/Accounting. This adds -Accounting and the contents of the directory to the recovery list.
Recovering Windows SYSTEM or VSS SYSTEM save sets	You must use the recover command in non-interactive mode.
Recovering NDMP data	<p>When restoring NDMP data, relocate the data to a directory that differs from the original location. The NDMP protocol does not support name conflict resolutions. NetWorker overwrites existing files that have the same name as the recovered file. To relocate the data, run the recover command and then type: relocate destination_directory_name. “Performing NDMP recoveries” on page 680 provides details about NDMP recoveries.</p>
Performing directed recoveries	<p>You must use the recover command in interactive mode. You cannot perform a save set restore. “Directed recoveries” on page 369 provides detailed information about a directed recovery.</p>
Performing save set recoveries	<p>Use user accounts with root (UNIX) or Administrator (Windows) permissions.</p> <p>You cannot perform a directed save set recover.</p> <p>Perform concurrent recoveries from an advanced file type by either using multiple -S options to identify multiple save sets, or starting multiple recover commands.</p>

To recover data by using the **recover** command, use the following syntax:

```
recover -a -s NetWorker_servername -c source_host -S ssid/cloneid -d
destination_directory -R destination_host -i_recover_option
[destination_name]
```

where:

- ◆ **-a** used to perform the recovery in non-interactive mode. You cannot use the -a option with the -R option. You cannot use the -a option to restore Windows SYSTEM or VSS save sets.

- ◆ *-s NetWorker_servername* specifies the name of the source client's NetWorker server.
When you do not specify *-s*, recover attempts to connect to the first host specified in the servers file. The servers file, located in */nsr/res*, contains an entry for each available server.
- ◆ *-c source_host*—Specifies the source host.
When you do not specify *-c*, NetWorker assumes that the source client is the host where you run the **recover** program.
- ◆ *-S ssid/cloneid* used to perform a save set recover. The *ssid/cloneid* specifies the save set to recover. To recover a clone save set, specify the *ssid* and *cloneid*. You cannot use the *-S* option with the *-R* option.
- ◆ *-d destination_directory*—Specifies the full path to the directory on the destination client for the recovered files. Ensure that you use proper syntax for platform of the destination client.
- ◆ *-R destination_host*—Specifies the destination host to receive the recovered data. When you do not use the *-R* option, the host where you run the **recover** program is the destination host. You cannot use the *-R* option with the *-S* option or the *-a* option. Requires the use of the *--i_recover_option* option.
- ◆ *-i_recover_option*—Specifies how NetWorker handles a naming conflict between a recovered file and an existing file. Required when you use the *-R* option.

For example:

- *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.
- ◆ *source_directory*—Specifies the initial directory in which to begin browsing.

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information about the **recover** command and available options.

Example 30 Performing a recover in interactive mode

1. Type: **recover**
2. To select the files or directories to recover:
 - a. Specify the directory to browse, for example:

```
recover> cd /var/adm
```
 - b. Select the file or directory for recovery:

```
recover> add file_name
```
3. To view the selected files or directory:

```
recover> list
```
4. To view the list of the volumes that NetWorker requires to recover the data:

```
recover> volumes
```

- To recover the files to a location that differs from the original location, type:

```
recover> relocate path
```

- 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 `jupiter'
Recover completion time: Tue Jan 21 08:33:04 2009
recover>
```

- To close the **recover** program, type **Quit**.

Example 31 Performing a recover in non-interactive mode:

To recover a directory /testdir on client mars to a new directory /newdir on client mars, type:

```
recover -d /newdir -a /testdir
```

Example 32 Save set recover from command line

To recover a directory /testdir from a save set with a ssid of 12345678 and overwrite any existing files with the same name, type this command:

```
recover -S 12345678 -iY /testdir
```

Example 33 Performing a directed recover

To recover a directory c:\mydir that was backed up on client mars to client jupiter, type:

- On client mars, type:

```
recover -R jupiter -iY
```

- Add the c:\mydir directory:

```
recover> add c:\mydir
```

- Recover the files:

```
recover> recover
```

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 *EMC NetWorker Module for Microsoft Applications Administration Guide* provides more information.

Browsable recover by using NetWorker User

Perform these steps in the **NetWorker User** program on the administering host.

1. Select the NetWorker server when prompted.
2. From the **Operations** menu, select **Recover/Directed**. To perform a save set recover, select **Save Set Recover**.
3. Select the source host that has the data you want to recover, then click **OK**.
4. Select the destination host for the recovered data, then click **OK**.
5. 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.

6. Select optional recover options. [Table 51 on page 383](#) summarizes the available recovery options.

Table 51 Optional browsable recovery options (1 of 2)

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: <ol style="list-style-type: none"> 1. Select View > Change Browse Time. 2. In the Change Browse Time window, select a new day within the calendar. Select Previous Month or Next Month to change from the current month. 3. 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. 4. 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: <ol style="list-style-type: none"> 1. Highlight the file or directory that you want to review. 2. Select View > Versions. 3. Select a previous version. 4. Select Change Browse Time. 5. When prompted to change the browse time, click OK. 6. 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: <ol style="list-style-type: none"> 1. From the File menu, select Find. 2. 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.

Table 51 Optional browsable recovery options (2 of 2)

Recover option	Details
Relocate the recovered file system objects	<p>By default, NetWorker recovers file system objects to their original location. To relocate the files to a different location:</p> <ol style="list-style-type: none"> 1. Select Options > Recover Options. 2. In the Relocate Recovered Data To field, type the path on the destination host to recover the data, then click OK. <hr/> <p>Notice: For NDMP data restores, the target path is a literal 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. “Performing NDMP recoveries” on page 680 provides details about NDMP recoveries.</p>
View volumes required for recovery	<p>Before you start the recovery operation, monitor which volumes NetWorker requires to recover the selected file system objects.</p> <p>To view the required volumes, select View > Required Volumes.</p> <hr/> <p>Notice: Ensure the listed volumes are available or NetWorker to mount into an available device.</p>
Resolve name conflicts	<p>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:</p> <ol style="list-style-type: none"> 1. Select Options > Recover Options. 2. Select a conflict resolution option: <ul style="list-style-type: none"> • 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 reboot the destination host. <hr/> <p>Notice: 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. “Performing NDMP recoveries” on page 680 describes how to perform NDMP recoveries</p>

7. 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 additional information. VSS recovery error messages are also written to the NetWorker log file

Save set recover by using NetWorker User

Use the NetWorker User program to perform a save set recover when the administering host is a Windows system.

NOTICE

Only members of the Windows Administrators group have permission to perform a save set recovery.

Perform these steps in the NetWorker User program, on the administering host.

1. Select the NetWorker server when prompted.
2. Select **Operation > Save Set Recover**.
3. Select the source host that has the data you want to recover, then click **OK**.
4. In the **Save Sets** window, select the name of the save set from the **Save Set Name** list.
5. Select the version of the save set (if there are multiple versions). When required, select the cloned version of a save set if one is listed.
6. Select optional recover options. [Table 52 on page 385](#) summarizes the recover options available with a save set recovery.

Table 52 Optional save set recovery options (1 of 2)

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: <ol style="list-style-type: none"> 1. Click Files... 2. Specify the files and directories to recover, one full path per line. 3. 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 . <p>Notice: Ensure the listed volumes are available for NetWorker to mount into an available device.</p>

Table 52 Optional save set recovery options (2 of 2)

Recover option	Description
Relocate the recovered file system objects	<p>By default, NetWorker recovers file system objects to their original location. To relocate the files to a different location:</p> <ol style="list-style-type: none"> 1. Select Recover Options. 2. In the Relocate Recovered Data To field, type the full path of the directory where the data should be relocated and then click OK. <hr/> <p>Notice: For NDMP data restores, the target path is a literal 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. “Performing NDMP recoveries” on page 680 provides details about NDMP recoveries.</p>
Resolve name conflicts	<p>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:</p> <ol style="list-style-type: none"> 1. Select Options > Recover Options. 2. Select a conflict resolution option: <ul style="list-style-type: none"> • 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 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 reboot the destination host. <hr/> <p>Notice: 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. “Performing NDMP recoveries” on page 680 describes how to perform NDMP recoveries</p>

7. Click **Ok** 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 the following appears:

```
Received 1 file(S) from NSR server server
Recover completion time: Tue Jan 21 08:33:04 2009
```

Using the scanner program

You can use the **scanner** program 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 of the source client.

NOTICE

You cannot use the **scanner** command recover data from a NetWorker Module, NDMP or DSA save set.

To recover data by using the **scanner** program:

1. Ensure the value in the **Idle device timeout** attribute of the device that contains the volume is 0. [“Automatic unmounting of volumes \(idle device timeout\)” on page 231](#) provides more information.
2. 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.

3. 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 -S ssid -f mediafile -r mediarec device | path\uasm -rv  
where:
```

- *ssid* specifies the save set ID value obtained from the **mminfo** output.
- *mediafile* specifies the starting file number of the save set, obtained from the **mminfo** output.
- *mediarec* specifies the starting file record number of the save set, 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** file.

For example, on Windows:

```
C:\Program Files\EMC NetWorker\bin
```

Example 34 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.

Example 35 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
```

Example 36 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 -r mediarec device -x path/uasm -rv -m  
source_dir=dest_dir filename
```

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information about the **scanner** program.

Using VSS file level recovery (FLR)

Currently, the system state files are filtered out using VSS directives created when VSS is initialized and processed. Not applying the VSS directives adds the system state back into the volume backups. The directives for exclude file list and un-supported writers (application writers) generated from VSS are applied.

For any volume backup used in online incremental system state or offline incremental disaster recovery, every file system backup must be VSS consistent. In order to create VSS consistent backups, system state writers that are affected by this backup must participate in the file system backup, even if the backup includes a single file. For example, when backing up C:\dir1\file1, the VSS writers which protect files on C:\ must participate in the snapshot, but only a snapshot of C:\ is created.

This behavior may not be appropriate for certain circumstances. Use the command line flag `VSS:*=off`, to remove this VSS requirement. However, to create a valid backup of any system critical volume, the Windows VSS mechanism must be used.

For incremental backups of any level other than full, time of the last change is used to determine whether a file needs to be included or not. This is not an issue for most system state files. For certain database files such as SQL files used by WID writer, the time of the last change may not reflect the latest updates to the file.

A new option, replace at reboot, is available during the recovery procedure. When this option is selected, file conflicts replace the target file with the contents of the source file. If any of the files are in use, they are replaced at reboot. This option is disabled if the target computer is different from that of the source.

This functionality is available for both browse based restore and BBB backup restores.

File Level Recovery is available from any supported web browser. You select a date and time of a backup to mount, and then select the option for specific files from the selected backup. You may restore a file to the original destination, or to a new destination on the VM. You can also restore files from a backup to a New VM.

VSS FLR restore of system state files using directed recovery renders the target host unstable. Restoring system state data, for example, registry C:\system32\config, or side by side files to another system, for example, C:\windows\winsxs, via directed recovery leaves the target host unstable or not bootable.

Recovering deduplication data

The *NetWorker Avamar Integration Guide* and the *NetWorker Data Domain Deduplication Devices Integration Guide* provides more information on how to recover deduplication data.

Recovering with BMR

[“Windows Bare Metal Recovery” on page 719](#) provides information on performing a Windows BMR recovery with NetWorker.

Recovering ACL files

NetWorker now provides the ability to browse and recover files with associated ACLs (Access Control Lists) in directories for which the user is not the primary owner.

To recover files with associated ACLs, **ACL passthrough** must be checked in the **Recover** section of the **NetWorker Server Properties** window. The feature is enabled by default.

If ACL passthrough is disabled, the following message is displayed when a non-owner attempts to browse ACL files in the directory:

```
Permission denied (has acl)
```

Note: ACLs and extended attributes for files are not recovered when files are recovered to a different operating system file system as can be the case in a directed recovery.

Recovering encrypted data

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, follow the procedure in this section.

The current pass phrase is listed in the **Datazone Pass Phrase** attribute of the NetWorker server. [“Set the Datazone pass phrase for a NetWorker server” on page 108](#) provides more information.

To recover AES encrypted data that was not encrypted with the current pass phrase:

- ◆ Use the `-p` option with the command that is being used to recover data. For example:

```
recover -p pass_phrase
```

```
winworkr -p pass_phrase
```

To enter multiple pass phrases with the `-p` option, type:

```
recover -p pass_phrase1 -p pass_phrase2 -p pass_phrase3
```

- ◆ Specify the pass phrase(s) in the **Pass phrase** field in the **NMC Recovery Wizard**.

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.

Recovering the Windows system configuration

To recover a Windows client operating system configuration, recover the SYSTEM or VSS SYSTEM save sets. Recover all SYSTEM or VSS SYSTEM save sets for a client at the same time to prevent conflicts.

When All is typed in the **Save Set** attribute of the Client resource, these SYSTEM save sets are backed up if VSS is disabled:

- ◆ SYSTEM STATE
- ◆ SYSTEM FILES
- ◆ SYSTEM DB
- ◆ SHAREPOINT (only if installed on the client to be backed up)

NOTICE

Non-VSS save sets are not supported with Microsoft Windows Vista or Windows Server 2008 or later.

These VSS SYSTEM save sets are backed up if VSS is enabled (default setting):

- ◆ VSS SYSTEM BOOT
- ◆ VSS SYSTEM FILESET
- ◆ VSS SYSTEM SERVICES
- ◆ VSS USER DATA (Windows Server 2003 only)
- ◆ VSS OTHER (Windows Server 2003 only)

When VSS is enabled (default setting), at a *minimum*, back up and recover VSS SYSTEM BOOT, VSS SYSTEM FILESET, VSS SYSTEM SERVICES, and all boot/system volumes to properly recover the entire system.

NOTICE

To back up and recover SYSTEM or VSS SYSTEM save sets by using the NetWorker User program or from the command prompt, you must have local Windows Administrator privileges.

After recovery of the SYSTEM STATE, SYSTEM FILES, and SYSTEM DB save sets, restart the operating system. Additionally, restart the system if instructed to do so after recovery of any of the other VSS SYSTEM save sets. This ensures complete recovery of the components.

Do not attempt to restore the SYSTEM STATE or VSS SYSTEM BOOT save set twice in succession without restarting after each operation. When you attempt to restore the VSS SYSTEM BOOT save set a second time without restarting after the first restore, this error message will appear:

```
regsw:cannot replace Registry key. Access is denied
```

When you install a service or driver *after* a backup, then restore the backup, the newly installed service or driver may not be in the state you expect. Reinstall the service or driver, or use Control Panel to reconfigure the startup type.

NOTICE

NetWorker software does not determine the Windows operating system version during recovery of the SYSTEM or VSS SYSTEM save sets. When you attempt to recover SYSTEM or VSS SYSTEM save sets to a different operating system, the system may be inoperable after the recovery.

For example, if you back up the SYSTEM or VSS SYSTEM save sets, then upgrade the Windows software to a new operating system. Do not recover the SYSTEM or VSS SYSTEM save sets that were backed up under the previous operating system.

Consider the following when recovering data on a Windows NetWorker client:

- ◆ [“Temporary disk space” on page 391](#)
- ◆ [“Disable Antivirus For Windows System Drive Recovery” on page 392](#)
- ◆ [“Recovering the Windows SYSTEM from the command prompt” on page 392](#)
- ◆ [“Point-in-time recovery of the SYSTEM and VSS SYSTEM save sets” on page 394](#)
- ◆ [“Point-in-time recovery of Microsoft SQL Server or Exchange Server” on page 395](#)
- ◆ [“Preparing to recover the SYSTEM DB save set” on page 396](#)

Temporary disk space

Restoring the SYSTEM or VSS SYSTEM save sets requires extra disk space on the system drive for temporary files that are created during the recovery. The process might require as much extra space as the total size of the SYSTEM or VSS SYSTEM save sets.

Before restoring the SYSTEM or VSS SYSTEM save sets, run the **mminfo** command from the command prompt, to monitor the size of the SYSTEM or VSS SYSTEM save sets to be restored:

- ◆ For a Windows 2003, 32-bit version NetWorker client, approximately 500 MB of extra disk space is usually sufficient.
- ◆ For a Windows 2003, 64-bit version NetWorker client, approximately 1GB of extra disk space is usually sufficient.
- ◆ For a Windows 2008, 32-bit or 64-bit version NetWorker client, approximately 10 GB of extra disk space is usually sufficient.

The default location for the restored system temporary files is a system drive where the original files reside. For the VSS SYSTEM save sets, the temporary files are placed in the temp directory on the system drive. For information on how to expand the available space in the temp directory by moving it to another large partition, refer to the Microsoft documentation.

When recovering SYSTEM or VSS SYSTEM save sets by using the NetWorker User program, verify that all save sets were recovered successfully. Do so by reviewing the messages in the Recover Status window (or the networkr log file) after the recovery is complete, but before restarting the client host.

Disable Antivirus For Windows System Drive Recovery

During recovery, antivirus programs may not be able to distinguish between a recovery and an attack and may therefore block the recovery of certain files. Prior to Windows system drive recovery, disable the antivirus program's protection properties. Consult the vendor-specific documentation for more information. After recovery, re-enable the protection properties.

Recovering the Windows SYSTEM from the command prompt

Before attempting to recover Windows SYSTEM components by using the commands described in this section, be aware of the following limitations:

- ◆ You cannot recover SYSTEM or VSS SYSTEM save sets by using the recover command in interactive mode. Instead, use the command line recovery procedures described in the following sections.
- ◆ You cannot perform a directed recovery of a VSS SYSTEM save set from the command prompt.
- ◆ A maximum of one SYSTEM or VSS SYSTEM save set can be included in the same **recover** command. To recover multiple SYSTEM or VSS SYSTEM save sets in one operation, use the NetWorker User program.
- ◆ File system directories cannot be specified in the **recover** command.
- ◆ A maximum of one SYSTEM or the VSS SYSTEM save set can be specified in an input file.
- ◆ File system directories cannot be specified in an input file. An input file is specified in a **recover** command with the **-l** option.

Examples of valid command line entries include:

```
recover -iY -s servername -N "VSS SYSTEM BOOT:"
recover -iY -s servername -N "VSS SYSTEM SERVICES:"
```

Examples of invalid command line entries include:

```
recover -iY -s servername -N "SYSTEM DB:" "SYSTEM STATE:"
recover -iY -s servername -N D:\letters "SYSTEM DB:"
```

Recover the SYSTEM save sets from the command prompt

To recover the SYSTEM save sets from the command prompt:

1. Recover the SYSTEM save sets in this order:

- SYSTEM DB
- SYSTEM FILES
- SHAREPOINT
- SYSTEM STATE

The command that is used to recover each save set should look similar to:

```
NetWorker_install_path\bin\recover.exe -iY
[-s NetWorker_server_name] -N "saveset_name"
```

2. Restart the host. The original SYSTEM files is replaced by restored files.

NOTICE

If the recovery process stops responding, terminate the process and perform the recovery operation again.

Recover the VSS SYSTEM save sets from the command prompt

To recover the VSS SYSTEM save sets from the command prompt:

1. Recover the VSS SYSTEM save sets in this order:
 - VSS SYSTEM SERVICES
 - VSS SYSTEM FILESET
 - VSS USER DATA (Windows Server 2003 only)
 - VSS OTHER (Windows Server 2003 only)
 - VSS SYSTEM BOOT

The command that is used to recover each save set should look similar to this:

```
NetWorker_install_path\bin\recover.exe -iY
[-s NetWorker_server_name] -N "saveset_name"
```

2. Restart the host. The original VSS SYSTEM files is replaced by restored files.

NOTICE

When the recovery fails, an error message appears. Use the Windows Event Viewer to examine the event logs for additional information. VSS recovery error messages are also written to the NetWorker log file.

Recover VSS SYSTEM save set components from the command prompt

You can recover individual components (writers) within these VSS SYSTEM save sets:

- ◆ VSS SYSTEM SERVICES
- ◆ VSS USER DATA (Windows Server 2003 only)
- ◆ VSS OTHER (Windows Server 2003 only)

NOTICE

You cannot recover individual components of the VSS SYSTEM save sets by using the **recover** command in interactive mode. Instead, use the procedure described in this section or use the NetWorker User program.

To recover selected components:

1. Type:

```
NetWorker_install_path\bin\recover.exe -iY [-s
NetWorker_server_name] [-t browse_time] -N
"VSS_SYSTEM_SAVESET_NAME:\component_name"
```

2. Place a semicolon (;) between multiple component names.

For example, to recover the Event Log Writer writer and the WMI Writer writer, type:

```
<NetWorker_install_path>\nsr\bin\recover.exe -iY
-s jupiter -N "VSS SYSTEM SERVICES:\Event Log Writer;WMI Writer"
```

NOTICE

Windows Server 2008 and Windows Vista do not have an event log writer. The event logs will not be backed up as part of the VSS system save sets. The event logs are backed up as part of the file system. To back up the event logs, you should perform a regular (non-VSS) back up of the system32\winevt\logs folder.

Point-in-time recovery of the SYSTEM and VSS SYSTEM save sets

To recover the SYSTEM or VSS SYSTEM save sets to a specific point in time from the command prompt, specify the ID of the save set to be restored. To browse a list of valid save set IDs:

1. From the **NetWorker Console Administration** window, click **Media**.
2. In the expanded left pane, select **Volumes**.
3. In the right pane, right-click a volume, then select **Show Save Sets**.
4. In the **SSID** column, note the appropriate save set ID.
5. To restore the system state or system database to a particular point in time, type:

```
NetWorker_install_path\bin\recover.exe -iY [-s
NetWorker_server_name] -S SSID
```

where *SSID* is the save set ID that was noted in [step 4](#).

NOTICE

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 additional information. VSS recovery error messages are also written to the NetWorker log file.

Point-in-time recovery of Microsoft SQL Server or Exchange Server

Use the following EMC NetWorker modules for backup and recovery of Microsoft servers and server applications:

- ◆ EMC NetWorker Module for Microsoft Exchange to back up and recover the Exchange Server.
- ◆ EMC NetWorker Module for Microsoft SQL Server to back up and recover the SQL Server.
- ◆ NetWorker Module for Microsoft Applications to back up and recover Exchange Server, SQL Server, Office Sharepoint Server, and Data Protection Manager Server.

Preparing to recover the Windows SYSTEM STATE save set

[Table 53 on page 395](#) describes components of the SYSTEM STATE save set that require special preparation before being recovered.

Table 53 Preparing to recover the SYSTEM STATE save Set

Component	Recover preparation
Active Directory (if installed)	<ol style="list-style-type: none"> 1. When the host is restarting, “Directory Services Restore Mode” must be specified. 2. On any domain controller that is a DNS server, ensure that the %SystemRoot%\system32\drivers\etc\hosts file includes the name and IP address of the NetWorker server. <p>The NetWorker <i>Procedure Generator</i> provides complete details.</p>
Certificate Server (if installed)	<ol style="list-style-type: none"> 1. Reinstall the Certificate Server after reinstalling the operating system. 2. Specify the same name for the Certificate Server database, and the same paths for the database and log files, as when the system was backed up. 3. Copy the EFS keys. For information about EFS keys. “Encrypting file system” on page 886 provides more information.
Cluster Server (if installed)	<p>Shut down the Cluster Service on any nodes in the cluster on which the service is started, except for the node on which the recover is performed. To shut down the cluster service, performing <i>one</i> of the following:</p> <ul style="list-style-type: none"> • Type the net stop clussvc command at the command prompt. • Use the Microsoft Computer Management program.
COM+ Database	Set the TEMP environment variable to a valid temporary directory.
SYSVOL (if installed)	None
Internet Information Server	None
Performance Counters	None

Preparing to recover the SYSTEM DB save set

[Table 54 on page 396](#) describes components of the SYSTEM DB save set that require special preparation prior to recovery.

Table 54 Preparing to recover the SYSTEM DB save set

Component	Recovery Preparation
Disk Quota Database	<ul style="list-style-type: none"> Subsystem for drive being recovered must be created with same drive letter as the original. Subsystem must be enabled. Drive must be in NTFS format.
Removable Storage	Removable Storage database backup and recovery is not supported.
Terminal Services Licensing (if installed)	Terminal Services Licensing must be started.
WMI	None

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.

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. These procedures do not apply if Automated System Recovery (ASR) is used. [Chapter 24, "Recovery Support for Windows XP and 2003 Automated System Recovery"](#) provides information about support for ASR.

Recovering a mount point and its data

To recover a Windows mount point and its data:

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 382](#) provides more information about performing data recoveries.

Recovering nested mount points

To recover nested mount points and their data:

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 special Windows databases

This section describes how to recover Windows system databases:

- ◆ [“Recovering Windows DHCP and WINS databases” on page 397](#)
- ◆ [“Restoring Windows Content Index Server on Windows” on page 398](#)

Recovering Windows DHCP and WINS databases

The DHCP and WINS databases are automatically included when performing a back up of the save set All. When the save set All is backed up, these procedures are not required.

Recover a DHCP database

To recover the DHCP database:

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.

Recover a WINS database

To recover a WINS database:

1. Use the NetWorker **User** program to recover the backup configured in the WINS backup procedure. [“Backing up Windows DHCP and WINS databases” on page 115](#) provides more information.
2. Use Microsoft WINS administrative tools to restore the **WINS** database.

NOTICE

For detailed instructions about Microsoft WINS administrative tools, refer to the Microsoft documentation.

Restoring Windows Content Index Server on Windows

The Windows Content Index Server (CIS) indexes the full textual contents and property values of files and documents stored on the local host.

The backup and recovery of the CIS occurs as part of the SYSTEM DB save set. The CIS is automatically regenerated upon system restart.

Note: When you delete a nondefault CIS database and then try to restore it, the restored database cannot be active until the registry is restored as part of a SYSTEM STATE save set recovery.

During a CIS restore:

- ◆ When you are using the CIS in a program that provides search capability (for example, a website), the search functionality will not be available.
- ◆ Queries can be issued, but response time might be slow.

After a restore, the CIS automatically updates the catalogs to reflect the current data. Therefore, if it has been a long time since the last backup, it may be more efficient to rebuild the catalog than to restore it. However, if the catalog is very large, restoring it might be faster than rebuilding it.

Note: After a restore, all catalog directories automatically restart, even if they were stopped at the time of the backup.

When a restore of the CIS fails, rebuild the CIS:

1. Right-click **My Computer** and select **Manage** to open the **Computer Management** console.
2. Expand **Services and Applications**.
3. Right-click the catalog to be rebuilt and select **All Tasks>Empty Catalog**.
4. Manually stop the CIS and restore the CIS again.
5. Restart the CIS.

When you restart the service, the CIS re-indexes the entire catalog.

Recovering expired save sets

Each NetWorker client, including the NetWorker server has a client file index. This database contains information about the files that are in a save set. When a save set browse policy expires, it is no longer available for recovery by file selection, that is, it is no longer browsable for recovery. Some applications such as the NetWorker Module for Databases and Applications programs require that the save set is browsable to perform the recovery.

You can make expired save set files browsable for recovery by adding the save set information back into the client file index. The procedure to recover the save set entry into the client file index depends on the state of the save set in the media database.

Use one of the following procedures to recover the save set information back into the client file index:

- ◆ [“Recovering a recyclable or recoverable save set entry in the online indexes” on page 399](#)
- ◆ [“Recovering a save set entry into the client file index and media database” on page 403](#)

Recovering a recyclable or recoverable save set entry in the online indexes

This section describes how to recover save set information into the client file index when the save set is the media database and save set status is recyclable or recoverable.

The `ssflags` attribute identifies the status of a save set. The `mminfo` command displays the `ssflags` attribute:

- ◆ A save set is recoverable when `ssflags` contains a `r`. This save set has exceeded its defined browse policy.
- ◆ A save set is eligible for recycling when `ssflags` contain an `E`. This save set has exceeded its defined retention policy. This is also referred to as an expired save set.

The *NetWorker Command Reference Guide* describes how to use the `mminfo` command.

[“Recovering a save set entry into the client file index and media database” on page 403](#) describes how to add save sets to the media database and client file index if they no longer exist.

Perform the following steps to recover the client file index entries for a recyclable or recoverable save set:

- ◆ [“Task 1: Use mminfo to generate save set information” on page 399](#)
- ◆ [“Task 2: Use nsrmm to modify the save set properties” on page 400](#)
- ◆ [“Task 3: Repopulate the client file index” on page 401](#)

Task 1: Use mminfo to generate save set information

To generate the save set information:

1. Login to the NetWorker server as root or the Windows administrator account.
2. Query the media database on the NetWorker server for the NetWorker client that requires modification:

```
mminfo -avot -c client_name
```

where *client_name* is the name of the recovery client.

3. Record the following values for the save sets to add into the client file index:

- `ssid` column
- date of the backup
- `lvl` column. When the `lvl` value is anything other than *full*, ensure that you record the `ssid` and dates for the previous full backup and all level backups in between.

For example:

```
<NetWorker_install_path>\nsr\bin>mminfo -avot -c swift
Volume          clientdatetime      sizessid          fl lvl name
```

```

snapimagewin1.001 swift11/5/200912:02:18 PM83 KB4294078835cEfull
c:\bkup
snapimagewin1.001 swift11/5/200912:02:23 PM4 KB4277301623crfull
index:swift
snapimagewin1.001 swift11/5/200912:02:25 PM 141 KB4260524409crfull
bootstrap

```

In this procedure, and all of the following examples, the ssid 4294078835 is used for the client swift11.

4. Obtain the cloneid for the recorded save sets:

```
mminfo -q ssid=ssid_number -r cloneid
```

where *ssid_number* is the ssid number provided in the previous **mminfo** command.

For example:

```
mminfo -q ssid=4294078835 -r cloneid
```

```
1257402739
```

When more than one ssid was recorded, repeat this step for all ssids.

Task 2: Use nsrmm to modify the save set properties

Refer to the following section to modify the save set properties.

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 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.

When the status of the save set selected is expired for example **cE**, perform [step 1 on page 400](#) before proceeding to the next step.

NOTICE

When the index save set is not recoverable, the save set will expire when the NetWorker software crosschecks the indexes. For example **nsrim -X**.

4. Record the values in the date and time columns.

Task 3: Repopulate the client file index

Repopulate the client file index with information about files in a save set one of two ways:

- ◆ Repopulate the client file index with information about all save sets for the client up to the date and time specified. “[Repopulate the client file index by using the nsrck program](#)” on page 401 provides more information.
- ◆ Repopulate the client file index with information about files and directories for a specific save set. “[Repopulate the client file index by using the scanner program](#)” on page 402 provides more information.

Repopulate the client file index by using the nsrck program

To repopulate the client file index by using the **nsrck** program:

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 saveset 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 saveset 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**.

Repopulate the client file index by using the scanner program

Use the **scanner** program to restore entries to the client file index. 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 using the **scanner** program:

1. Ensure the **idle device timeout** value of the device containing the volume is **0**. Refer to [“Automatic unmounting of volumes \(idle device timeout\)” on page 231](#) 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.

NOTICE

The `-i` option is not supported for cloud devices.

Recovering a save set entry into the client file index and media database

When a volume contains a save set that does not appear in the media database, the **scanner** command is used to restore save set information into the media database and optionally, the client file index.

To rebuild the save set's entry in the media database and the client file index:

1. Log in as root or Windows Administrator.
2. At the command prompt, run the **scanner** program on the volumes that contain the appropriate file or files:

```
scanner device_name
```

3. Use the output from the **scanner** program to determine:
 - Whether the save set to be rebuilt is on this volume.
 - Whether to reintroduce the contents of this volume into the online indexes.
 - Whether the save set spans multiple volumes.
4. Load the first volume containing the save set information into an available device. Ensure the **Idle Device Timeout** value for the device is **0**. Refer to [“Automatic unmounting of volumes \(idle device timeout\)” on page 231](#) for details.
5. Use the **scanner** command to repopulate the NetWorker databases:

- To repopulate the media database with the save set information:

```
scanner -m -S ssid device_name
```

- To repopulate the media database and client file index with the save set information:

```
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.

6. Recover the data.

Recovering client files from an old NetWorker server

This section describes how to move a NetWorker client to a new NetWorker server without losing the ability to recover the client files that were backed up on the old NetWorker server.

To move a client to a new NetWorker server:

1. Record the **Client ID** attribute of the NetWorker client on the old server.
 - a. From the **Administration** window, click **Configuration**.
 - b. In the left pane, click **Clients**.
 - c. In the right pane, right-click the client to be renamed, then select **Properties**.
 - d. Click the **Globals (1 of 2)** tab.
 - e. Record the **Client ID** attribute listed for the client, then click **Cancel** to close the **Properties** window.

2. On the new NetWorker server, create a new client:

- a. In the **Name** attribute, type a name for the client.

This can be the same name that was used on the old server, but it cannot be the same name as an existing client on the new server. When a client with the same name exists on the new server, use this format for 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
```

- b. Click the **Globals (1 of 2)** tab.
 - c. In the **Client ID** attribute, type the client ID determined in [step 1](#).
 - d. Complete other attributes as necessary, and click **OK**.
3. Ensure the pool resource used when the save set was created exists on the new NetWorker server.
 4. Import the client file index entries by using the **scanner** command:

```
scanner -i -c client_name device_name
```

where *client_name* is the name of the client that was set up on the old NetWorker server.

You can now recover data that was backed up when the NetWorker client was set up on the old NetWorker server.

NOTICE

When **scanner -i** or **scanner -m** is used to import data before the **Client** resource is configured on the new server, the client ID for the imported save sets is maintained in the media database. When a client of the same name already exists on the new server, **scanner** stores the client name in the format described in [step 2](#). You can then create the client based on the client ID by completing that step. However, run the **scanner -i** command again after creating the **Client** resource to import save set information into the client file index.

Recovering critical NetWorker server databases

Protecting a NetWorker server including its critical databases requires careful planning and preparation. The recovery methods described in this section may not work if the NetWorker server is not adequately protected. Information about protecting a NetWorker server is provided in the *NetWorker Server Disaster Recovery and Availability Best Practices Guide*.

The databases that are critical to the recovery of a NetWorker server include the bootstrap and the client file indexes.

A bootstrap includes the:

- ◆ Media database
Contains the volume location of each save set.
- ◆ Resource files
Contains all of the resources, such as NetWorker clients and backup groups, that are defined on the NetWorker server.

The client file indexes include tracking information for each file that belongs to a client's save sets. There is one client file index for each NetWorker client.

Starting in NetWorker 8.1 there is a command line program named **nsrdr** that simplifies the recovery of the NetWorker server's media database, resource files, and client file indexes. Previously, you had to use the **mmrecov** command to recover the media database and resource files, and the **nsrck** command to recover client file indexes. These commands are still available. The *NetWorker Command Reference Guide* and the UNIX man pages contain information about these commands and the **nsrdr** command.

Use the procedures in this section to recover lost or corrupted bootstrap or client file indexes. If your server databases are not corrupted and you only need to restore expired save set entries into the client file index or the media database, you can use the procedures in [“Recovering expired save sets” on page 398](#). Save sets are removed from the client file index when their browse policy time has expired. Save set entries are removed from the media database when their retention policy time expires.

The **nsrdr** command is flexible. You can run the **nsrdr** program in fully interactive mode and respond to questions or you can run the program silently with command line options. You can recover the media database, resource files, and all client file indexes in one operation, or recover just one item by itself. If you are recovering client file indexes, you can also recover the indexes for just one or a small number of NetWorker clients instead of recovering all client file indexes for all clients in one operation.

To help troubleshoot issues with the wizard, messages are logged to the following locations:

- ◆ On UNIX, `/nsr/logs/nsrdr.log`
- ◆ On Windows, `<NetWorker_install_path>\nsr\logs\nsrdr.log`

The following topics are included in this section:

- ◆ [“Prerequisites to recover the NetWorker server databases” on page 406](#)
- ◆ [“Consider your recovery options” on page 408](#)
- ◆ [“Options for running the nsrdr command” on page 414](#)
- ◆ [“Setting nsrdr tuning parameters” on page 416](#)

Prerequisites to recover the NetWorker server databases

Depending on the state of your NetWorker server, you may have to do some preparation before you can recover the bootstrap and client file indexes. There are two main scenarios to consider, this guide covers scenario 1 only:

- ◆ Scenario 1: lost bootstrap or client file indexes.
In this scenario you just need to recover the NetWorker server bootstrap or client file indexes because they have been lost or deleted. The NetWorker server software and hardware is intact but you notice that some bootstrap data such as the media database or NetWorker server resources are missing or incomplete. Additionally, you may notice that some clients are no longer browsable for recovery even though they have not exceeded their browse retention time policies; this indicates missing or incomplete client file indexes.
- ◆ Scenario 2: disaster recovery
In this scenario, the NetWorker server host has suffered some damage, such as a disk or power supply failure, and the base operating system might have been removed or corrupted. In this scenario, the hardware must be replaced and a fresh install of the software is required. The steps in this section are beyond the scope of a disaster recovery. The NPG (NetWorker Procedure Generator) provides disaster recovery steps. Additionally, you should follow the practices described in the *NetWorker Server Disaster Recovery and Availability Best Practices Guide* to reduce the likelihood of encountering a disaster recovery scenario and to maximize the likelihood of successfully recovering from a disaster.

The NPG and all user documentation can be downloaded from EMC Online support. To access the NPG, log on to <https://support.emc.com/> and search for NetWorker Procedure Generator. You must have a service agreement to use this site.

Before recovering lost bootstrap or client file indexes, ensure that the following prerequisites are met.

Is the NetWorker server is installed?

If you need to reinstall the NetWorker server software, refer to the disaster recovery steps in the NPG (NetWorker Procedure Generator). These procedures are located in the NPG by selecting the options titled *Recovering with NetWorker > Disaster Recoveries > Server*.

Is the bootstrap report available?

Bootstrap report information includes the following:

- ◆ Bootstrap SSID (Save Set Identification Number)
- ◆ Volume name containing the bootstrap
- ◆ File-number and record-number of the tape media (if used) where the bootstrap information starts

To obtain the bootstrap:

1. Use one of the following methods to obtain the bootstrap:

- Save Group Completion Report. [“Generating and printing bootstrap reports” on page 259](#) describes how to send the Save Group Completion report to an email address on a regular basis.

- Locate the bootstrap SSID and volume name in the messages log file.

- On UNIX, /nsr/logs/messages

- On Windows, <NetWorker_install_path>\nsr\logs\messages

- If the media database is not lost and the volume list is available then obtain bootstrap information by running this command:

```
mminfo -av -B -s server_name
```

- Let the **nsrdr** command scan the device for the bootstrap information. For existing devices, **nsrdr** will detect the latest bootstrap on a volume that contains the bootstrap information.

If the bootstrap is on a disk volume such as an AFTD volume, and you need to create the corresponding AFTD device to access the volume, special precautions are required to prevent the inadvertent destruction of the bootstrap data. Follow the considerations in [“Is a local device available?” on page 407](#) to ensure that you do not destroy the bootstrap data on the volume when adding the device.

- If you cannot locate the bootstrap volume using any of the previous methods, refer to the disaster recovery procedures in the EMC NetWorker SolVe Desktop (formerly known as the NetWorker Procedure Generator).

Is a local device available?

The NetWorker server requires a local device resource to back up the bootstrap data. You can use the same device resource to recover the bootstrap data. In the unlikely case where you need to add a local device because the original device resource is lost, keep the following considerations in mind:

- ◆ Do not relabel the volume when you create the device. Relabeling a volume with bootstrap backups, or any other backups, will render the data unrecoverable.
- ◆ Additional requirements for disk based devices such as AFTD.

- Do *not* allow the device wizard to label the disk volume. The **Label and Mount** option on the wizard's Device Label and Mount window has this option selected by default. Uncheck the **Label and Mount** option.
- Specify the local path to the AFTD volume in the device wizard Select Storage Node window. Ensure that this is the same path on which the bootstrap data is stored.

Consider your recovery options

The **nsrdr** command is flexible and can be run in a variety of ways. However, the major options to consider before running the **nsrd** command are outlined in this section. For a complete list of advanced options, refer to [“Options for running the nsrdr command” on page 414](#).

Do you need to recover all client file indexes?

Recovering all client file indexes can take a long time. If you only need to recover the client file indexes for a limited set of clients, use the **nsrdr -I** option, for example:

```
nsrdr -c -I clientA clientB clientD
```

[“Options for running the nsrdr command” on page 414](#) provides more options for recovering specific client file indexes with the **nsrdr** command.

Were save sets backed up after the last bootstrap backup?

If save sets were backed up after the last bootstrap backup, then these backup records might be overwritten after the bootstrap is recovered. This situation can only occur when a manual backup is taken. A manual backup does not trigger a bootstrap backup immediately, therefore the manual backup will not be recorded in the bootstrap until the next scheduled backup. To protect against losing save sets that were backed up after the last bootstrap backup, use the **nsrdr -N** or **-N -F** options, for example:

If using a tape volume:

```
nsrdr -N
```

If using a disk based device such as an AFTD:

```
nsrdr -N -F
```

If you know that manual backups were not taken after the last bootstrap backup or you are not concerned about losing these backups, do not use the **-N** or **-N -F** options. These options can increase the time and complexity of the recovery considerably.

Recovering the NetWorker server databases

The steps in this section assume that you are running the NetWorker server disaster recovery command, **nsrdr**, in fully interactive mode. This is the recommended mode to use in a typical NetWorker server recovery operation. You can also run the wizard with various command line options depending on what you need to accomplish. [“Options for running the nsrdr command” on page 414](#) provides more information.

To run the NetWorker server disaster recovery command:

Note: Steps one and two are required only if using the **-N** option with the **nsrdr** command in step three. [“Were save sets backed up after the last bootstrap backup?” on page 408](#) provides more information.

1. Unmount all volumes including tape, file type, advanced file type devices, and cloud volumes.
 - a. In the NetWorker Administration interface, click **Devices**.
 - b. Select **Devices** in the navigation tree. The Devices detail table appears.
 - c. Right-click a device and select **Unmount**.
2. Enable the CDI (Common Device Interface) attribute on all tape devices.

Note: NDMP, AlphaStor, and optical devices do not support CDI.

- a. In the NetWorker Administration interface, click **Devices**.
 - b. From the View menu, select **Diagnostic Mode**.
 - c. Select **Devices** in the navigation tree. The Devices detail table appears.
 - d. Double-click a device in the Devices table.
 - e. Select the **Advanced** tab. In the Device Configuration area, locate the CDI settings and select **SCSI commands**. The EMC NetWorker Administration Guide provides more details about CDI considerations.
 - f. Stop and restart the NetWorker server services/daemons.
3. Log in as root or Administrator and type one of the following commands at a command prompt:

nsrdr

If your backups are on tape and you want to prevent the possibility of overwriting manual backups that were taken after the last bootstrap backup, type the following command:

nsrdr -N

If your backups are on a disk device such as an AFTD (Advanced File Type Device), and you want to prevent the possibility of overwriting manual backups that were taken after the last bootstrap backup, type the following command:

nsrdr -N -F

Using the **-N** or the **-N -F** options set the Scan Needed (scan volflag) flag on ALL appendable (non read-only) volumes that are listed in the recovered bootstrap's media database. When the **-N** option is specified and you attempt to write data to a tape-based device that has newer save sets than what is recorded in the media database, a message displays that explains how to update the media database to avoid the possibility of overwriting the newer data. When the **-N -F** option is used for disk devices such as an AFTD, you can still write to the disk, however, recover space

operations will be suspended until the Scan Needed flag is removed. A recover space operation purges the disk device of any save sets that do not have a corresponding entry in the media database.

If you are sure that no backups were done after the last bootstrap backup or you do not need to recover that data, omit the **-N** or **-N -F** options.

4. When asked to continue, type **Y** for yes.
5. Select the device that contains the NetWorker server bootstrap save set.
6. Type the save set ID of the latest bootstrap.

If you do not know the save set ID of the latest bootstrap, leave this entry blank and press **Enter**.

- a. Select **Yes** when given the option to scan the device for the latest bootstrap save set ID.

Note: The option to scan for a bootstrap save set ID is not supported for non-English locales. In this case, use the **scanner** command to find the bootstrap ID.

- b. When the latest bootstrap save set is located, select **Yes** to recover the bootstrap save set.
- c. If you are recovering from tape, you are given the option to input the tape file number and record location number of the bootstrap save set. This information can speed up the bootstrap recovery. Enter this information if you have it, otherwise press **Enter**.
- d. If you are recovering from tape, you are prompted to load the volume.

The scanner program is run and the bootstrap save set is recovered. The media database is merged with the recovered media database and the recovered resource database is saved to a temporary folder named *res.R*. The NetWorker server services are also shut down because the resource database cannot be overwritten while these services are running.

7. Select **Yes** when asked if you want to replace the resource database folder (*res*) with the recovered resource database.

The NetWorker server services are restarted after the resource database folder is replaced with the recovered resource database. The replaced folder is renamed to *res.timestamp*.

8. If you want to recover all client file indexes, select **Yes** when asked if you want to do a client file index recovery. Select **Yes** again when asked to confirm your choice.

You will recover one client file index for each NetWorker client that was backed up including the client file index for the NetWorker server. The wizard ends after the client file indexes are recovered.

If you want to recover the client file index for selected clients only:

- a. Select **No** and exit the wizard.
- b. Re-enter the **nsrdr** command with the **-c -l** options and provide a list of client names with each name separated by a space, for example:

```
nsrdr -c -I clientA clientB clientD
```

The bootstrap recovery is skipped and you are prompted to complete the recovery of the specified client file indexes.

The wizard exits after the client file indexes are recovered.

9. Open the NetWorker server's Administration window and check that all NetWorker server resources appear as expected. "[NetWorker Management Console interface](#)" on [page 35](#) describes how to open the NetWorker Administration window.
 - a. Click the **Configuration** icon and check that all resources appear as expected prior to recovery.
 - b. Click the **Devices** icon and check that all devices appear as expected prior to recovery.
 - c. Click the **Media** icon and check that all media resources appear as expected prior to recovery.
 - d. Select **Disk Volumes** from the Media screen. Check the volume's *mode* status, which is shown in the window on the right. All disk volumes should have the same mode that existed prior to the recovery. All devices that will be written to should be in the appendable mode.

Note: If **nsrdr** was used with the **-N** or **-N -F** options in [step 3](#) then all recovered devices will be set to the scan needed mode (displayed as Mode = Scan Needed). In this case, complete the following steps in this procedure. Otherwise, skip the following steps and resume regular operations with NetWorker.

10. If the **-N -F** option was used in [step 3](#) all disk volumes that are appendable (non read-only) and that are listed in the recovered bootstrap's media database are set to **Scan Needed**. If you suspect that the disk volumes have savesets that were saved after the last bootstrap backup, you can run the **scanner -i** command to populate the recovered bootstrap and the client file indexes with the missing saveset information.

A manual save operation is the only way a saveset can get backed up without triggering a save of the bootstrap and CFI data. If a manual backup was performed before the next scheduled backup, which always backs up the bootstrap and client file indexes, then the last saved bootstrap and CFI will not have a record of the savesets that were backed up manually.

NOTICE

The **scanner -i** command can take a very long time to complete, especially on a large disk volume. For volumes that you do not suspect have save sets that were backed up after the last bootstrap backup or for volumes where you do not need to keep these manual backups, skip this step and complete [step 11](#) where you will remove the volume's Scan Needed flag.

- For AFTD volumes that you suspect may have savesets that were saved after the last bootstrap backup, enter the following command:

```
scanner -i device_name
```

where *device_name* is the AFTD device name *not* the AFTD volume name.

If you do not know the AFTD device name that corresponds to the AFTD volume, use the **nsrmm** command with the **-C** option, for example:

nsrmm -C

Output similar to the following is displayed:

```
32916:nsrmm: file disk volume_name mounted on device_name,
write enabled
```

where *device_name* is the device that corresponds to the AFTD *volume_name*.

- For cloud volumes, enter the following command:

scanner -i -V *cloud_volume* -Z *datazone_ID* *cloud_device*

where *datazone_ID* is the NetWorker server datazone ID if it is in a different datazone than the cloud device.

11. For AFTD devices, remove the Scan Needed status so that recover space operations are enabled for the device:
 - a. Unmount the AFTD volume:
 - From the NetWorker server’s Administration window, click the **Devices** icon and then click **Devices** in the left panel.
 - Identify the device in the right panel to be unmounted. Note the volume associated with the device.
 - Right click the device and select **Unmount**.
 - Repeat for all devices that require the Scan Needed status to be removed. This should be the status of all devices if **nsrdr** was used with the **-N -F** options.
 - b. Remove the Scan Needed status:
 - From the NetWorker server’s Administration window, click the **Media** icon and then click **Disk Volumes** in the left panel.
 - Identify the volume in the right panel that is associated with the device in the previous step.
 - Right click the volume and select **Mark Scan Needed**.
 - Select **Scan is NOT needed** and click **OK**.
 - Repeat for all volumes that require the Scan Needed status to be removed.
 - c. Mount the AFTD volume
 - From the NetWorker server’s Administration window, click the **Devices** icon and then click **Devices** in the left panel.
 - Identify the device in the right panel to be mounted.
 - Right click the device and select **Mount**.
 - Repeat for all devices that were unmounted.
 - Ensure that all devices are mounted and that the Scan Needed status has been removed for the associated volumes.

12. If the **-N** option was used in [step 3](#) and you attempt to mount a tape volume that has savesets that are newer than what is recorded in the media database, a message similar to the following appears:

```
nw_server nsrd media info: Volume volume_name has save sets unknown
to media database. Last known file number in media database is ###
and last known record number is ###. Volume volume_name must be
scanned; consider scanning from last known file and record numbers.
```

- a. Make a note of the file number and record number that is displayed in the message and then enter the following command to update the media database and thus, avoid a potential loss of data:

```
scanner -f file -r record -i device
```

- b. After the scanner operation completes, remove the Scan Needed flag from the tape volume by using the **nsrmm** command:

```
nsrmm -o notscan volume_name
```

You can now use regular recovery procedures to recover application and user data on the NetWorker server.

NOTICE

If the recovered NetWorker server was protecting virtual cluster clients or a NMM protected virtual DAG Exchange server, the `nsrd.log` file will contain false error messages related to the CFI recovery of the underlying physical hosts. Using a NMM protected virtual DAG Exchange server as an example, you would see messages similar to the following:

```
9348:nsrck: The index recovery for 'EXCH2010-2.vl11.local'
failed.9431:nsrck: can't find index backups for
'EXCH2010-2.vl11.local' on server 'sa-wq.vl11.local'
```

You can ignore error messages related to the physical hosts because NetWorker does not backup the underlying physical host in a virtual environment.

Options for running the nsrdr command

You can run the NetWorker server disaster recovery wizard command (nsrdr) with various command line options instead of running the wizard in fully interactive mode. [Table 55, “Command line options for the nsrdr command,”](#) includes a brief description of the **nsrdr** command line options. For a complete description of the nsrdr command and its options, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

Table 55 Command line options for the nsrdr command (1 of 2)

Option	Description
-a	Runs the command line wizard in non-interactive mode. At a minimum, the -B and -d options must be specified with this command. Notice: Be sure to specify a valid bootstrap ID with the -B option when running this command in non-interactive mode. Otherwise, the wizard will exit as though it was cancelled without providing a descriptive error message.
-B <i>bootstrap_ID</i>	The saveset ID of the bootstrap to be recovered.
-d <i>device_name</i>	The device from which to recover the bootstrap.
-K	Use the original resource files instead of the recovered resource files.
-v	Verbose mode. Generates debugging information.
-q	Quiet mode. Display only error messages.
-c	Recover client file indexes only. If specified with the -a option, you must also specify the -I option.
-I -I <i>client1 client2...</i>	Specify which CFIs (client file indexes) to recover. Each client name must be entered on the command line and separated with a space. If no client names are specified, all client file indexes are recovered. Note: When the -I option is specified, ensure that it is the last option in the command string because any entries after the -I option are interpreted as client names.
-f <i>path/file_name</i>	Specify which CFIs to recover by using an ASCII text file. Place each client name on a separate line in the file. Must be used with the -I option. Ensure that each client name is entered correctly because there is no validation of client names.

Table 55 Command line options for the `nsrdr` command (2 of 2)

Option	Description
<code>-t date/time</code>	Recover CFIs from the specified date or date and time. You must enter a date and optionally, a time, format that is accepted by the <code>nsr_getdate</code> program. The <i>EMC NetWorker Command Reference Guide</i> or the UNIX man pages provide more information about <code>nsr_getdate</code> .
<code>-N</code>	If tapes have save sets that are newer than what is recorded in the recovered bootstrap backup, they will be marked as Scan Needed to prevent the possibility of losing backed up data.
<code>-F</code>	Prevents recover space operations on disk devices such as AFTDs until the Scan Needed flag is removed. A recover space operation purges the disk device of any save sets that do not have a corresponding entry in the media database. Must be used with the <code>-N</code> option.

Examples

The following examples depict some common command line usages with the `nsrdr` command and its options.

- ◆ To recover the bootstrap data and selected client file indexes only:


```
nsrdr -I client1 client2 client3
```

 where each client name is separated with a space.
- ◆ To recover the bootstrap data and selected client file indexes by using an input file:


```
nsrdr -f path\file_name -I
```

 where `file_name` is an ASCII text file with one client name on each line.
- ◆ To skip the bootstrap recovery and recover selected client file indexes by using an input file:


```
nsrdr -c -f path\file_name -I
```

 where `file_name` is an ASCII text file with one client name on each line.
- ◆ To skip the recovery of bootstrap data and recover all client file indexes:


```
nsrdr -c -I
```
- ◆ To skip the recovery of bootstrap data and recover selected client file indexes:


```
nsrdr -c -I client1 client2
```
- ◆ To skip the recovery of bootstrap data and recover selected client file indexes from a specified date:


```
nsrdr -c -t date/time -I client1 client2
```

 where the `date/time` is the date and/or time from which the client file indexes are recovered. The `date/time` format is specified in MM/DD/YYYY format or any date and time accepted by the `nsr_getdate` command. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information about `nsr_getdate`.
- ◆ To run `nsrdr` in non-interactive mode and to recover the bootstrap data and all client file indexes:


```
nsrdr -a -B bootstrap_ID -d device -I
```

Setting nsrdr tuning parameters

You can specify the following tuning parameters for the NetWorker Server Disaster recovery wizard.

- ◆ The path to the NetWorker services, such as **nsrdr**, if the default path was not used during the installation.
 - The default path on Solaris and Linux is `/etc/init.d/networker`, the default path on HPUX is `/sbin/init.d/networker`, and the default path on AIX is `/etc/rc.nsr`.
 - The default path on Windows is `C:\Program Files\EMC NetWorker\nsr\bin`
- ◆ The number of parallel threads that can be spawned when recovering CFIs (client file indexes) for multiple NetWorker clients. The default value is 5, which means that up to five parallel threads are spawned to recover CFIs. If you are recovering a large number of client CFIs, increasing this value can shorten the disaster recovery time.

If you do specify any of these parameters, they must be set up before running the wizard. You set up these parameters by creating an ASCII plain text file named `nsrdr.conf`, entering the parameter values in the file, and placing the file under the debug folder of the NetWorker installation.

To set up wizard tuning parameters:

1. Create a text file and give it the name `nsrdr.conf`.

Note: Some text editors append `.txt` to the end of the file name. If this occurs, remove the `.txt` extension so that the file name is `nsrdr.conf`.

2. To specify a non-default path to the NetWorker services on UNIX or Linux, add the following entry:

```
NSRDR_SERVICES_PATH = /non_default_path/nsr
```

Where *non_default_path* is the path to the NetWorker services. On Windows, the path would look similar to the following:

```
NSRDR_SERVICES_PATH = drive:\non_default_path\EMC NetWorker\nsr\bin
```

3. To specify the number of parallel threads that can be spawned when recovering CFIs for multiple clients, add the following entry:

```
NSRDR_NUM_THREADS = number
```

Where *number* is a value that is greater than 1. If a value of zero (0) or a negative value is entered, the default value of 5 is used instead.

Ensure that a space is added before and after the equals (=) sign. If you specify both tuning parameters, ensure that each value is entered on a separate line.

4. Save the `nsrdr.conf` file as a plain text file and place it in the following directory:

On UNIX or Linux:

```
/nsr/debug/
```

On Windows:

```
NW_install_path\nsr\debug
```

The tuning parameters will take effect the next time the wizard is run.

Recovering the NMC server database

The NMC server database contains management data such as report information. The *NetWorker Procedure Generator* provides information about recovering a NMC server database.

CHAPTER 15

Enterprise reporting and events monitoring

This chapter covers these topics:

- ◆ Enterprise data reporting 420
- ◆ Enterprise events monitoring 456

Enterprise data reporting

To facilitate trend analysis, capacity planning, and problem detection, NetWorker software automatically collects data on a continual basis from the NetWorker enterprise. 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 421](#).

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 422](#) 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 or not 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.

Note: The NetWorker Console is unable to generate reports when deployed in a pure IPv6 environment due to a Sybase iAnywhere 9 limitation.

Enabling/Disabling the gathering of report data

The Gather Reporting Data feature is set by default when a host is added to the enterprise. If the NetWorker server to be monitored is not yet in the enterprise, you can enable the Gather Reporting Data feature when adding the NetWorker server to the enterprise. [“Adding a managed host” on page 544](#) provides details.

To enable or disable the gathering of report data:

1. From the **Console** window, click **Enterprise**.
2. Select the **NetWorker** server for which the collection of report information is to be enabled.

3. Right-click the NetWorker managed application, then select **Properties**. The **Managed Application Properties** dialog box appears.
4. Under **Features**, select **Gather Reporting Data**, then click **OK**.

To disable the gathering of reporting data, clear the **Gather Reporting Data** checkbox, then click **OK**.

Data retention and expiration policies

NetWorker Console provides separate expiration policies for retaining different types of data, in the NetWorker Console database, to meet the needs of the environment as described in [Table 56 on page 421](#). Only a Console Application Administrator can modify these policies.

Table 56 Data retention policies

Retention policy	Type of data to be retained	Default
Completion Data (in Backup Status reports, <i>except</i> 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 (in Backup Status reports, <i>only</i> in the save set output). Retention policy for completion messages can affect multiple reports.	Messages, such as error messages for failed save sets.	Two weeks
Backup statistics (in all Backup Statistics reports). Retention policy for backup statistics data can affect multiple reports.	Backup statistics, such as	One year
Audit Data (in User Audit reports). 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
Recover Statistics Save Set Data in Recover Statistics reports.	Save set records.	One year

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 422](#). These different policies give administrators the flexibility to retain certain types of information for less time than others, as demonstrated in [Example 37 on page 421](#).

Note: Reports not mentioned in [Table 56 on page 421](#) have no retention policies.

Example 37 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 time period. But it might be useful to save the completion data for a somewhat longer period to help with load balancing and trends.

The longest time 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. Once data is purged because of the retention policy, it *cannot* be retrieved except by recovering the full database.

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.

Setting expiration policies for data retention

Note: Only a Console Application Administrator can perform this procedure.

To set expiration policies:

1. From the **Console** 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 Console database to hold the data. If the data retention policy settings cause the Console database to run out of storage space, it stops running. The *NetWorker Installation Guide* provides information about estimating the size of the Console database.

Report categories

[Table 57 on page 423](#) describes the various report categories included in NetWorker software. Each of these categories is discussed in detail in [“Preconfigured reports” on page 435](#).

Report categories appear as folders within the Reports window. These reports can be run either from the Console window or from the command prompt.

Table 57 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 backed up. “NetWorker backup statistics reports” on page 435 provides more information.
NetWorker Backup Status	Provide status information about savegroup completion and save set backups. “NetWorker backup status reports” on page 437 provides more information.
Clones	Provides the history of automatic and scheduled clone operations performed on NetWorker servers on version 7.6 Service Pack 2 and later. “NetWorker clone reports” on page 439 provides more information.
Events	Provide summary and detailed information about NetWorker events. “Event reports” on page 445 provides more information.
Hosts	Provide a listing of NetWorker servers in the Enterprise, including information about event and reporting features. “Host reports” on page 447 provides more information.
Users	Provide lists of defined NetWorker Console users, logout and login reports, audit reports, and users with restricted views. “User reports” on page 447 provides more information.
Devices	Provide information about the way devices are being used. “Device reports” on page 447 provides more information.
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. “Inactive files” on page 450 provides more information.
NetWorker Recover	Provide the history of recovery operations that have been performed by NetWorker servers. “NetWorker recovery reports” on page 441 provides detailed information about the reports in this category.
Avamar Statistics	Provide deduplication backup statistics for each selected NetWorker client. “Avamar Statistics reports” on page 444 provides more information.
Cloud Backup and Recover	Provide information on the Cloud usage for scheduled backups and recovers that are performed by the NetWorker server to and from the Cloud storage device. “Cloud backup and recover reports” on page 448 provides more information.
Data Domain Statistics	Provides deduplication backup statistics for each selected NetWorker client. <i>NetWorker Data Domain Deduplication Devices Integration Guide</i> provides more information.
NetWorker Data Protection Policy	Provides details and summaries for Data Protection Policies. The <i>EMC NetWorker and VMware Integration Guide</i> provides more information.

Report 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 424 and “[Drill-down reports](#)” on page 424 provide detailed information.

Different icons represent the different types of reports:

Table 58 Report icons

Icon	Description
	Basic report
	Shared basic report
	Drill-down report
	Shared drill-down report

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.
- ◆ A Priority Summary of Events provides a report in an attribute-centric manner.

Select the basic report that best provides the information you need.

Drill-down reports

Drill-down reports are preset sequences of basic reports, and can be saved as customized reports in shared mode.

Move up and down through a sequence to compare the information provided by the different focal points. For example, from the NetWorker Backup Status category, it is possible to select the Group Status by a Server drill-down report. This report starts at the server level, then drills down to display a summary report for each of the following:

- ◆ A selected group
- ◆ A selected monthly summary
- ◆ A selected daily summary

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 currently displayed. Also note that drill-down reports cannot be run from the command prompt.

Customized reports

A report that is included with NetWorker software is known as a [“canned report”](#) 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, and then saved and rerun whenever necessary, without having to be configured again. This 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 451](#) and [“Sharing reports” on page 452](#) provide more information.

Configuring reports

Each type of report includes its own configuration parameters that act as filters limiting the data used to build the report output. By default, these parameters are set to include all the information available in the report, which means that the filters are turned off to begin with.

For example, the NetWorker Backup Statistics Server Summary report includes these configuration parameters:

- ◆ Server name
- ◆ Backup type
- ◆ Backup level
- ◆ Save time

In this example, accepting the default configuration of selected parameters results in a report that includes backup statistics 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 the data available.

Note: For Drive Utilization reports, the time range cannot exceed eight days. [“Date and time formats” on page 426](#) provide more information about this limitation, or for details on how to set this range.

The scope of a report can be limited by filtering out one or more parameter options, for example:

- ◆ To exclude certain servers in the enterprise from the report, remove selected server names from the Server Name Selected box.
- ◆ To select only full backups, remove the other backup types from the Server Name Selected box.
- ◆ To include only the statistics for the past month, specify that time range. Time ranges are localized. The input format follows the format specified in locale settings of the operating system.

When a parameter is removed from the Server Name Selected box, it goes into the Server Name Available box. To include that parameter again, click **Add** (➤).

How to configure a report

To configure a report:

1. From the **Console** window, click **Reports**.
2. Expand a report category folder, then select an available report type.
 - When a report type has been selected, the **Configure** tab for that report appears in.
 - The possible parameters for that report appear by default in the **Selected** boxes.
3. To limit the scope of the report, click any of the parameters in the **Selected** box, then click **Remove** (◀).
 - To remove all of the parameters from the **Selected** box, click **Remove All** (◀◀).
 - Removed parameters appear in the **Available** boxes.
4. To return:
 - A single parameter to the **Selected** box, select it from the **Available** box and click **Add** (➤).
 - All available parameters to the **Selected** box, click **Add All** (➤➤).
5. 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 Console server window, clear the Java Cache on the Console client, then generate the report again. The *NetWorker Installation Guide* describes how to clear the Java Cache.

Date and time formats

If a report includes a date-and-time-range parameter, specify the beginning and end date and time in the To and From text boxes. Clicking the arrow of a time input field displays a calendar and clock selector, including adjustment arrows for setting values.

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.

The option of displaying times in 12- or 24-hour formats is determined by the Regional and Language Settings on the system.

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 [Table 59 on page 427](#).

Table 59 Date and time input formats for common locales

Language	Date formats	Time formats
US English	<ul style="list-style-type: none"> • 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) 	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • DD MMMM YYYY 08 March 2009 • DD-MMM-YYYY (08-Mar-2009) • DD/MM/YY (08/03/07) 	<ul style="list-style-type: none"> • HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
French	<ul style="list-style-type: none"> • 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) 	<ul style="list-style-type: none"> • HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
German	<ul style="list-style-type: none"> • 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) 	<ul style="list-style-type: none"> • HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
Japanese	<ul style="list-style-type: none"> • YYYY/MM/DD (2009/03/08) • YY/MM/DD (07/03/08) 	<ul style="list-style-type: none"> • HH:mm:ss z (23:27:30 JST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
Simplified Chinese	<ul style="list-style-type: none"> • YYYY-M-D (2009-3-8) • YY-M-D (07-03-8) 	<ul style="list-style-type: none"> • HH:mm:ss (23:27:30)

In [Table 59 on page 427](#), note that:

- ◆ 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.

Viewing reports

Before displaying a report, select the scope of the report by specifying Configuration parameters. [“How to configure a report” on page 426](#) provides information on configuring reports.

In addition, reports can be printed or exported to various file formats. [“Printing reports” on page 455](#) and [“Exporting reports” on page 453](#) provide information about printing and exporting.

The administrative user can restrict a user’s view of the enterprise to certain servers, affecting the user’s view and scope of his or her reporting.

View reports

To view reports:

1. From **Console** window, click **Reports**.
2. Select a report type.
3. To limit the scope of the report, specify parameters on the **Configuration** tab.
4. Click the **View Report** tab. Most reports display initially in interactive mode and table format.
5. To modify the current view, right-click on the **View Report** tab and select the appropriate view option.

Reports can be displayed in either interactive or document mode. Depending on the report, you may also select to display the content of the report as a table or as a chart. NetWorker supports the following report modes:

- ◆ [“Interactive mode” on page 428](#)
- ◆ [“Document mode” on page 430](#)
- ◆ [“Interactive and document mode chart types” on page 431](#)

In addition, NetWorker includes these restrictions and processing considerations:

- ◆ [“Restricting report views” on page 434](#)
- ◆ [“Background processing of reports” on page 435](#)

Interactive mode

Interactive mode displays a report with dynamic components. The effect of the dynamic components depends on whether a report is viewed as a table or as a chart.

Interactive mode allows access to drill-down reports. Drill-down reports conveniently group related reports to make it easier to view increasing levels of granularity in report data.

Interactive Mode also offers a set of chart selection choices. These choices limit the data in a report by including or excluding certain parameters. Examples of chart-selection parameters include:

- ◆ Duration
- ◆ Save set size
- ◆ Number of files
- ◆ Amount of data
- ◆ Number of save sets

Not all parameters apply to each chart type.

Interactive mode table view

In the table view of the interactive mode, you can:

- ◆ Scroll through rows of the table.
- ◆ Sort, rearrange, or resize columns in the table.
- ◆ Choose which columns to display, and the order in which to display them.

Note: In interactive mode, tables can be sorted just as they can be sorted within other Console windows.

Interactive mode chart view

When a chart is displayed in interactive mode, you can:

- ◆ Switch back and forth between different chart formats by selecting a format from the Chart Type list.

A simplified list of chart formats is provided in [Table 60 on page 432](#).

Note: When viewing a Drive Utilization report as a chart, it automatically displays as a Gantt chart. The chart type cannot be changed.

- ◆ Change selections by using the Data Selector, where applicable.

The Data Selector is available in select reports, and includes control-column information that works in conjunction with a graph of numerical data. While the Data Selector is useful in table format, it can also be used to display interesting and useful data groupings in chart format.

For example, in a Group Summary by Server report 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.

- ◆ Limit the set of X and Y axes in the report by clearing one or more options from the Chart Selector checkboxes. 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 a static report that resembles the view in Print Preview as shown by a PDF file viewer. Within Document mode, these options are available:

- ◆ Orientation (portrait or landscape)
- ◆ Table or chart format
- ◆ Size (zoom level)

Note: In Document mode, for any chart type that displays X - Y axes, two graphs are displayed. If the top graph contains excessive Y -axis data, the data displayed in both graphs could be truncated.

In Document mode, the columns of a tabular report cannot be sorted, rearranged, or resized. In addition, you cannot choose which columns to display, and the order in which to display them. Likewise, the chart format cannot be modified 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 offers a set of parameters for chart selection that limit the data that is displayed, a report in Document mode displays all of 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 options for paging through the output:

- ◆ First
- ◆ Previous
- ◆ Next
- ◆ Last

Document mode table view

Document-mode reports displayed as a table contain several columns of information:

- ◆ One or more *control* columns represent report information that *cannot* be summed as quantitative data (for example, Server name, Save set name, Backup type, and so on). The control columns topics are generally shown as *X*-axis data in charts.
- ◆ One or more *data* columns represent report information that can be summed as quantitative data (for example, Amount of data, number of files, number of save sets, and duration). The data columns topics are generally shown as *Y*-axis data in charts.

The bottom line of each report gives subtotals and totals of all the columns of quantitative data shown in the report.

For example, a report on Save Set Details by Client:

- ◆ Lists each client.
- ◆ Provides:
 - 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.

This makes it easy to parse the data, visually, on a per-client basis, on a save set-per-client basis, and for all clients in the report.

Interactive and document mode chart types

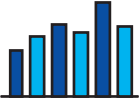
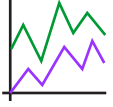
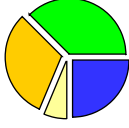
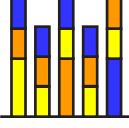

These chart types are available in both interactive and document mode:

- ◆ Bar chart
- ◆ Pie chart
- ◆ Plot chart
- ◆ Stacking bar chart
- ◆ Gantt chart (for Drive Utilization reports only -- more information is provided in the section [“Types of Device reports and configuration” on page 448](#)).

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.

Table 60 on page 432 shows a simplified version of chart format options.

Table 60 Report chart formats

Format	Example
Bar	
Plot	
Pie	
Stacking Bar	
Gantt	

Bar chart

A bar chart 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 backed up by each server. The additional lines show the corresponding numbers of files and save sets backed up by each server.

The set of axes displayed in the report depends on the type of report.

To select various elements for display, select or clear the appropriate checkboxes in the Chart Selector.

Plot chart

Plot charts display data graphed as points along X and Y axes.

To select various elements for display, select or clear the appropriate checkboxes in the Chart Selector.

Pie chart

Pie charts 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 chart

The stacking bar charts are most appropriate for reports where the data is grouped and measured according to more than one category. 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*.

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 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 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 Console window, 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 this chart type is specified 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 shown.

[Example 38 on page 433](#) describes appropriate use of the stacking bar chart type.

Example 38 Appropriate usage of the stacking bar chart

To appreciate the different ways in which a stacking bar chart may be used, consider these reports:

- ◆ A NetWorker Backup Statistics Group Summary by Server shows statistics broken down by savegroup for each server. Different blocks of color are used for the amounts of data backed up by each savegroup within the vertical bars that represent the amount of data that are backed up by servers.
- ◆ A 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. There would, however, 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.

Gantt chart

When viewing a Drive Utilization report as a chart, it automatically displays as a Gantt chart, and the chart type cannot be changed. The Drive Utilization report is the only report that displays 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 this information:

- ◆ Drive
- ◆ Start time
- ◆ End time
- ◆ Throughput value

Chart axis selection

Document mode can display more than one chart in the document. Any or all available *Y* axes can be inserted into the report. When a user changes to document mode, prints or exports a report, or saves a configuration, the axis selection currently set in the Chart Selector section of the Configuration tab is used.

The exceptions to this are stacked bar and pie charts, which display all axes when the **gstclreport** command is used to generate a report.

Restricting report views

When a NetWorker Console user is added or reconfigured, the user's views of NetWorker servers, groups, and clients within the enterprise determines the content of reports that he or she can produce.

Since each user can have different access restrictions, the view of each report can potentially be different. This applies to all report types, whether customized, private, or shared.

For example, a shared backup summary report entitled "Building C Backups" will show different data for different users if the users' access permissions include different NetWorker servers. This is so even if the reports are run at the same time.

In the Reports function, report parameters for a specific user display only the allowed NetWorker servers, groups, and clients as sources of report information. The resulting reports contain data only from those servers. A user may only run reports for servers he or she is allowed to manage.

Note: If no data is available for a given server, that server will not appear in any lists, regardless of the user's view or access.

Background processing of reports

When the **View Report** tab is selected, the report data is removed from the server. This process happens in the background and could take awhile. Other portions of the interface are usable while the report data is being processed. The requested report appears upon returning to the **View** tab.

NOTICE

Do not request multiple reports at the same time. Since reports are run sequentially in the background, a user can navigate around in the user interface while a report is running. If a new report is started before an earlier report is finished, then the earlier report is terminated and deleted. A report is either complete or deleted, the results are never partial.

Preconfigured reports

To facilitate the dissemination of information, NetWorker software includes a variety of reports:

- ◆ [“NetWorker backup statistics reports” on page 435](#)
- ◆ [“NetWorker backup status reports” on page 437](#)
- ◆ [“NetWorker clone reports” on page 439](#)
- ◆ [“NetWorker recovery reports” on page 441](#)
- ◆ [“Data Domain statistics reports” on page 443](#)
- ◆ [“Avamar Statistics reports” on page 444](#)
- ◆ [“Event reports” on page 445](#)
- ◆ [“Host reports” on page 447](#)
- ◆ [“User reports” on page 447](#)
- ◆ [“Device reports” on page 447](#)
- ◆ [“Cloud backup and recover reports” on page 448](#)
- ◆ [“Inactive files” on page 450](#)

NetWorker backup statistics reports

The different types of reports 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 backed up.
- ◆ Number of files backed up.
- ◆ Number of save sets 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 [Table 61 on page 436](#). The specific parameters available depend on which NetWorker Backup Statistics report is selected.

Table 61 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
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. Note: 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 6 months is specified, NetWorker software cannot query the database for a time range that extends back more than 6 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 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. [“How to configure a report” on page 426](#) provides information on selecting and removing parameters.

Except for the Avamar Backup Summary report, these basic reports do not distinguish between regular and deduplication clients.

Backup statistics drill-down reports

The drill-down reports consist of multiple NetWorker Backup Statistics basic reports, connected in a predetermined sequence. [“Drill-down reports” on page 424](#) 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. [“How to configure a report” on page 426](#) 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 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

These backup status reports cover both regular and deduplication clients.

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 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 the Module for Oracle and 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 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 [Table 62 on page 438](#).

Table 62 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.	<ul style="list-style-type: none"> • Full • Incremental • Skip • Level 1-9 (Partial list of options)
Status	Selects status.	<ul style="list-style-type: none"> • 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 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. [“How to configure a report” on page 426](#) provides information on selecting and removing parameters.

Backup status drill-down reports

The drill-down reports are comprised of multiple NetWorker Backup Status basic reports, connected in a predetermined sequence. [“Drill-down reports” on page 424](#) 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. [“How to configure a report” on page 426](#) provides information on selecting and removing parameters.

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.

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 [Table 63 on page 440](#). The specific parameters available depend on which clone report is selected.

Table 63 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 used for cloning.	Selected clone resource
Save Set	Cloned saveset name.	Selected save set names
Level	Backup level of the clone.	<ul style="list-style-type: none"> • Full • Incremental • Skip • Level 1-9 (Partial list of options)
Status	Final completion status of the clone.	<ul style="list-style-type: none"> • Successful • Failed • No save sets found
Type	Type of clone operation.	<ul style="list-style-type: none"> • Scheduled • Manual
Start / End Time	Limits the report to a specified time range. The default range is one day for save set details reports. Note: 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 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. [“How to configure a report” on page 426](#) 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, connected in a predetermined sequence. [“Drill-down reports” on page 424](#) 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 will be 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 on 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 your viewing mode:

- ◆ When in Table mode, double-click on 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.

NetWorker recovery reports

The Recovery reports, available from the Reports task pane in the Console window, allow you to view the history of recovery operations that have been performed by NetWorker servers for any server version 7.3 and later. Additionally, NMC checks for new recovery operations and stores the recover statistics in the Console database every 12 hours and every time a scheduled savegroup backup completes.

Note: Since NMC gathers reporting data from pre-7.6 servers, and allows for recovery jobs for pre-7.6 clients, there may be missing fields in the recover statistics for these jobs. However, NMC still populates the Console database with information about those jobs and generates the reports, leaving any missing fields empty.

Reports can be viewed in both Chart and Table modes, with the Table mode set as the default mode. Four different types of recover reports can be generated:

- ◆ Server Summary
- ◆ Client Summary
- ◆ Recover Details
- ◆ Recover Summary Over Time

Be aware that recovery reports may not be up-to-date because recover job history is gathered by the console server every 12 hours and on completion of every scheduled backup.

Types of NetWorker recovery reports and configuration

The NetWorker recovery report category includes basic and drill-down reports. The different types of reports 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 [Table 64 on page 442](#). The specific parameters available depend on which NetWorker Recovery Statistics report is selected.

Table 64 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 initiated the recover.	
User	Name of the user who initiated the recover.	Selected usernames
Size	The size of the recover	
Number of files	For file system recoveries, the number of files in the recover.	
Start time/End time	Limits the report to a specified time range. Note: 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	<ul style="list-style-type: none"> • 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 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. [“How to configure a report” on page 426](#) provides information on selecting and removing parameters.

Recovery Statistics drill-down report

This drill-down report consists of multiple NetWorker Recovery Statistics basic reports, connected in a predetermined sequence. [“Drill-down reports” on page 424](#) 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. [“How to configure a report” on page 426](#) 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 of time.

To generate the Recover Summary Over Time report, you must first specify the same parameters as those in the Server Summary report, which will be the first report displayed in the sequence.

To drill down to the client level, perform one of the following, depending on your viewing mode:

- ◆ When in Table mode, double-click on 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 your viewing mode:

- ◆ When in Table mode, double-click on 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 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 1 year.

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 *EMC NetWorker Data Domain Deduplication Devices Integration Guide* provides more information.

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 *EMC NetWorker and VMware Integration Guide* provides more information.

Avamar Statistics reports

The NetWorker Avamar Statistics reports, available from the **Reports** task pane in the Console window, provide deduplication backup statistics for each selected NetWorker client.

Reports can be viewed in Table mode. There are four different types of reports that can be generated from the deduplication statistics:

- ◆ Client Summary
- ◆ Save Set Summary
- ◆ Save Set Details
- ◆ Backup Summary

Types of Avamar Statistics reports and configuration

The Avamar 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 Deduplication Statistics report category are described in [Table 65 on page 444](#). The specific parameters available depend on which NetWorker Deduplication Statistics report is selected.

Table 65 Avamar Statistics parameters

Parameter	Description	Options
Server Name	Selects one or more servers	Selected server names
Client Name	Selects one or more clients.	Selected client names
Group Name	Selects one or more groups	Selected group names
Save Set Name	Selects one or more save sets.	Selected save set names
Save Time	Limits the report to a specified time range. The default range is one day for save set details reports. Note: The date/time format available depends on the language locale of the operating system.	save time (range)
Backup Level	Select one or more backup levels.	<ul style="list-style-type: none"> • Full • Incremental • Skip • Level 1-9 (Partial list of options)

The parameters available for each report type in the Avamar Statistics report category are listed in the user interface.

Avamar Statistics basic reports

Within the Avamar Statistics report category, choose any of the basic reports listed in the user interface. Once a report is chosen, the Configure tab displays boxes with lists of the selected parameters for that report. To exclude unwanted parameters from the report, delete them from the list. [“How to configure a report” on page 426](#) provides information on selecting and removing parameters.

Avamar Statistics drill-down reports

The drill-down report, Backup Summary, consists of multiple NetWorker deduplication Statistics basic reports, connected in a predetermined sequence. [“Drill-down reports” on page 424](#) provides general information about drill-down reports.

The configuration parameters for the 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 Client Deduplication Summary report, the configuration parameters are the same as they would be for the basic report, Client Deduplication.

When a report is chosen, the Configure 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. [“How to configure a report” on page 426](#) provides information on selecting and removing parameters.

Event reports

These reports provide summary information about current events on NetWorker, Avamar, 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 (and Avamar server system events), the Event reports provide additional features. The reports enable you to organize, export, and print the event data.

The *NetWorker Avamar Integration Guide* provides information on Avamar.

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 [Table 66 on page 446](#). 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 66 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 Avamar
Priority	Selects only priority events. Priority represents the relative severity of the event. Table 71 on page 458 provides descriptions of the priorities.	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 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. [“How to configure a report” on page 426](#) provides information about selecting and removing parameters.

Additional information

These resources provide more information about the contents of Event reports:

- ◆ [“Working with notes” on page 458](#) provides information on NetWorker Console notes.
- ◆ The *EMC NetWorker Error Message Guide* provides descriptions of NetWorker software error messages and troubleshooting procedures.
- ◆ [Chapter 28, “Troubleshooting”](#) provides information on troubleshooting NetWorker software issues.

Event drill-down reports

The drill-down reports consist of multiple Event basic reports, connected in a predetermined sequence. [“Drill-down reports” on page 424](#) 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. [“How to configure a report” on page 426](#) 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 [Table 67 on page 447](#).

Table 67 Host reports

Report name	Purpose	Configuration parameters	Default
Host List	Provides an overview of servers in the enterprise, including: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Start from Enterprise folder. • Start from selected folder. 	Enterprise Path	Start from Enterprise folder

[“Enterprise” on page 542](#) provides a description of the Enterprise and its folders.

User reports

The Users report category provides information on NetWorker Console user activity. [Chapter 17, “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.

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 of time, 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. 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 included 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.

Cloud backup and recover reports

Cloud backup and recover reports display information on the Cloud usage for scheduled backups and recovers that are performed by the NetWorker server to and from the Cloud storage device.

Types of Cloud backup and recover reports and configuration

The Cloud backup and recover reports category includes basic and drill-down reports. [“Drill-down reports” on page 424](#) provides general information about drill-down reports.

The Configure tab allows you to limit the scope of the report that was selected. The parameters available within the Cloud backup and recover report are described in [Table 68 on page 449](#). The specific parameters available depends on which Cloud backup and recover report is selected.

Table 68 Cloud backup and recover parameters

Parameter	Description	Options
Server Name	Selects managed hosts within the enterprise.	Selected server names
Start Time	Limits the report to a specified time range. The default range is one day for the Backup Details report. Note: The date/time format available depends on the language locale of the operating system.	start time (range)
Device Name	Selects the devices used for backup and recover.	Selected device names

Cloud backup and recover reports

Within the Cloud backup and recover report category, choose any of the basic reports 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. [“How to configure a report” on page 426](#) provides information on selecting and removing parameters.

Cloud backup and recover reports can include this information:

- ◆ NetWorker server — Name of the server.
- ◆ Device name — Name of the device used for backup or recover.
- ◆ Device type — Type of the Cloud storage device, for example, Atmos.
- ◆ Login Account — Cloud username used for logging in.
- ◆ Size — Backup or Recover size.
- ◆ Compression ratio — Ratio of the bytes of information written to or read from the Cloud to the total size of the backup or recover.
- ◆ Bytes transferred — Total number of bytes written to or read from the Cloud.
- ◆ Start time — Start time for the backup or recover.
- ◆ End time — End time for the backup or recover.
- ◆ Save Set Name — Displayed only for backup.
- ◆ User name — Name of the user who initiated the recover. Displayed only for recover.

- ◆ Client name — Displays the name of the client that was backed up. In case of recover, source client name is displayed.
- ◆ Status — Displays the status of backup or recover. For example: succeeded, failed, and so on.

The Device Backup Summary and Device Recover Summary reports can be viewed in both Chart and Table modes. The other reports can be viewed in Table mode. [“Interactive and document mode chart types” on page 431](#) provides general information on charts.

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.

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.

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.

A customized report can be rerun exactly 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:

- ◆ Delete
- ◆ Rename
- ◆ Save
- ◆ Save As...
- ◆ Share

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).

Note: For NetWorker reporting purposes, the terms *customized* report and *saved* report are synonymous.

Customized reports appear alphabetically in the report hierarchy below the canned report from which they were created. They are stored in the Console database, which means that users can access them from wherever they are logged in to the NetWorker Console. This also makes them accessible by the command line reporting feature. [“Command line reporting” on page 454](#) provides more information about command line reporting.

These types of information are stored in customized reports:

- ◆ All options from the report’s Configure tab
- ◆ Column display preferences for tables
- ◆ Orientation (portrait or landscape)
- ◆ Current view type (table or chart)

If the view type is Chart, then the current chart type (bar, pie, plot, or stacked bar) is also saved. For charts, the current chart axis selection is also saved. [“Chart axis selection” on page 434](#) 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 reports

By default, customized reports are stored as private for each user. This means that if a user saves a report, it appears only in that user’s report hierarchy. A report’s owner or the Console Application Administrator may, however, enable it for sharing.

Only the original owner of a customized report or the Console Application Administrator may select:

- ◆ **Delete**, to delete the report.
- ◆ **Rename**, to rename the report.
- ◆ **Save**, to resave the report.
- ◆ **Share**, to add sharing to, or remove sharing from, the report.
- ◆ If the Console Application Administrator removes sharing, the report becomes private again to the original owner, the report’s creator.

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.
- ◆ Chose the **Hide Other Users' Reports** option to toggle the view of reports between only those owned by the user (both private and shared) and all shared custom reports.

If a user copies a sharable report with the **Save As** command, that user becomes the owner of the new report, which is initially set as not shared.

Sharing a report

To enable sharing of a customized report:

1. From the **Console** window, click **Reports**.
2. Expand the report folder that contains the customized report 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 .

Once a report has been enabled for sharing, all users can see it in the report hierarchy.

Note: The Share option is a toggle. To disable sharing, right-click the shared report and select Share.

Exporting reports

Reports can be converted into other file formats and shared with others. [Table 69 on page 453](#) lists the file formats available when exporting reports.

Table 69 Report export formats

Format	Purpose
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 into other programs (such as spreadsheets) that accept the comma separated values (CSV) format. Does not show data totals. Use for raw data only.

Exporting a report

You can choose to export a report as a file in a different format (for example, HTML, PDF, CSV, or PostScript). To export a report to a different file format:

1. From the **Console** window, click **Reports**.
2. Expand the report folder that contains the report to export, then click the report.
3. Click the **View Report** tab to display the report.

4. Right-click the **View Report** tab, select **Export**, then select a format.

Note: To sort or rearrange table columns in a report, export the report to CSV format. The columns then can be sorted or rearranged in a spreadsheet program.

5. In the **Save** dialog box, specify the filename and file location and then click **Save**.

Exporting non-ASCII characters

Due to a limitation in the embedded reporting tool, reports that contain multibyte characters cannot be exported to PostScript or PDF formats. Such characters are replaced by a “?” character.

To obtain a printed version of such a report, print directly from the Console window, or export to HTML format.

Note: ISO8859-1 characters can be exported to PostScript or PDF formats.

Command line reporting

Command line reporting offers these features:

- ◆ Allows reports to be run offline, either as needed or by using scheduling software that makes reports available at predetermined times.
- ◆ Makes use of 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. Additionally, 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 *EMC NetWorker Command Reference Guide* or the UNIX man pages provide a complete description of the command and its options.

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.

Additionally, you must add an environment variable named JAVA_HOME to your 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.

Printing reports

All reports can be printed. This allows for the sharing of report data with users who are unable to view it online.

To print a report:

1. From the **Console** window, click **Reports**.
2. Expand the report folder that contains the report to print, then click the report.
3. Click the **View Report** tab to display the report.
4. Right-click anywhere on the **View Report** tab, then select **Print**.
5. From the **Print** dialog box, select the appropriate options on each tab, then click **Print**.

The **-x print** option in the **gstclreport** command is also available.

Enterprise events monitoring

The NetWorker Management Console (NMC) provides the ability to view details of current NetWorker, Avamar, and Data Domain systems. Information that can be monitored includes activities and operations related to devices and libraries, and events that require user intervention. NMC makes the administration of servers more efficient by providing a centralized means of monitoring activity throughout an enterprise. [“Managing various servers in the enterprise” on page 544](#) provides details on adding hosts to be monitored.

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 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 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.

Polling for System Events

From the System Options dialog box, you can set the poll interval for events and activities generated at system-level for the following:

- ◆ Events and reporting (in seconds)
- ◆ NetWorker activities (in seconds)
- ◆ Data Domain events (in seconds)
- ◆ NetWorker libraries (in hours)
- ◆ Avamar events (in hours)

Note: Event polling for NetWorker libraries and Avamar events can only be done to a maximum of once per hour. [“Setting system options” on page 531](#) provides information on setting polling intervals.

Enabling or disabling the Capture Events option

The Capture Events option must be enabled for a given server before NetWorker software can monitor that server for events. This option is selected by default when a host is added.

To disable or reen able the **Capture Events** option:

1. From the **Console** window, click **Enterprise**.
2. Select the host for which the capturing of events is to be disabled or enabled.
3. Right-click the appropriate application, then select **Properties**.
4. Complete one of these steps as required:
 - To enable captured events, select **Features** > **Capture Events**.
 - To disable captured events, select **Features**, clear the **Capture Events** checkbox.

For Avamar servers, the **Capture Events** option monitors only system-level events. The Avamar documentation provides other event information.

5. If the host is a Data Domain system, select the **Configure SNMP Monitoring** tab.
 - a. Enter **public** in the SNMP community string field.
 - b. Enter the value of the SNMP process port. The default port is **162**.
 - c. Select the **SNMP Traps** (Data Domain system events) to be monitored by NetWorker.
6. Click **OK**.

Viewing events

To view events, from the **Console** window, click **Events**. If any events exist, they are displayed in the Console window.

The **Console** window includes columns that provide specific types of information about each event. [Table 70 on page 457](#) describes the various columns and the information they provide for NetWorker events.

Table 70 Events columns (1 of 2)

Column	Description
Priority	Represents the relative severity of the problem by displaying one of seven icons. Table 71 on page 458 describes each priority.
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, Avamar, and Data Domain.








Table 70 Events columns (2 of 2)

Column	Description
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 savegroup fails at 11:00 A.M. in New York, a Console server in Los Angeles reports the event for the savegroup as occurring at 8:00 A.M. Note: The time format presented depends on the current locale setting. “Date and time formats” on page 426 provides more information.
Category	Classifies the source of the problem.
Message	Displays the text of the error message that generated the event.
Annotation	Displays an icon when an annotation has been made. An annotation is a log associated with an event. “Working with annotations” on page 459 provides more information.
Note	Provides an editable field for making brief notes associated with an event. “Working with notes” on page 458 provides more information.

Event priorities

Each event is designated one of seven possible priorities. [Table 71 on page 458](#) lists the event priorities and the information they provide. When the Console window sorts events by priority, it lists the events in alphabetical order, with *Emergency* between *Critical* and *Information*.

Table 71 Event priorities

Icon	Priority	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 may cause NetWorker software to fail unless corrected immediately. This icon represents the <i>highest</i> priority.
	Information	Information about the current state of the server. This icon represents the <i>lowest</i> 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.

Working with notes

The **Note** column of the **Console** window provides a place to record brief administrative information about an event. For example, you can enter:

- ◆ Name of the NetWorker administrator or operator assigned to the event.
- ◆ Letters or numbers that allow the sorting of events into a preferred order.

A note can contain up to 30 characters, and can be edited or deleted.

Adding a note

To add a note for an event:

1. From the **Console** window, click **Events**.
2. Double-click the cell of the **Note** column corresponding to the appropriate event, then type the text of the note in the cell.
3. After entering the text, click outside the cell.

Editing a note

To edit a note for an event:

1. From the **Console** window, click **Events**.
2. Double-click the note to edit, then change the text as appropriate.
3. After editing the note, click outside the cell.

Deleting a note

To delete a note from an event:

1. From the **Console** window, click **Events**.
2. Double-click the note, highlight the text in the cell, then press **Delete**.
3. After deleting the note, click outside the cell.

Working with annotations

The Annotation column provides a place to record comments associated with an event, and can accommodate more information than the Note column. Each annotation can be up to 12 KB in size. For example, use annotations to log steps taken to resolve an event.

When an annotation has been added to an event, an icon appears in the Annotation column of the Events window. Multiple annotations can be added to a single event, and unlike notes, they cannot be edited or deleted.

Viewing annotations

To view an annotation:

1. From the **Console** window, click **Events**.
2. Right-click the event with the annotation to be viewed, then select **Annotation**. Annotations are listed in descending order, with the most recently added annotation at the top of the list.
3. After viewing the annotation, click **Cancel** to close the dialog box.

Adding an annotation

To add an annotation:

1. From the **Console** window, click **Events**.
2. Right-click the event to be annotated, then select **Annotation**. The **Event Annotation** dialog box appears.

3. Type the text of the annotation.
4. To clear the text just entered, click **Reset**.
5. Click **OK**.

Dismissing an event

After an event has been viewed and acted on, it can be dismissed from the Console window. This helps prevent other users from acting unnecessarily on events that have already been resolved.

Note: Dismissing an event makes it disappear from the Console window for all NetWorker users.

To dismiss an event:

1. From the **Console** window, click **Events**.
2. Right-click the event to dismiss, then select **Dismiss**.
3. Click **Yes** to confirm the dismissal.

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.

System events from an Avamar server (deduplication node) are *not* automatically dismissed in the Console window when the problem that triggered the event is solved. These events must be manually dismissed in the Console window.

CHAPTER 16

NetWorker server events reporting and monitoring

This chapter covers these topics:

- ◆ [Monitoring NetWorker server activities](#) 462
- ◆ [Notifications](#) 479
- ◆ [Reporting group status and backup job status](#) 493
- ◆ [Reporting recover job status](#) 500

Monitoring NetWorker server activities

Monitor the activities of an individual NetWorker server by using the **NetWorker Administration** application.

To access the **NetWorker Administration** application to monitor a NetWorker server:

1. From the **NMC Console** window, click **Enterprise**.
2. In the **EMC NetWorker Management Console Enterprise** view, select the appropriate NetWorker server.
3. Highlight the host's Managed Application, then right-click and select **Launch Application**.

The **Administration** window appears.

Figure 34 on page 462 shows how to select the NetWorker managed application.

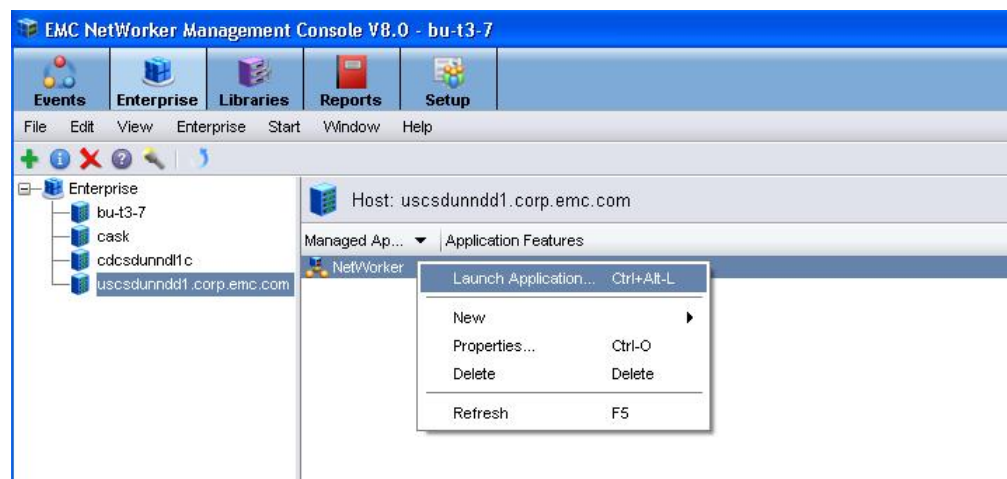


Figure 34 Selecting a NetWorker Managed Application

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

- ◆ Automatic and manual savegroups
- ◆ Archiving, cloning, recovering, synthetic backups, and browsing of client file indexes
- ◆ Alerts and log messages, and operations 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 group backup, as well as, starting and monitoring save set clones.

Figure 35 on page 463 shows the **Monitoring** window.

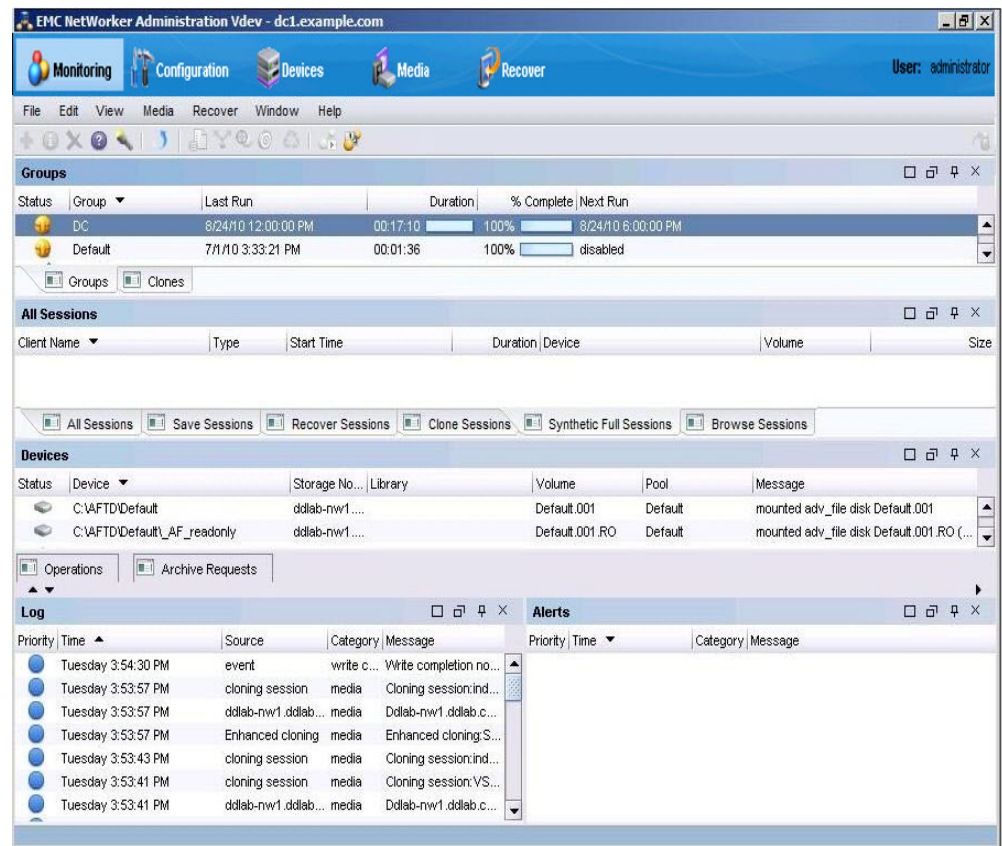


Figure 35 Monitoring window

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, 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 top-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 your view. You can select multiple types from the panel to create multiple floating windows that can be viewed at the same time. [Table 72 on page 464](#) describes the various types of information available in the docking panel and the details each one provides.

Table 72 Monitoring window panel

Window	Information provided
Groups	Lists all groups related to the server, the backup status, the time the last backup was run, the duration of the backup, the completion percentage, and the next time the backup will run. “Groups window” on page 465 provides more information.
Clones	Lists all scheduled clone jobs with the last start time, the last end time, and additional details on the included save sets. “Clones window” on page 467 provides more information.
Sessions	Allows you to customize whether to display all session types, or only certain session types. The information 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 468 provides more information.
Alerts	Lists the priority, category, time, and message of any alerts. “Alerts window” on page 469 provides more information.
Devices	Lists devices, device status, storage nodes, libraries, volumes, pools, and related messages. “Devices window” on page 470 provides more information.
Operations	<p>Lists the status of all library and silo operations, including nsrjb operations 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.</p> <hr/> <p>Notice: When displaying Show Details from the Operations window, the length of time that the window is displayed is controlled by the value entered 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.</p> <hr/>
Log	Lists messages 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 472 provides more information.
Archive Requests	Lists the status of all Archive Requests configured on the server, including the last time the data was archived, the date and time of the next scheduled archive, and any annotations. “Archive Requests window” on page 473 provides more information.

Groups window

The Groups window displays groups that are in the process of completing, or have completed, their backup.

You can use this window to:

- ◆ Identify which groups backed up successfully
- ◆ Identify which groups failed,
- ◆ Start, stop, or restart group backups.

The backup of a client group may fail for one of the following reasons:








- ◆ The NetWorker server failed.
- ◆ The NetWorker client failed.
- ◆ The network connection failed.

To find out more about a backup failure, check Group Backup details. [“Viewing group backup details” on page 466](#) provides more information.

Groups window backup status

The backup status is represented by an icon. [Table 73 on page 465](#) lists and describes each of the icons.

Table 73 Groups window icons

Icon	Label	Description
	Being cloned	The group backup is being cloned.
	Failed	The group backup failed.
	Interrupted	The group backup was interrupted.
	Never ran	The group backup never ran.
	Running	The group backup is running.
	Successful	The group backup successfully completed.
	Probing	The group is in a probing state.

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

Consider the following when a group 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 back to its state prior to 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.

Groups backup operations

This section describes how to use the Monitoring window to start, stop, and restart group backups.

Starting a group immediately

You can override the scheduled backup start time and start the group manually. This is equivalent to selecting Start Now in the Autostart attribute of the Group resource.

When a group backup is started manually, the NetWorker server runs the backup at the level of the next scheduled backup, such as full, levels 1 through 9, incremental, or consolidated.

To manually start a group backup:

1. From the **Administration** window, click **Monitoring**.
2. Click **Groups** in the docking panel.
3. Right-click the group to start, then select **Start**.
4. Click **Yes** to confirm the start.

The NetWorker server immediately backs up the clients in the group.

Stopping a group

To stop a group backup:

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

Restarting groups

To restart a group backup:

1. From the **Administration** window, click **Monitoring**.
2. Click **Groups** in the docking panel.
3. Right-click the group to restart, then select **Restart Group**.
4. Click **Yes** to confirm the restart.

The backup (including synthetic full backup) continues from the point at which it was stopped.

Viewing group backup details

To view detailed information about a group backup:

1. From the **Administration** window, click **Monitoring**.

2. Click **Groups** in the docking panel.
3. Right-click the group to view, then select **Show Details**. The **Group Backup Details** dialog box appears.
4. View detailed information related to the group backups. If any messages were generated, the **Show Messages** button is enabled. Click **Show Messages** to view the messages.
5. Click **OK** to close the **Group Backup Details** dialog box.

Clones window

The **Clones** window displays the scheduled clone jobs and their completion status. This window also identifies which client save sets are cloned successfully and which save sets are not cloned successfully. You can also use this window to start a scheduled clone job immediately.

Scheduled clone operations

This section describes how to use the Monitoring window to start a scheduled clone operation and how to view the clone details for a client's save set.

Starting a scheduled clone immediately

You can start a scheduled clone job at any time instead of waiting for the scheduled start time.

To start a scheduled clone job immediately:

1. From the **Administration** window, click **Monitoring**.
2. Click **Clones** in the docking panel.
3. Right-click the scheduled clone to start, then select **Start**.
4. Click **Yes** to confirm the start.

The NetWorker server immediately starts the scheduled clone job.

Viewing the save sets for a scheduled clone

You can view the NetWorker clients and their save sets that are included in a schedule clone job. You can also determine which save sets were cloned successfully and which ones were not.

To view the clients and save sets for a scheduled clone job:

1. From the **Administration** window, click **Monitoring**.
2. Click **Clones** in the docking panel.
3. Right-click the scheduled clone to view, then select **Show Details**. The **Clone Details** dialog box appears.
4. Click **OK** to close the **Clone Details** dialog box.

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:

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

The default setting for the **Sessions** window is to display save sessions. [“Changing displayed session types” on page 468](#) provides instructions on viewing other session types.

Changing displayed session types

To change the type of sessions displayed on the Sessions window:

1. From the **Administration** window, click **Monitoring**.
2. Click **Sessions** in the docking panel.
3. 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**.

The column headings displayed on this window will differ depending on the type of sessions you chose to display.

Stopping a session

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 **savegrp**.

To stop a session, right-click the session in the window and select **Stop** from the drop-down.

The following table provides a list of actions that can be stopped from NMC.

Table 74 Sessions that can be stopped from NMC

Session type	Stop from NMC?
Save by Save Group	Yes
Synthetic Full by Save Group	Yes
Clone by Save Group	Yes
Schedule Clone	Yes
Manual Save	No
Manual Clone via NMC	No
Manual Clone via CLI	No

Table 74 Sessions that can be stopped from NMC

Session type	Stop from NMC?
Winworker and CLI Recovery	No
Recovery started from Recover wizard	Yes
VMware Backup Appliance Save and Recover	No

NOTICE








When a session is cancelled from NMC, this does not impact any other group operations running.

Alerts window

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

The alert priority is represented by an icon. [Table 75 on page 469](#) lists and describes each of the icons.

Table 75 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 <i>highest</i> priority.
	Information	Information about the current state of the server. This icon represents the <i>lowest</i> priority.
	Notification	Important information.
	Warning	The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.
	Warning	Non-fatal error has occurred.

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

Removing alerts

Individual messages can be deleted from the **Alerts** and **Events** tables by removing them from the **Events** table. The two views show the same messages. To delete a message in the **Events** table, right-click the message and select **Dismiss**.

Devices window







The Devices window 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.

If the current server configuration includes a shared device, a Shared Device Name column appears on the Devices window. 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. Additionally, since only a single device per hardware ID can be active at any given moment, the information for inactive shared devices is filtered out, so only one device per hardware ID is presented on the window at any time.

If the current server uses an AlphaStor library, then a Logical Name column is added to the Devices window to accommodate logical devices.

The device status is represented by an icon. [Table 76 on page 470](#) lists and describes each of the icons.

Table 76 Devices window 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 items on the Devices window are sorted by the Status column, they are sorted in alphabetical order based on the label of the icon.

Operations window

The Operations window displays information about device operations. It includes this information:

- ◆ Status of the operation.
- ◆ Name of the library.
- ◆ Whether or not 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. [“Supply user input” on page 472](#) 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 will appear 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. [Table 77 on page 471](#) lists and describes each of the icons.

Table 77 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.

View operation details

To view detailed information about an operation:

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

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

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.

Stop an operation

Certain operations can be stopped from the Operations window. To stop an operation:

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.

NOTICE

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.

Supply user input

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

To supply input:

1. Right-click the operation, then select **Supply Input**.
2. Confirm whether or not to supply input.
 - If **Yes**, and input is supplied, the icon in the User Input column disappears.
If two users attempt to respond to the same user input prompt, the input of the first user will take precedence, and the second user will receive 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.

NOTICE

If a particular log file is no longer available, check the log file on the NetWorker server. The log files are located in this directory: <NetWorker_install_path>\logs. [“Viewing log files” on page 803](#) provides information about viewing log files.

The **log priority** is represented by an icon. [Table 78 on page 473](#) lists and describes each of the icons.

Table 78 Log 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 <i>highest</i> priority.
	Information	Information about the current state of the server. This icon represents the <i>lowest</i> 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 items on the Log window are sorted by the Priority column, they are sorted in alphabetical order based on the label of the icon.

Archive Requests window

The Archive Requests window displays the current status of all archive requests that are scheduled on the NetWorker server. Use this window to identify which archive requests are running, completed, or failed, as well as when they were last run, and when they are scheduled to run next.

The archive status is represented by an icon. [Table 79 on page 473](#) lists and describes each of the icons.

Table 79 Archive requests window icons

Icon	Label	Description
	Disabled	The scheduled archive is disabled.
	Failed	The archive failed.
	Running	The archive is running.
	Scheduled	The archive is scheduled to run.
	Successful	The archive completed successfully.

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

Viewing details of an archive operation

From the Monitoring window, you can view the details of an archive request, including the start time, the most recent completion time, and other information such as the pool and clone pool to which the archive request will write its data.

To view details of an archive operation:

1. From the **Administration** window, click **Monitoring**.
2. Click **Archive Requests** in the docking panel.
3. Right-click the appropriate archive request, then select **Show Details**.

The Archive Request Details dialog box opens, providing information about the completion of the archive request. The Completion Time displays the time the archive finished. The time it took to complete the archive is the difference between the completion and start times of the archive.

To save archive request information to a file, click Save in the Archive Request Details dialog box. When prompted, identify a name and location for the file.

Archive request operations

Use the Monitoring window to perform a number of archive request operations, such as canceling manual clone jobs, or starting, stopping, restarting, and disabling archive requests. The Monitoring window can also be used to schedule archive requests to start at a specific time in the future. These operations are equivalent to changing the Status attribute of the Archive Request resource, described in [“Scheduling data archives” on page 331](#).

Starting an archive immediately

You can start an archive immediately from within the Monitoring window. This will override and disable any scheduled archive for the selected archive request.

To start an archive immediately from the Monitoring window:

1. From the **Administration** window, click **Monitoring**.
2. Click **Archive Request** in the docking panel.
3. Right-click the appropriate archive request, then select **Start**.
4. Click **Yes** to confirm the start.

Stopping an archive in progress

To stop an archive in progress:

1. From the **Administration** window, click **Monitoring**.
2. Click **Archive Request** in the docking panel.
3. Right-click the appropriate archive request, then select **Stop**.
4. Click **Yes** to confirm the stop.

Scheduling an archive to start automatically

You can also use the Monitoring window to schedule an archive to start automatically at a later time.

To schedule an archive to start automatically at a later time:

1. From the **Administration** window, click **Monitoring**.
2. Click **Archive Request** in the docking panel.
3. Right-click the appropriate archive request and select **Schedule Archive**.
4. In the **Schedule Archive Request** dialog box, type the time that the archive should start, by using the *hh:mm* format.
5. Click **OK**. The **Next Run** column on the **Archive Requests** window displays the entered time.

Disabling a scheduled archive

If an archive request has a scheduled start time, you can disable the scheduled archiving.

To disable a scheduled archive:

1. From the **Administration** window, click **Monitoring**.
2. Click **Archive Requests** in the docking panel.
3. Right-click the appropriate archive request, then select **Disable Archive**.
4. Click **Yes** to confirm the disable.

Recover window

The Recover window displays information about recover configurations created with the NMC Recovery Wizard.

You can use this window to:

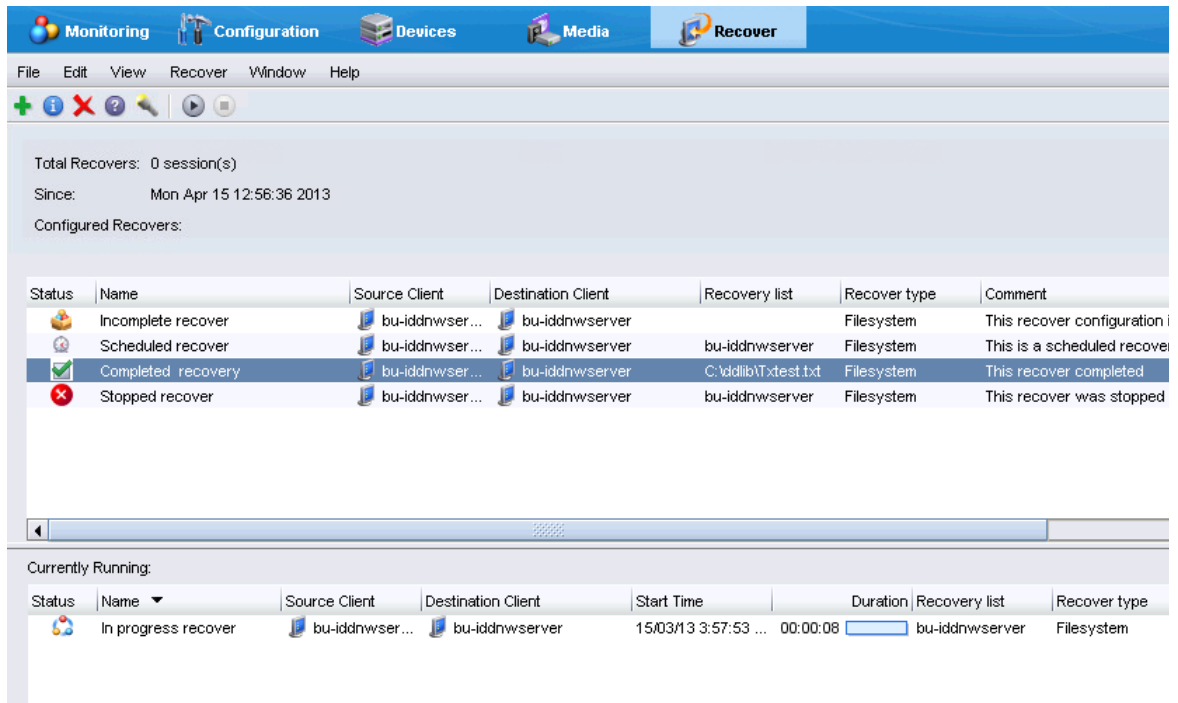
- ◆ Start the NMC Recovery Wizard to create new recover configurations or modify saved recover configurations.
- ◆ Identify the status of a recover configuration created with the NMC Recovery Wizard.
- ◆ Start and stop a recover job.

The Recover window is divided into five sections:

- ◆ Toolbar
- ◆ Summary
- ◆ Configured Recovers
- ◆ Currently Running

A splitter separates the Configured Recovers section from Currently running area. You can click and move the splitter to resize these two windows. [on page 476](#) shows an example of the Recover window.

Figure 36 Recovery Window



Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. [Table 80 on page 476](#) summarizes the function of each toolbar button.

Table 80 Recovery toolbar options (1 of 2)








Button	Function
	Starts the NMC Recover wizard to create a new recover configuration.
	Displays the properties window for the saved recover configuration selected in the Configured Recover window.
	Displays the Find window at the bottom of the Recover window.
	Deletes the saved recover configuration selected in the Configured Recover window.

Table 80 Recovery toolbar options (2 of 2)

Button	Function
	Displays online help for the Recover 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 the in progress recover operation that you selected in the Currently Running window.

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—[Table 81 on page 478](#) summarizes the job status of a saved recover configuration.
- ◆ Name
- ◆ Source client
- ◆ Destination client
- ◆ Recovery list
- ◆ Recover type—For example, filesystem 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 81 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, filesystem 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. [Table 82 on page 479](#) summarizes the available search options.

Table 82

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.

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 executed 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.

The following sections provides information to manage event notifications:

- ◆ [“Preconfigured notifications” on page 480](#)
- ◆ [“Customizing notifications” on page 484](#)
- ◆ [“Logging event notifications” on page 490](#)
- ◆ [“Creating a custom notification” on page 490](#)
- ◆ [“Editing a notification” on page 491](#)
- ◆ [“Copying a notification” on page 491](#)
- ◆ [“Deleting a custom notification” on page 491](#)
- ◆ [“Owner notifications” on page 492](#)
- ◆ [“Savegroup completion and failure notifications” on page 493](#)

Preconfigured notifications

NetWorker is preconfigured to provide most of the event notifications required to monitor NetWorker events. [Table 83 on page 480](#) lists these preconfigured notifications and the associated actions performed by the NetWorker server.

Table 83 Preconfigured notifications (1 of 5)

Notification	Default action
Bootstrap	<p>Windows: Provides the syntax for the smtpmail command to send an email to the administrators with the results of the bootstrap backup. The action attribute must be modified to replace mailserver with the actual hostname of the mail server. “Using smtpmail to email notifications” on page 488 describes how to customize the smtpmail command.</p> <p>UNIX: Sends an email to the root account with the results of the bootstrap backup.</p>
Bootstrap backup failure	<p>Windows: Provides the syntax for the smtpmail command to send an email to the administrator account stating that the bootstrap backup has failed. The action attribute must be modified to replace mailserver with the actual hostname of the mail server. “Using smtpmail to email notifications” on page 488 describes how to customize the smtpmail command.</p> <p>UNIX: Sends an email to the root account with the results of a failed bootstrap backup.</p>
Bus/Device Reset	<p>Windows: Provides the syntax for the smtpmail 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 mailserver with the actual hostname of the mail server. “Using smtpmail to email notifications” on page 488 describes how to customize the smtpmail command.</p> <p>UNIX: Sends an email to the root account stating that a bus or device reset has been detected.</p>
Cleaning cartridge expired	<p>Windows: Reports to the <i>NetWorker_install_path</i>\nsr\logs\media.log file that a cleaning cartridge has expired.</p> <p>UNIX: Sends an email to the root account stating that an expired cleaning cartridge has been detected.</p>
Cleaning cartridge required	<p>Windows: Reports to the <i>NetWorker_install_path</i>\nsr\logs\media.log file that a device cleaning is required.</p> <p>UNIX: Sends an email to the root account stating that a cleaning cartridge is required.</p>
Client install	<p>Windows: Reports the hostname and NetWorker client software version information to the <i>NetWorker_install_path</i>\nsr\logs\media.log file.</p> <p>UNIX: Sends an email to root account: host <i>host_name</i> installed <i>product_version</i>.</p> <p>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.</p>

Table 83 Preconfigured notifications (2 of 5)

Notification	Default action
Device cleaned	Windows: Reports that a device has been cleaned to the <code><NetWorker_install_path>\nsr\logs\media.log</code> file. UNIX: Sends an email to the root account stating that a device cleaning operation is complete.
Device cleaning required	Windows: Reports that a device requires cleaning to the <code><NetWorker_install_path>\nsr\logs\media.log</code> file. UNIX: 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 <code><NetWorker_install_path>\nsr\logs\media.log</code> file. UNIX: Sends an email to the root account stating that a device has been automatically disabled.
Device ordering issue detect	Windows: Provides the syntax for the smtmail command to send an email to the administrator account with the message “Check system device ordering. Moving device on <i>NetWorker_server</i> to service mode. To correct, scan for devices in NMC and re-enable the device. Refer to the section “Devices -> Replace a drive” in NetWorker Procedure Generator for full details. The action attribute must be modified to replace mailserver with the actual hostname of the mail server. “Using smtmail to email notifications ” on page 488 describes how to customize the smtmail command. UNIX: Sends an email to the root account with the message “Check system device ordering. Moving device on <i>NetWorker_server</i> to service mode. To correct, scan for devices in NMC and re-enable the device. Refer to the section “Devices -> Replace a drive” in NetWorker Procedure Generator for full details.”
Event log (Windows only)	Logs notification events triggered by events and priorities to the Event Log.
Filesystem full - recovering adv_file space	Launches the nsrim program to remove aborted and expired save sets. Used with advanced file type devices only.
Filesystem 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. UNIX: 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\logs\messages</code> log file. Unix: 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\logs\index.log</code> file. UNIX: Sends this email to root: “Check the size of the client file index because it will soon exceed the space available.”
Log default	Windows: Sends data about NetWorker events to the <code>C:\Program Files\EMC NetWorker\logs\messages</code> log file. UNIX: 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 83 Preconfigured notifications (3 of 5)

Notification	Default action
NetWorker Daemons Not Running	<p>Windows: Provides the syntax for the smtpmail 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 mailserver with the actual hostname of the mail server. “Using smtpmail to email notifications” on page 488 describes how to customize the smtpmail program.</p> <p>UNIX: Sends an email to the root account stating that NetWorker daemons are not running on the NetWorker server.</p>
New Virtual Machine	<p>Windows: Reports a message that new virtual machines have been detected to the <code><NetWorker_install_path>\nsr\logs\messages</code> log file.</p> <p>UNIX: Sends an email to the root account stating that new virtual machines have been detected.</p>
Policy completion	<p>Windows: Sends an event notification to the <code><NetWorker_install_path>\nsr\logs\policy.log</code> file with a message that a VMware protection policy has been completed.</p> <p>UNIX: Sends an email to the root account with a message that a VMware protection policy has been completed.</p>
Registration	<p>Windows: Sends messages about the registration status of your NetWorker products to the <code><NetWorker_install_path>\nsr\logs\messages</code> log file.</p> <p>UNIX: Sends this email to root: check the registration status.</p>
Resource File Corruption	<p>Windows: Provides the syntax for the smtpmail 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 mailserver with the actual hostname of the mail server. “Using smtpmail to email notifications” on page 488 describes how to customize the smtpmail program.</p> <p>UNIX: Sends an email to the root account stating that resource file corruption has been detected on the NetWorker server</p>
Savegroup completion	<p>Windows: Reports the degree of success in completing all of the scheduled backups, cloning, and archive operations for a group to the <code><NetWorker_install_path>\nsr\logs\savegrp.log</code> file.</p> <p>UNIX: Sends an email to the root account of the NetWorker server to report the degree of success in completing all of the scheduled backups, cloning, and archive operations for a group.</p>
Savegroup failure	<p>Windows: Reports when a group backup fails to start at the scheduled time in the <code><NetWorker_install_path>\nsr\logs\savegrp.log</code> file.</p> <p>UNIX: Sends an email to root to report when a group backup fails to start at the scheduled time.</p> <p>Possible reasons include that the previously scheduled backup is still running.</p>

Table 83 Preconfigured notifications (4 of 5)

Notification	Default action
Save set marked suspect	<p>Windows: Provides the syntax for the smtpmail program to send an email to the administrator account when a save set has been marked suspect. The action attribute must be modified to replace mailserver with the actual hostname of the mail server. “Using smtpmail to email notifications” on page 488 describes how to customize the smtpmail program.</p> <p>UNIX: Sends an email to the root account when a save set has been marked suspect.</p>
Scheduled clone completion	<p>Windows: Sends an event notification to the <code><NetWorker_install_path>\nsr\logs\clone.log</code> file with a message that a scheduled clone operation has completed.</p> <p>UNIX: Sends an email to the root account with a message that a scheduled clone operation has completed.</p>
Scheduled clone failure	<p>Windows: Sends an event notification to the <code><NetWorker_install_path>\nsr\logs\clone.log</code> file with a message that a scheduled clone operation has failed.</p> <p>UNIX: Sends an email to the root account with a message that a scheduled clone operation has failed.</p>
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 694 provides details on SNMP notifications
Tape mount request 1	<p>Windows: Requests media be mounted in a device and displays a pending message in the <code><NetWorker_install_path>\nsr\logs\messages</code> log file.</p> <p>UNIX: Sends a request message to the system logger to mount a backup volume, using a local0 facility and an alert level.</p>
Tape mount request 2	<p>Windows: Requests media be mounted in a device and displays a critical message.</p> <p>UNIX: Sends a request message to the system logger to mount a backup volume, using a local0 facility and an alert level.</p>
Tape mount request 3	<p>Windows: Sends a request to mount a backup volume with a priority of Alert, to the <code><NetWorker_install_path>\nsr\logs\media.log</code> file.</p> <p>UNIX: Sends an email to the root account requesting that the tape be mounted.</p>
Tape mount request 4	<p>Windows: Provides the syntax for the smtpmail 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 mailserver with the actual hostname of the mail server. “Using smtpmail to email notifications” on page 488 describes how to customize the smtpmail program.</p> <p>UNIX: Sends an email to the root account stating that a Tape mount request 4 event has occurred.</p>
Verify Label failed on unload	<p>Windows: Provides the syntax for the smtpmail 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 mailserver with the actual hostname of the mail server. “Using smtpmail to email notifications” on page 488 describes how to customize the smtpmail program.</p> <p>UNIX: Sends an email to the root account stating that a label verification on unload operation has failed.</p>

Table 83 Preconfigured notifications (5 of 5)

Notification	Default action
VMware Protection Policy Failure	Windows: Sends an event notification to the <code><NetWorker_install_path>\nsr\logs\policy.log</code> file with a message that a VMware protection policy failed. UNIX: Sends an email to the root account with a message that a VMware protection policy failed.
Volume Marked full	Windows: Provides the syntax for the smtpmail 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 488 describes how to customize the smtpmail program. UNIX: 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><NetWorker_install_path>\nsr\logs\media.log</code> file with a message that a volume with the Scan needed flag is detected. UNIX: 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” on page 484](#)
- ◆ [“Actions” on page 486](#)
- ◆ [“Priorities” on page 489](#)

Events

Events are the activities on the NetWorker server that can trigger a notification

[Table 84 on page 484](#) lists the events that trigger a notification.

Table 84 Events (1 of 2)

Event	Description
adv_file	The file system is full and is waiting for additional space.
Bootstrap	The bootstrap backup failed.
Bus/Device Reset	A SCSI bus or device reset has occurred.
Cleaning cartridge expired	A cleaning cartridge has expired and needs replacing.
Cleaning cartridge required	Mount the cleaning cartridge.
Client	NetWorker client software has been installed on a host.
Deleted media	A media device has been deleted.
Device cleaned	A device has been cleaned.
Device cleaning required	A device requires cleaning.
Device disabled	A device has been automatically disabled.

Table 84 Events (2 of 2)

Event	Description
Hypervisor	New virtual clients have been auto-discovered, or auto-discovery failed.
Index	An index needs attention.
License expiration	A license has expired.
Media	A media related event occurred. For example, a volume may require mounting.
Media attention	Media needs operator attention to mount or unmount backup volumes.
Media capacity	A volume has almost reached the maximum number of save sets allowed in the media database.
Media request	Media needs operator attention to mount backup volumes.
Potential device ordering issue	A device ordering or serial mismatch error has occurred.
Resource File	A resource file corruption has occurred
Registration	Product registration needs attention.
Savegroup	A backup group has completed the backup.
Savegroup failure	A backup group has completed with failures.
Server	Other server events (for example, restarting the NetWorker server).
Storage node	A storage node has been installed.
Volume scan needed	A volume with the scan needed flag has been detected
Write completion	A write operation is complete.

Actions

The Actions attribute defines the action that the NetWorker server takes after an event notification occurs. [Table 85 on page 486](#) provides a summary of actions.

Table 85 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 f to identify a specific file. For example: <pre>nsrlog -f log file path</pre> <p>If no option is specified, then messages go to the <code>/nsr/logs/messages</code> file.</p>
nsrlpr	Windows only, prints information to a printer. Use option P to identify a specific printer. For example: <pre>nsrlpr -P printer_name</pre> <p>The printer can also be a remote print server such as LAN manager printer. In this case, use the following syntax <pre>nsrlpr -P \\server_name\printer_name</pre> <p>“Using nsrlpr to print notifications” on page 487 provides more information about nsrlpr.</p> </p>
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.
smtmail	Windows only, sends an email to the specified user.
nsrtrap	Sends notifications to an SNMP management console. Use with the following options: <ul style="list-style-type: none"> • -c <i>community</i> (if not specified, then the default public is used) • -f <i>file</i> (reads message from a file and sends as snmp trap.) • -i <i>version</i> (if not specified, then the default version is SNMPV2) • -s <i>specific</i> (default is NetWorker enterprise assignment, which is 1) • -t <i>trap</i> (default trap is #6 which is the enterprise-specific trap) • -u <i>snmp uptime</i> • -v <i>verbose</i>

Third-party programs can also be used for the action, as long as 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 **smtmail** 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 nsrlpr to print notifications

NetWorker Server, on Windows systems only, looks at the printer named in each of these two sources to determine which to use to print notifications with **nsrlpr**:

- ◆ The Printer attribute in the Group resource. This entry is ignored if a printer is named in the Action attribute for a notification.
- ◆ The printer named in the Action attribute for a notification, or that you specified by using the -P option of the **nsrlpr** program.

Designating a printer for a Notification Resource

To designate a printer for a Notification resource:

1. From the **Administration** window, click **Configuration**.
2. Click **Notifications**.
3. Right-click a notification, then select **Properties**. The **Properties** dialog box appears.
4. In the **Action** attribute, type:

```
nsrlpr -P printer_name
```

where *printer_name* is the name of the designated printer.

If the printer name has spaces, such as *eng printer one*, then enclose the printer name in double quotes, as shown here:

```
nsrlpr -P "eng printer one"
```

If the printer is associated with a particular server, as is the case with Microsoft LAN Manager printers, use this syntax:

```
nsrlpr -P \\server_name\printer_name
```

where:

- *server_name* is the name of the server to which the printer is attached.
- *printer_name* is the name of the printer to use.

Send the bootstrap notification printout to its group's printer

To send the bootstrap notification printout to the printer defined in the Printer attribute of the Group resource:

1. From the **Administration** window, click **Configuration**.
2. Click **Notifications**.
3. Right-click **Bootstrap**, then select **Properties**. The **Properties** dialog box appears.
4. In the **Action** attribute, type:

```
nsrlpr -P %printer
```

Testing the nsrlpr program

To test the **nsrlpr** program, type:

```
nsrlpr -P printer_name text_file
```

where:

- ◆ *printer_name* is the name of the printer to use.
- ◆ *text_file* is the name of a text file to print.

The printer name was typed incorrectly if you receive this error message:

```
Error: print server did not accept request. Job aborted.
```

Once you can print from the command prompt, enter this command to change the Action attribute to the print command:

```
nsrlpr -P printer_name
```

NOTICE

Print jobs sent by the NetWorker Backup and Recover Server service run in the Local System context. Under certain conditions, it may not have access to network print queues. Microsoft Knowledge Base articles 132679 and 143138 on the Microsoft web site provide more information.

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 e-mail addresses.

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 message sent from standard input.

The message is terminated in one of the following ways:

- ◆ An EOF.
- ◆ CTRL-Z on console.
- ◆ Aline consisting of a single period (.).

To use the **smtpmail** program to email event notifications:

1. From the **Administration** window, click **Configuration**.
2. Click **Notifications**.
3. Right-click on the notification, then select **Properties**. The **Properties** dialog box appears.

4. In the **Action** attribute, type:

```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- **-s subject**— includes a standard e-mail header with the message and specifies the subject text for that header. Without this option, the **smtpmail** program assumes that the message contains a properly formatted e-mail header and nothing is added.
- **-h mailserver** — specifies the hostname of the mailserver 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 assigned to the message is Waiting. The priority of the second request is Alert. The priority of the third request is Critical.

[Table 86 on page 489](#) lists the priorities upon which notifications are based.

Table 86 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.

NOTICE

Event priorities are sorted alphabetically, rather than by severity.

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.

[“Viewing log files” on page 803](#) provides information about viewing log files.

To better access and use these event logs on 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 listed in [Table 87 on page 490](#).

Table 87 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 (savegroup, server, registration, and so on).
Severity	Mapped from NetWorker notification priority: <ul style="list-style-type: none"> • 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.

Creating a custom notification

NetWorker also provides preconfigured notifications. [“Preconfigured notifications” on page 480](#) provides a complete list of preconfigured notifications.

To create a custom notification:

1. From the **Administration** window, click **Configuration**.
2. Right-click **Notifications**, then select **New**. The **Create Notification** dialog box appears.
3. In the **Name** attribute, enter 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, enter a command to execute in response to the selected events and priorities. [Table 85 on page 486](#) provides command options.
7. Click **Ok**.

Editing a notification

To edit a notification:

1. From the **Administration** window, click **Configuration**.
2. Click **Notifications**.
3. Right-click the notification to edit, then select **Properties**. The **Properties** dialog box appears.
4. Make any required changes, then click **Ok**.

NOTICE

You cannot change the name of a notification.

Copying a notification

To copy a notification:

1. From the **Administration** window, click **Configuration**.
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, except for **Name** attribute.
4. In the **Name** attribute, enter a name for the new notification.
5. Edit any other attributes as appropriate, then click **Ok**.

Deleting a custom notification

To delete a custom notification:

1. From the **Administration** window, click **Configuration**.
2. Click **Notifications**.
3. Right-click the notification to delete, then select **Delete**.
4. When prompted, click **Yes** to confirm the deletion.

NOTICE

You cannot delete any of the preconfigured notifications.

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.

For Windows NetWorker servers, use the **smtpmail** program to send the owner notification email. [“Using smtpmail to email notifications” on page 488](#) 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.

To configure an owner notification:

1. From the **Administration** window, click **Configuration**.
2. Select **Clients** in the left navigation pane.
3. Right mouse click the client and select **Properties**.
4. Select **Globals (2 of 2)**.
 - For a Windows NetWorker server, use the **smtpmail** program to configure the email notification. [“Using smtpmail to email notifications” on page 488](#) 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**.

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.emc.com]
Sent: Thursday, March 22, 2012 12:45 PM
To: debbie@mymailhost.com
Subject: Backup status for client xyz in group abc

cdcsdunndl1c, savefs, "succeeded:full:savefs"
* cdcsdunndl1c:savefs savefs cdcsdunndl1c: succeeded.
cdcsdunndl1c, C:\cmdcons\system32, "<NULL>:full:save"
* cdcsdunndl1c:C:\cmdcons\system32 cdcsdunndl1c:C:\cmdcons\system32
  aborted
* cdcsdunndl1c:C:\cmdcons\system32 Termination request was sent to job
  64006 as requested; Reason given: Aborted
```

Reporting group status and backup job status

When you perform a backup, clone or archive activities for clients in an active group, NetWorker records the status of the group completion and job activities. There are three way to report job activities:

- ◆ In the Monitoring window for the NetWorker server in NMC. [“Monitoring NetWorker server activities” on page 462](#) describes how to view the backup group completion status in the Monitoring window.
- ◆ Through predefined savegroup completion notifications. [“Savegroup completion and failure notifications” on page 493](#) provides more information.
- ◆ By querying the job status. [“Querying the job status” on page 497](#) provides more information.

Savegroup completion and failure notifications

By default, for a UNIX NetWorker server, an email of the Savegroup completion and the Savegroup failure report are sent to the root account after all client activities of a group are finished. On Windows, the Savegroup completion and Savegroup failure details are written to a savegrp.log file located in the `NetWorker_install_dir\logs` directory. [“Notifications” on page 479](#) describes how to customize Savegroup completion and Savegroup failure notifications.

The following sections provide more information about the Savegroup completion and Savegroup failure notifications:

- ◆ [“The format of the Savegroup completion and Savegroup failure notifications” on page 493](#)
- ◆ [“Customizing the save sets status in the Savegroup completion and Savegroup failure notifications” on page 496](#)
- ◆ [“Filtering the savegroup completion report messages” on page 496](#)

The format of the Savegroup completion and Savegroup failure notifications

The Savegroup completion and Savegroup failure notifications are divided into a number of sections that describe the job activities for a savegroup.

The sections of the Savegroup completion and Savegroup failure notifications are:

- ◆ The summary line — Provides a summary of the overall status of the savegroup. The Savegroup completion and the Savegroup failure notifications provide summaries that includes the name of the savegroup, the total number of clients that are in the savegroup, the number of clients that experienced a job failure, and the status of the clone job, when automatic cloning is enabled for the group.

For example:

```
NetWorker savegroup: (alert) Default completed, Total 2 client(s), 2
Failed. 1 Succeeded(Save Set Cloning Failed). Please see group
completion details for more information.
```

The Savegroup failure notification provides an additional summary line, reporting the number of:

- Clients in the savegroup.
- Disabled clients in the savegroup.
- Clients whose hostname cannot be resolved.
- Successful and unsuccessful BMR client backups.
- Successful and unsuccessful checkpoint-enabled client backups.

For example:

```
NetWorker savegroup failure: (alert) Default Completed/Aborted,
Total 7 client(s), 1 Clients disabled, 1 Hostname(s) Unresolved, 2
Failed, 1 Succeeded, 1 CPR Failed, 1 CPR Succeeded, 1 BMR Failed, 1
BMR Succeeded.
```

- ◆ The summary body — Provides more detailed information about the jobs in a savegroup including:
 - The Start time of the savegroup.
 - The Restart time of the savegroup. This only appears if the savegroup was restarted.
 - The Clone start time of the savegroup. This only appears if automatic cloning is enabled for the savegroup.
 - The End time of the savegroup.
 - A list of hostnames that experienced a job failure and the type of failure.

For example:

```
hostname resolution failed for 1 client(s)
Unresolved: myfailedclient1.emc.com
Failed: myfailedclient1.emc.com, myfailedclient2.emc.com
Disabled: mydisabledclient.emc.com
Failed after CPR: mycprfailed.emc.com
BMR Failed:myfailedBMR.emc.com
Succeed with warning(s):mywarnings.emc.com
Succeeded: mysuccess.emc.com
```

```
Start time: Thu Apr 19 16:45:55 2012
Clone Start: Thu Apr 19 16:50:55 2012
End time: Thu Apr 19 16:55:22 2012
```

Automatic cloning of save sets to pool Default Clone failed.

- ◆ The save set status section is divided into six status categories:
 - Never started save sets — Provides a summary of save sessions that were not started.
 - Unsuccessful Save Sets — Provides a summary of finished save sessions that are considered unsuccessful or incomplete. Save sets that complete with warnings are considered unsuccessful when the Success threshold attribute for the savegroup resource is set to Success. [“Customizing the save sets status in the Savegroup completion and Savegroup failure notifications” on page 496](#) provides detailed information about configuring the Success threshold attribute.

- **Successful Save Sets** — Provides a summary of finished save sessions that are considered successful.
- **Succeeded With Warnings** — Provides a summary of finished save sessions that are considered successful but encountered warnings. Save sessions that finish with warnings are considered successful when the Success threshold attribute for the savegroup resource is set to Warning. [“Customizing the save sets status in the Savegroup completion and Savegroup failure notifications” on page 496](#) provides detailed information about configuring the Success threshold attribute.
- **Cloned Save Sets** — Provides the status of the clone job. This section only appears when automatic cloning is enabled for the savegroup.
- **Succeeded In Previous Runs** — Provides a summary of save sessions that succeeded in a previous backup attempt of the savegroup. This section only appears when a save group is restarted after a previous failure.

For example:

```
--- Never Started Save Sets ---
* client1:D:\ save was never started

--- Unsuccessful Save Sets ---
* myfailed_client:index 90018:save: Cannot open a save session with
NetWorker server 'bu-t3-7.lss.emc.com': no matching devices for save
of client `bu-t3-7.lss.emc.com'; check storage nodes, devices or
pools
* myfailed_client:index cdcsdunndllc:index: retry #1

* myfailed2_client:bootstrap 90018:save: Cannot open a save session
with NetWorker server 'myfailed2_client': no matching devices for
save of client `bu'; check storage nodes, devices or pools
myfailed2_client:bootstrap: retry #1

--- Successful Save Sets ---
* mysuccessful_client:savefs savefs mysuccesful_client: succeeded.
* mysuccessful_client C:\cmdcons          level=full,  7949 KB
00:00:05    183 files
* mysuccessful_client:C:\cmdcons completed savetime=1334782210

--- Cloned Save Sets ---
Automatic cloning of save sets to pool Default Clone failed.
```

Customizing the save sets status in the Savegroup completion and Savegroup failure notifications

How the status of the save sessions is reported in Savegroup completion and Savegroup failure notifications is determined by the value that is defined in Success threshold for the group. Consider the following:

- ◆ The Success threshold value is defined under the Advanced tab of the Group Properties window.
- ◆ When the Success threshold is set to:
 - Warning — Any save set that completes with warnings will be reported as successful. This is the default value.
 - Success — Any save set that completes with warnings is considered to have failed and is reported as unsuccessful. Failed save sets are retried according to the value defined in the Client Retries attribute, also located under the Advanced tab of the Group Properties window.
- ◆ The Success threshold attribute also applies to the save sets displayed in the Completed successfully and Failed sections of the Group details window. [“Viewing group backup details” on page 466](#) provides more information about viewing group backup details.

Filtering the savegroup completion report messages

In NetWorker 8.0 and later, the **nsrscm_filter** command provides the ability to filter the savegroup completion messages on a UNIX or Linux NetWorker server based on a user defined filter file.

There are a number of ways to use the **nsrscm_filter** command to filter savegroup completion messages in NetWorker:

- ◆ Modify the action attribute of the savegroup completion notification to direct the output to a savegrp log file. Use the **nsrscm_filter** program to filter the savegrp log file.

The **nsrscm_filter** program is located in the following directory by default:

- Solaris and Linux: /usr/sbin/
- HPUX: /opt/networker/bin
- AIX: /usr/bin

- ◆ Modify the action attribute of the savegroup completion notification to use the **nsrscm_filter** program and generate a filtered savegrp log file.

For example:

```
name: savegroup completion report;
action: /usr/sbin/nsrscm_filter -f /nsr/res/filter_msgs -l 2 -D 1
-s /nsr/logs/scm/savegrp.$$ -o /nsr/logs/scm/scmfilter.$$ ;
event: Savegroup;
priority: alert, notice;
```

where *filter_msgs* is the user defined filter file. A template filter file **filter_msgs.templ** is provided in the same location as the **nsrscm_filter** program.

“[Editing a notification](#)” on page 491 describes how to modify preconfigured notifications.

- ◆ Modify the **action** attribute of the savegroup completion notification to call a script that runs the **nsrscm_filter** command and sends the filtered output in an email.

A template script file, **scm-notification.sh** is provided in the same location as the **nsrscm_filter** program.

The *NetWorker 8.0 Command Reference Guide* and the UNIX man pages provide **nsrscm_filter** usage information.

Querying the job status

When a backup, clone or archive job is run within an active savegroup, job information is stored in savegrp log files and the jobs database (jobsdb) on the NetWorker server host.

Before you query job status, review this information:

- ◆ NetWorker logs details of the corresponding savegroup’s child jobs into text files located in:

```
NetWorker_install_dir\logs\sg\savegroup_name\Job_ID
```

where *Job_ID* is the file whose name corresponds to a savegroup child job ID.

Note: In NetWorker 7.5.x and 7.6.x, Savegroup log by job id was a configurable option and not enabled by default. In NetWorker 8.0 and higher, Savegroup log by job id is the default behaviour and is not configurable.

- ◆ The job files are purged based on value of the *Jobsdb retention in hours* attribute in the properties of the NetWorker server resource. Prior to NetWorker 8.0, the job files were overwritten each time the group started.
- ◆ Each jobs file contains the output that the job returns to stderr or stdout.
- ◆ The jobsdb stores the savegroup completion information. Prior to NetWorker 8.0 NetWorker stored the Savegroup completion information in the Completion attribute of the Group resource in the resource database of the NetWorker server.

The NetWorker software provides two command line programs to query job information.

- ◆ “Using jobquery” on page 498 describes the **jobquery** program that is used to locate and retrieve information about all jobs including the child jobs of a savegroup.
- ◆ “Using nsrsgrpcmp” on page 499 describes how to use the **nsrsgrpcmp** program to query savegroup specific information from the job database. Use this program to retrieve completion information that was stored in the resource database in previous versions of NetWorker.

The man pages or the *EMC NetWorker 8.0 Command Reference Guide* provides more information on the **jobquery** and **nsrsgrpcmp** commands.

Using jobquery

The **jobquery** program provides an CLI similar to the **nsradmin** program. The jobquery program contacts the nsrjobd process to query job information 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.

In the query, the attribute name (for example, 'type') is preceded by a '.', and optionally followed by a ':' 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 ';'. 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 the following command:

```
jobquery -s NetWorker_server
```

If the **-s NetWorker_server** option is not used, **jobquery** attempts 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 will 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
```

The above example would return information on all savegroup jobs from the machine mars that are either in progress or in completed state.

- ◆ **show** — restricts the list of attributes returned for each resource descriptor that matches the query. For the above example, specifying the following:

`show name; job id; completion status; completion severity`

will return the names, job ids, completion status and completion severity for all matched completed and active savegroups.
- ◆ **print** — executes 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 *EMC NetWorker Command Reference Guide* provides detailed information about the **jobquery** program.

Using nsrsgrpcomp

Similar to the **jobquery** program, the **nsrsgrpcomp** program queries information stored in the jobsdb. The **nsrsgrpcomp** program differs from the **jobquery** program because it also queries savegrp log files and is limited to savegrp job information only.

Use the **nsrsgrpcomp** program to:

- ◆ Provide savegroup completion output.
- ◆ Retrieve all job details for a group including save set status.
- ◆ Retrieve all job details for a client in a group.
- ◆ Retrieve all savegroup job details in the jobs database.

Example 39 View the completion report

To generate the completion report that was previously stored in the **completion** attribute for a savegroup in the resource database of the NetWorker server, type:

```
# nsrsgrpcomp -b -1 group_name
```

where **-b -1** is optional and used to override the default 2kb limit for job output.

Example 40 Summary of last savegroup backup

To generate a summary of the savegroup jobs for the last run of a savegroup run, type:

```
# nsrsgrpcomp -H group_name
```

For example, to generate a summary report of the last run of the Default savegroup:

```
nsrsgrpcomp -H Default
```

```
bu-t3-7.lss.emc.com, bootstrap, "failed:full:bootstrap"
cdcsdunnd11c, savefs, "succeeded:full:savefs"
cdcsdunnd11c, C:\cmdcons\system32, "succeeded:incr:save"
cdcsdunnd11c, index, "failed:full:index"
```

Example 41 Query the savegroup status

To query the jobs database for a status of a previous run of a savegroup:

1. Use **nsrsgrpcomp -L** to report a summary of all jobs stored in the jobs database.
2. For example:

```
nsrsgrpcomp -L Default
name, start time, job id, previous job id, completion status
Default, Tue Apr 17 10:40:04 2012(1334673604), 160023, 0, failed
Default, Tue Apr 17 14:15:46 2012(1334686546), 160033, 0, failed
```

3. Use **nsrsgrpcomp** with the **-t** option and the **start time** value associated with the savegroup.

```
nsrsgrpcomp -b -1 -t "Tue Apr 17 10:40:04 2012" Default

bu-t3-7.lss.emc.com, bootstrap, "failed:full:bootstrap"
* bu-t3-7.lss.emc.com:bootstrap Failed with error(s)
* bu-t3-7.lss.emc.com:bootstrap 90018:save: Cannot open a save
session with NetWorker server 'bu-t3-7.lss.emc.com': no matching
devices for save of client `bu-t3-7.lss.emc.com'; check storage
nodes, devices or pools
cdcsdunnd11c, savefs, "succeeded:full:savefs"
* cdcsdunnd11c:savefs savefs cdcsdunnd11c: succeeded.
cdcsdunnd11c, C:\cmdcons\system32, "succeeded:incr:save"
* cdcsdunnd11c:C:\cmdcons\system32 cdcsdunnd11c:
C:\cmdcons\system32 level=incr,      0 KB 00:00:03      0 files
* cdcsdunnd11c:C:\cmdcons\system32 completed savetime=1334587461
cdcsdunnd11c, index, "failed:full:index"
* cdcsdunnd11c:index Failed with error(s)
* cdcsdunnd11c:index 90018:save: Cannot open a save session with
NetWorker server 'bu-t3-7.lss.emc.com': no matching devices for save
of client `bu-t3-7.lss.emc.com'; check storage nodes, devices or
pools
```

The man pages or the *EMC NetWorker Command Reference Guide* provides detailed information about the **nsrsgrpcomp** program.

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” on page 462](#) 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 501](#) 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 of given to 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 42 Summary report of recover jobs

To generate a summary report of each recover job in the jobsdb, type:

```
nsrrecomp -L
```

Example 43 Recovery job completion report

To generate a completion report for recover job, type:

```
nsrrecomp -b -1 recover_job_name
```

where **-b -1** is optional and used to override the default 2kb limit for job output.

Example 44 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 man pages or the *EMC NetWorker Command Reference Guide* provides detailed information about the **nsrrecomp** program.

CHAPTER 17

NMC Server Management

This chapter covers these topics:

- ◆ NMC server authentication 504
- ◆ Moving the NMC server 528
- ◆ Setting system options 531
- ◆ Setting environment variables..... 533
- ◆ Accessing the Console Configuration Wizard 535
- ◆ NetWorker NMC server maintenance tasks 535
- ◆ Displaying international fonts in non-US locale environments 538
- ◆ NetWorker License Manager 538

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. You do not require a secure http (https) connection because only the Java client transfers information and performs authentication between the NMC server and NMC client. The NMC server uses SSL to encrypt the username and password that you specify in the login window and authenticates the credentials in one of the following methods:

- ◆ “Native NMC-based authentication” on page 505.
The first time a NMC client connects to the NMC server, this authentication method is used.
- ◆ “An external authentication authority” on page 506.

Consider the following:

- ◆ The NMC server restricts user privileges based on the following three authorization roles. These roles cannot be deleted and their privileges cannot be changed. [Table 88 on page 504](#) describes the three Console authorization roles.

Table 88 Console roles

User Role	Privileges
Console Security Administrator	Add, delete, and modify Console Users. Configure login authentication such as configuring the NMC server to: <ul style="list-style-type: none"> - use LDAP authentication instead of native NMC authentication. - use native NMC authentication instead of LDAP authentication. Control user access to managed applications such as a NetWorker server. All tasks available to a ‘Console User’ role.
Console Application Administrator	Configure Console system options. Set retention policies for reports. View custom reports. Specify the NetWorker server to backup the Console database. Specify a NetWorker License Manager server. Run the Console Configuration wizard. All tasks available to a Console User role.
Console User	All tasks except for those tasks explicitly mentioned for the Console Security Administrator and the Console Application Administrator . Tasks include: <ul style="list-style-type: none"> Add/delete hosts, folders. Add/Delete Managed applications for NW, Data Domain, Avamar. Create/Delete their own reports. Set features for managed applications. Manage a NetWorker server with the appropriate privilege levels. Dismiss events.

- ◆ When an NMC user is authenticated by the NMC server, the user is granted privileges to manage all of the NetWorker hosts defined on the NMC server. [“Restricting a user’s view of managed servers” on page 526](#) describes how to restrict the NetWorker servers that a console user can view and manage.
- ◆ If the NetWorker server and the NMC server are on different hosts, ensure the administrators list attribute on the NetWorker server includes the appropriate NMC accounts before connecting to a NetWorker server. [“The administrator list” on page 558](#) provides more information.
- ◆ [“Troubleshooting authentication errors” on page 521](#) describes how to troubleshoot common authentication errors on the NMC server.
- ◆ [“Troubleshooting login errors” on page 526](#) describes how to troubleshoot common errors when attempting to log into the NMC server.

Native NMC-based authentication

Native NMC-based authentication uses a data store on the NMC server host to authenticate NMC users. Native NMC based authentication is enabled by default and is pre-configured the first time the NMC server is accessed from a NMC client.

NMC user names and passwords are maintained on the NMC server. By default, one NMC user with the login ID, administrator, is created when the NMC server is accessed for the first time. Additional set up is not required to enable Native NMC based authentication, however additional NMC user accounts can be added with each users assigned to different Console roles.

The following sections describe how to:

- ◆ [“Add NMC users” on page 505](#)
- ◆ [“Modify the NMC user” on page 506](#)
- ◆ [“Delete an NMC user” on page 506](#)

Add NMC users

To add additional NMC users when using Native NMC login authentication:

1. Log into the NMC server as a **Console Security Administrator**. The ‘administrator’ account is a **Console Security Administrator**.
2. From the **Console** window, click **Setup**.
3. In the left pane, right-click **Users**, then select **New**. The **Create User** dialog box appears, with the **Identity** tab displayed.
4. Enter a username.

The username cannot:

- Exceed 64 characters.
 - Use spaces, or any of these characters: : < > /
 - Use characters with an ASCII value less than or equal to 32.
 - Begin a username with an underscore (_) character.
5. Optional, enter the full name of the user and a user description.

6. Select the **Console user roles**. [Table 88 on page 504](#) provides more information about **Console user roles**.
7. Enter the user password. Passwords must be at least eight characters long and cannot be the same as the user name. This requirement is enforced for all newly created or modified users. For customers upgrading from a previous release, this requirement is enforced when user passwords are changed.
8. In the **Confirm Password** attribute, re-enter the password.
9. Click **OK**.

Modify the NMC user

The following attributes of an NMC user can be modified:

- The password
- The Descriptive information
- The Console roles

To modify an existing NMC user:

1. Log into the NMC server as a Console Security Administrator. The administrator account is a Console Security Administrator.
2. From the **Console** window, click **Setup**.
3. In the left pane, select **Users**.
4. Right-click the user and then select **Properties**.
5. Under the **Identification** tab, modify the attributes as required. [Table 88 on page 504](#) describes the privileges assigned to each Console role.

Delete an NMC user

To delete an existing NMC user:

1. Log into the NMC server as a Console Security Administrator. The NMC user administrator is a Console Security Administrator.
2. From the **Console** window, click **Setup**.
3. In the left pane, select **Users**.
4. Right-click the user and then select **Delete**.
5. Click **Yes** to confirm the deletion.
6. If the user had saved customized reports, a dialog box prompts for the username to which to reassign those reports. Otherwise, the reports can be deleted.

An external authentication authority

Using an external authentication authority enables you to log into the NMC server with user names and passwords that are maintained by a Lightweight Directory Access Protocol (LDAP), Lightweight Directory Access Protocol over SSL (LDAPS), or a Microsoft Active

Directory server (AD). User privileges are controlled by mapping LDAP or AD user roles or user names to Console user roles. There is no need to add user names and passwords on the NMC server.

The NetWorker software automatically distributes the LDAP or AD configuration file from the NMC server to the NetWorker servers that the NMC server manages. This distribute process automatically puts the managed NetWorker servers in LDAP or AD mode.

When an LDAP or AD user logs into the NMC server and connects to a NetWorker server:

- ◆ The NetWorker server performs a look-up to get the LDAP or AD group that the OS authenticated user belongs to, in the external authority. The NetWorker server does not authenticate the user against the LDAP authority.
- ◆ The privileges assigned to a user on the NetWorker server are based on the LDAP user or the group entries present in the external roles attribute of the User Group resource on the NetWorker server. [“Managing server access” on page 558](#) describes how to configure the external roles attribute for LDAP or AD groups on the NetWorker server.

The following sections describe tasks specific to configuring and managing an external authentication authority:

- ◆ [“Preparing the NMC server and NetWorker servers for LDAPS” on page 507](#)
- ◆ [“Configuring the NMC server for LDAP, LDAPS, or AD authentication” on page 508](#)
- ◆ [“Adding or removing LDAP or AD console users” on page 518](#)
- ◆ [“Modifying NMC roles for LDAP or AD users and groups” on page 520](#)
- ◆ [“Modifying an LDAP or AD console user” on page 520](#)
- ◆ [“Deleting an LDAP or AD console user” on page 520](#)

Preparing the NMC server and NetWorker servers for LDAPS

Before you configure the NMC and NetWorker servers to use LDAPS, ensure that a local copy of the CA Certificate, Client Certificate, and Client Key reside on each NMC and NetWorker server in the same file system path.

Note: Ensure that the LDAPS certificates use the PEM format.

When the operating system of the NMC server and any NetWorker server differs, perform the following tasks to ensure that each host can successfully communicate with the LDAP server:

1. On a UNIX NMC server, create a subdirectory for the certificates in the *NMC_installation_directory/cst* directory.

For example, on a Solaris NMC server, create a subdirectory called corpldap in the */opt/LGTONmc/cst* directory.

2. On a UNIX NetWorker server, create a subdirectory for the certificates in the */opt/nsr/cst* directory.

For example, create a subdirectory called corpldap in the */opt/nsr/cst* directory.

3. On Windows NMC server, create a subdirectory for the certificates in the *NMC_installation_directory\cst* directory.

For example, create a subdirectory called corpldap in the C:\Program Files\EMC NetWorker\Management\GST\cst directory.

4. On Windows NetWorker server, create a subdirectory for the certificates in the *NetWorker_installation_directory\cst* directory.

For example, create a subdirectory called corpldap in C:\Program Files\EMC NetWorker\nsr\cst.

5. Copy the CA Certificate to the new subdirectory on each host that will use LDAPS. If the LDAPS configuration requires a certificate from the client side, copy the Client Certificate and Client Key to the new directory on each host.

Note: To secure the subdirectory, you can restrict access to the directory. For a UNIX host, ensure that the root account on UNIX has access to the directory. For a Windows host, ensure that the Administrator and Local System accounts have access to the directory.

Configuring the NMC server for LDAP, LDAPS, or AD authentication

To configure an external authority for LDAP, LDAPS, or AD authentication:

1. Log into the NMC server as a Native NMC based authentication user who is assigned to the Console Security Administrator role. The NMC user administrator is assigned to the Console Security Administrator role, by default.
2. From the **Setup** menu, select **Configure Login Authentication**.
3. In the **Select Authentication Method** window, select **External Repository**.
4. Click on the **Add** button to add a new external authentication authority.
5. Define the attributes for your configuration in the **Parameters** section. [Table 89 on page 509](#) summarizes and defines each attribute.

Table 89 Authority configuration parameters (1 of 3)

Parameter name	Parameter definition	Configuration Information
Authority Name	Descriptive name for the LDAP or AD server.	Required. This is a user defined field. If you configured the LDAPS certificate directories, ensure that the authority name matches the name of the subdirectory you created on the NMC server and the NetWorker server. For example corpldap
Provider Search Name	Hostname or IP address of the LDAP or AD server.	Required. For LDAPS, ensure that you specify the hostname exactly as it appears in the ca.cert file. For example, if the ca.cert file contains the FQDN of the LDAPDS server, you must specify the FQDN in the Provider Search Name field.
Distinguished Name	The dn of an LDAP or AD account that you use to perform operations such as searching for users and groups in the LDAP or AD hierarchy.	Specify an account on the LDAP or AD server that has full read access to the directory from which the AD or LDAP server accesses its data. Required.
Password	Password of the LDAP or AD account.	Required.
User Search Path	The dn to use when searching for users on the LDAP or AD server.	Required.
Group search path	The dn to use when searching for groups on the LDAP or AD server.	Required.
Group Name Attribute	Identifies the LDAP or AD group name in the User Search Path dn.	Required. Default value: cn
LDAPTimeout(millisecond)	The time out for LDAP or AD calls.	Range is 0 to 2 000 000 000 ms. A value of 0 indicates that calls will never time out. Required. Default value: 30000.
User ID Attribute	The user ID associated with the users in the User Search Path dn.	For LDAP this attribute is usually uid. For AD, this attribute is usually cn. Required. Default value: uid.
User Object Class	The object class that identifies users in the dn defined in the User Search Path.	Required.

Table 89 Authority configuration parameters (2 of 3)

Parameter name	Parameter definition	Configuration Information
Group Object Class	The object class that identifies groups in the LDAP or AD hierarchy of the dn defined in the User Search Path.	For LDAP, depending on the configuration, use groupOfNames or groupOfUniqueNames. For AD, use group. Required. Default value: groupOfUniqueNames.
Group Member Attribute	The group membership of users in dn that is defined in the User Search Path.	For LDAP: -If the Group Object Class is groupOfNames the attribute is usually member. -If the Group Object Class is groupOfUniqueNames the attribute is usually uniquemember. For AD the value is usually member. The default value is uniquemember. Required.
LDAP Debug level	Level of debug messages to log in the gstd.raw file.	Increase this value for troubleshooting purposes only. Default value: 0.
Protocol	Communication protocol between the NetWorker server and authentication server.	For LDAP or AD, select LDAP. For secure communications, select LDAPS.
Server Certificate (LDAPS only)	The full path to the CA certificate on the NMC server.	Required for LDAPS. When the NMC server and NetWorker server are on different platforms, use a forward slash to specify the path. For example: C:/Program Files/EMC NetWorker/Management/GST/cst/corpldap/ca.cert

Table 89 Authority configuration parameters (3 of 3)

Parameter name	Parameter definition	Configuration Information
Client certificate (LDAPS only)	The full path to the Client certificate on the NMC server.	Required for LDAPS when the LDAPS server requires a client certificate. When the NMC server and NetWorker server are on different platforms, use a forward slash to specify the path. For example: C:/Program Files/EMC NetWorker/Management/GST/cst/corpldap/client.cert
Client key (LDAPS only)	The full path to the Client key on the NMC server.	Required for LDAPS when the LDAPS server requires a client certificate. When the NMC server is a Windows host, use a double backslash to specify the path. For example: C:/Program Files/EMC NetWorker/Management/GST/cst/corpldap/client.key
Port Value	Port number of the server.	Required. Default value for AD and LDAP: 389 Default value for LDAPS: 636

The following examples describe how to determine what attributes values to use when configuring LDAP or AD authorities.

- ◆ [Example 45 on page 512](#) — Configuring an LDAP authority
- ◆ [Example 46 on page 514](#) — Configuring an AD authority

Example 45 Configuring an LDAP authority

In this example, a third party LDAP management tool, LDAPAdmin is used to view the properties of the LDAP configuration.

[Figure 37 on page 512](#) provides an example of the values required to specify the following attributes:

- ◆ Provider Server Name
- ◆ Distinguished Name
- ◆ User ID Attribute
- ◆ User Search Path — a combination of the ADDistinguished name and User Container name.
- ◆ User Object Class

Attribute	Value
uid	alberta_user1
givenName	alberta_test
objectClass	top
objectClass	person
objectClass	organizationalPerson
objectClass	inetOrgPerson
sn	alberta_user1
cn	alberta_test alberta_user1
userPassword	{SHA}i+PJQ7Fgn/+xRqZm0KBK34PJ0=

Figure 37 LDAP values for user attributes — LDAP Admin

[Figure 38 on page 512](#) provides an example of the values associated with following LDAP group attributes:

- ◆ Group Search Path — a combination of the Distinguished Name and Group Container name
- ◆ Group Member Attribute
- ◆ Group Object Class

Attribute	Value
objectClass	top
objectClass	groupOfUniqueNames
cn	AlbertaTestGroup2
uniqueMember	uid=alberta_user1,ou=AlbertaPeople,dc=alberta,dc=emc,dc=com
uniqueMember	uid=alberta_user3,ou=AlbertaPeople,dc=alberta,dc=emc,dc=com

Figure 38 LDAP values for Group attributes — LDAP Admin

Figure 39 on page 513 provides an example of the Manage Authentication Authorities screen with configuration details related to an LDAP server installation specified in the attribute fields.

The screenshot shows a window titled "Configure Login Authentication" with a sub-section "Manage Authentication Authorities". Below the title is a table with two columns: "Authority Name" and "Provider Server Name". The table contains one entry: "Alberta" with "alberta.lss.emc.com" as the provider server name. To the right of the table are buttons for "Up", "Down", "Add", and "Delete".

Below the table is a "Parameters" section with the following fields:

- Authority Name: Alberta
- Type: LDAP-v3 / Active Directory
- Provider Server Name: alberta.lss.emc.com
- Distinguished Name: cn=AlbertaSuperAdmin,dc=alberta,dc=emc,dc=com
- Password: [masked]
- User Search Path: ou=AlbertaPeople,dc=alberta,dc=emc,dc=com
- Group Search Path: ou=AlbertaGroups,dc=alberta,dc=emc,dc=com
- Group Name Attribute: cn
- LDAP Timeout (millisecond): 30,000

Below the parameters is an "Advanced" section with the following fields:

- User ID Attribute: uid
- User Object Class: inetOrgPerson
- Group Object Class: groupOfUniqueNames
- Group Member Attribute: uniqueMember
- LDAP Debug Level: 0
- Protocol: ldap
- Server Certificate File: [empty]
- Client Certificate File: [empty]
- Client Key File: [empty]
- Port Number: 389

At the bottom right of the window are buttons for "< Back", "Next >", and "Cancel".

Figure 39 Manage Authentication Authorities — LDAP

Example 46 Configuring an AD authority

In this example, the Active Directory Services Interfaces Editor (ADSI Edit) program is used to view the properties of the AD configuration.

Figure 40 on page 514 provides an example of the values required to specify the following attribute fields:

- ◆ Distinguished Name—a combination of the AD Distinguished name, User container, and User ID Attribute.
- ◆ User Search Path — a combination of the Distinguished name and User Container name.
- ◆ User Object Class
- ◆ User ID Attribute




Figure 40 ADSI and AD client properties window

Figure 41 on page 515 provides an example of the values associated with following AD group attributes:

- ◆ Provider Service Name
- ◆ Group Container
- ◆ Group Member Attribute
- ◆ Group Object Class
- ◆ Group Search Path — a combination of the Distinguished Name and Group Container name.

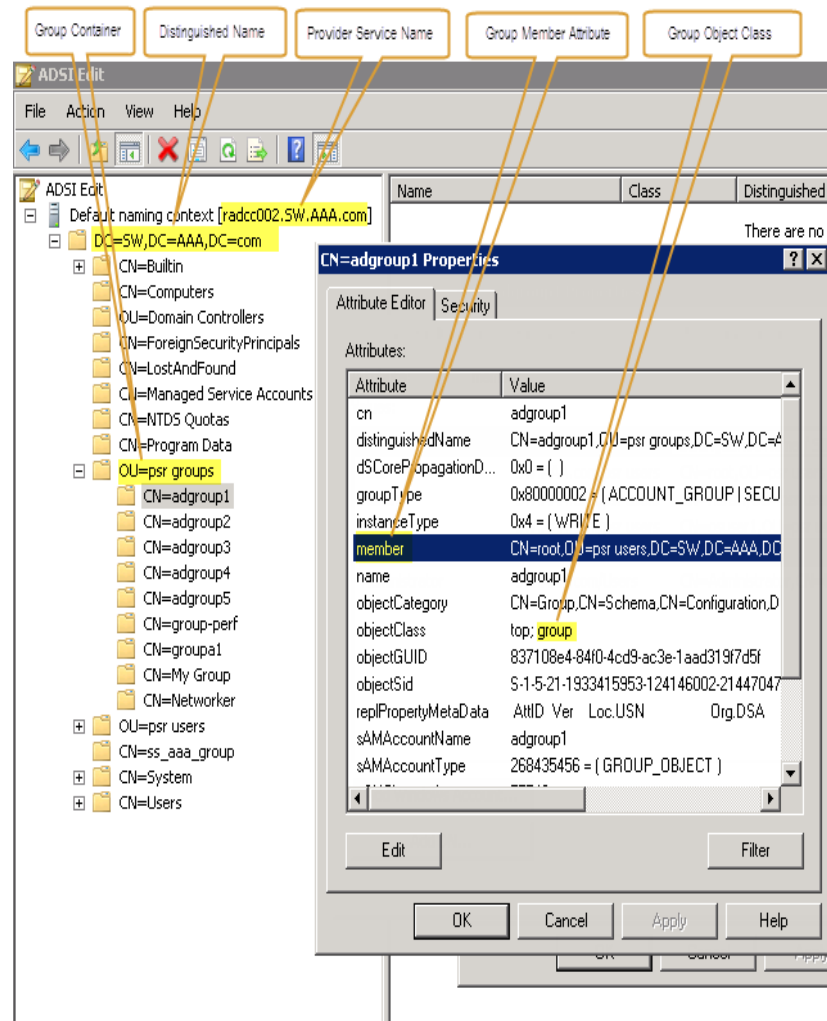


Figure 41 ADSI and AD group properties window

Figure 42 on page 516 provides an example of the Manage Authentication Authorities screen with configuration details related to an AD server installation specified in the attribute fields.




Figure 42 Manage Authentication Authorities — AD example

6. Click **Next**.

“[Troubleshooting authentication errors](#)” on page 521 describes common error messages that might display.

7. In the **External Roles** field, specify the LDAP or AD users and group to assign to the NMC Console **Security Administrator** role.

Using the values shown in [Example 45](#) and [Example 46](#):

- For LDAP, specify the group **AlbertaTestGroup1**.
- For AD, specify the group **adgroup1**.

8. Click **Next**.

If the user or group specified is not valid on the LDAP or AD server, the following message displays:

```
External role <user or group> is invalid
```

9. In the **Distributed Authority Configuration File** window, select the NetWorker servers that will use LDAP or AD. This will copy the LDAP configuration file from the NMC server to the `<NetWorker_install_path>\nsr\cst` directory on a Windows NetWorker server or the `<NetWorker_install_path>/nsr/cst` folder on a UNIX NetWorker server. The NMC server is selected by default.

10. Click **Distribute**.

11. If value specified in the **Distinguished Name** field is not valid the following error message displays:

```
Failed to validate authority option. Error code: -8, message: Search for user name failed.
```

To resolve this issue, return to the **Authority Configuration** window, correct the value in the **Distinguished Name** field and attempt to distribute the authority configuration file again.

12. In the **Monitor Distribution Progress** window, review the progress of the configuration file distribution. Ensure that the authority configuration file distribution succeeds for all of the NetWorker servers.

13. Click **Finish**.

The next time you use a NMC client to connect to the NMC server, you must specify the appropriate LDAP or AD user. If you cannot log in to the NMC server you can revert back to native NMC authentication mode and reconfigure AD/LDAP authentication. The *NetWorker Installation Guide* provides more information.

Consider the following:

- ◆ During the authority file distribution process, the LDAP and AD authenticated NMC users that are granted the NMC Console Security Administrator role are added to the Security Administrators User Group on all NetWorker servers that they have the privilege to manage. Members of the Security Administrators User Group have permissions to modify the Audit Log server and User Group resources only. [“Modifying User Group privileges” on page 567](#) describes how to add a manually created LDAP or AD user to a User Group on a NetWorker server.
- ◆ When an LDAP or AD user logs in for the first time, a user object is automatically created on the NMC server for the user.
- ◆ The user name Administrator is not supported when LDAP or AD authentication is used.
- ◆ The NMC server cannot perform LDAP and AD administrative functions. Perform LDAP and AD administrative functions such as creating new domain users and groups with the appropriate LDAP and AD tools.
- ◆ The External Roles field for the Console Security Administrator role is not populated until an LDAP or AD user logs in for the first time.
- ◆ [“Troubleshooting login errors” on page 526](#) provides detailed information to troubleshoot common login error messages.

Adding or removing LDAP or AD console users

When an LDAP or AD user logs into the NMC server for the first time:

- ◆ A user object is automatically created.
- ◆ The user is assigned the same Console role that its LDAP or AD group is assigned to. [Table 88 on page 504](#) describes the three console roles.

Use one of the following methods to add or remove the LDAP or AD users on a NMC server:

- ◆ [“Adding LDAP or AD users by using the Configure Login Authentication Wizard” on page 518](#) — Use this method to add LDAP and AD users that require:
 - The Console Security Administrator role on the NMC server.
 - Membership to the Security Administrator User Groups on the managed NetWorker servers.
- ◆ [“Adding LDAP or AD users to the NMC server manually” on page 519](#) — Use this method to add LDAP or AD users to manage the NMC server but restrict NetWorker server access.

Adding LDAP or AD users by using the Configure Login Authentication Wizard

LDAP or AD users and groups that are added by using the Configure Login Authentication wizard are automatically:

- ◆ Assigned to the **Console Security Administrators** role on the NMC server.
- ◆ Added to the **Security Administrators User Groups** on the managed NetWorker servers.

To add LDAP or AD users and groups by using the **Configuration Login Authentication** wizard:

1. Log into the NMC server with a user that has the **Console Security Administrator** role.
2. From the **Console** window, click **Setup**.
3. From the **Setup** menu, select **Configure Login Authentication**.
4. In the **Select Authentication Method** window, select **External Repository**.
5. Select the appropriate LDAP or AD **Authority Name** and click **Next**.
6. In the **External Roles** field, specify the new LDAP or AD users and groups and click **Next**.
7. In the **Distribute Authority Configuration** window, select the NetWorker servers that have the **Requires Update** status and click **Distribute**.
8. In the **Monitor Distribution Progress** window, review the progress of the configuration file distribution. Ensure that the configuration file distribution succeeds for all NetWorker servers.
9. Log out of the NMC server and log in with a user account in the new group. [“Troubleshooting login errors” on page 526](#) describes how to troubleshoot login errors.

NOTICE

Members of the Security Administrators group have permissions to modify the **Audit Log** server and **User Group** resources only. [“Modifying User Group privileges” on page 567](#) describes how to add a manually created LDAP or AD user to a User Group on a NetWorker server.

Adding LDAP or AD users to the NMC server manually

Add LDAP or AD users to the NMC server manually when NetWorker server access restriction is not required. For example, a user or group manually assigned to the **Console Security Administrator** role is not automatically assigned to the **Security Administrators User Group** on the managed NetWorker servers. [“Modifying User Group privileges” on page 567](#) describes how to add a manually created LDAP or AD user to a **User Group** on a NetWorker server.

To manually add LDAP or AD users:

1. Log into the NMC server with a user that has the **Console Security Administrator** role.
2. From the **Console** window, click **Setup**.
3. In the left pane, right-click **Users**, then select **New**. The **Create User** dialog box appears, with the **Identity** tab displayed.
4. In the **User Name** attribute, enter the LDAP or AD user name.
5. Optionally, enter the full name of the LDAP or AD user and a general description in the remaining attributes.
6. Click **OK**.

[Figure 43 on page 519](#) provides an example of the **Create User** window.

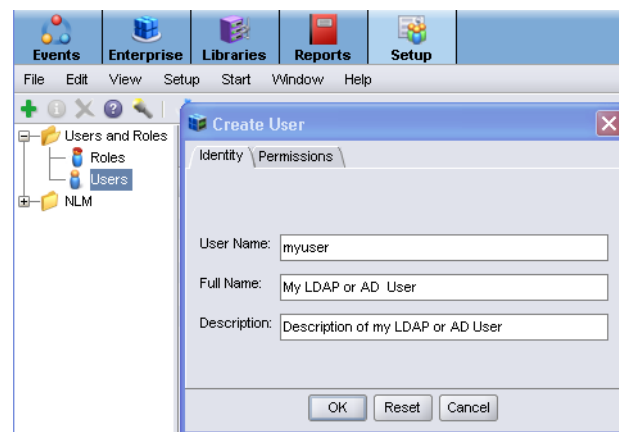


Figure 43 Create new LDAP or AD user manually

Modifying NMC roles for LDAP or AD users and groups

After configuring the NMC server to use LDAP or AD authentication, modify NMC console roles in the NMC console. “[Console roles](#)” on page 504 describes the available NMC console roles.

To modify NMC roles for LDAP or AD users:

1. Log into the NMC server with an account that has the **Console Security Administrator** role.
2. From the **Console** window, click **Setup**.
3. In the left pane, select **Roles**, then right-click the role and select **Properties**.
4. In the **External Roles** field, specify the LDAP group, username, or user role.
5. Click **Ok**.

Consider the following:

- ◆ If the user or group specified is not valid on the LDAP or AD server, the following message appears:

```
External role group_name is invalid
```

Modifying an LDAP or AD console user

Once you create an LDAP or AD user and assign it to an NMC console role, you can modify the descriptive information about the user in the NMC console.

To modify the descriptive information about an existing LDAP or AD user:

1. Log into the NMC server as a **Console Security Administrator**. The administrator account is a **Console Security Administrator**.
2. From the **Console** window, click **Setup**.
3. In the left pane, select **Users**.
4. Right-click the user and then select **Properties**.
5. In the **Identity** tab, modify the attributes as required.
6. Click **Ok**.

Deleting an LDAP or AD console user

Once you create an LDAP or AD user and assign NMC console roles to the user, you can delete the user in the NMC console.

To delete a user:

1. Log into the NMC server with a user that has the **Console Security Administrator** role.
2. From the **Console** window, click **Setup**.
3. In the left pane, click **Users**.
4. Right-click a username, then select **Delete**.
5. Click **Yes** to confirm the deletion.

6. If the user saved customized reports, a dialog box prompts for the username to which to reassign those reports. Otherwise, delete the reports.
7. Remove the user from the LDAP user role on the LDAP server, if required.

Troubleshooting authentication errors

This section provides a list of possible causes and resolutions for authentication configuration error messages:

- ◆ “Authority definition must specify external authority attribute name” on page 521
- ◆ “LDAP bind failed due to invalid credentials” on page 522
- ◆ “Failed to propagate external roles to NetWorker server” on page 522
- ◆ “No entry in hierarchy ‘ou=orgname, dc=domain_component1, dc=domain_component2 dc=domain_component3...’ on page 522
- ◆ “User Search Path hierarchy ou=orgname,dc=domain_component1,dc=domain_component2’ dc=domain_component3’ does not exist or is empty.” on page 523
- ◆ “No ldap search path for usernames” on page 523
- ◆ “Group Search Path hierarchy ou=orgname,dc=domain_component1,dc=domain_component2’ dc=domain_component3’ does not exist or is empty.” on page 523
- ◆ “Error querying for user groups” on page 523
- ◆ “LDAP bind failed because the server is down” on page 523
- ◆ “networker_server (Permission denied, user 'LDAP_user' on 'NMC_server' does not have 'Configure NetWorker' OR 'Change Application Settings' privilege to configure this resource) - NSR.” on page 523
- ◆ “Failed to retrieve authentication control attributes from NetWorker server [NetWorker_server]” on page 524
- ◆ “Could not validate external authority Failed to get status of file (clientCertificate) 'full_path_to_client_certificate': No such file or directory. Provide valid path or copy the certificates/key to the specified path” on page 524
- ◆ “NSR Could not validate external authority LDAP bind failed because the server is down” on page 525
- ◆ “You do not have privileges to use NetWorker Management Console” on page 526
- ◆ “Could not authenticate this user name and password, try again!” on page 526
- ◆ “The specified user name is restricted and cannot be used to log into the system” on page 526

Authority definition must specify external authority attribute name

Appears in the **Configure Login Authentication** wizard when the **Authority Name** field is blank.

LDAP bind failed due to invalid credentials

Appears in the **Configure Login Authentication** wizard when:

- ◆ The LDAP or AD user specified in **Distinguished Name** field is incorrect.
- ◆ The password specified for the LDAP or AD user is incorrect.

Failed to propagate external roles to NetWorker server

Appears when the distribution of the authority file fails for a NetWorker server. Distribution fails because the NMC user used to distribute the file is not a member of the **Application Administrators User Group** on the NetWorker server.

To resolve this issue:

1. Close the **Configure Login Authentication** wizard.
2. Connect to the NetWorker server with a NMC user that is a member of the **Security Administrators User Group**.
3. Add the appropriate LDAP or AD group to the **Application Administrators User Group**.
4. Launch the **Configure Login Authentication** wizard and configure the new LDAP or AD authority.

No entry in hierarchy ‘ou=orgname, dc=domain_component1, dc=domain_component2 dc=domain_component3...

These error messages appear in the **Configure Login Authentication** window when the attribute value referenced in the error message is incorrect or the attribute value cannot be validated to the LDAP or AD authority. [Table 90 on page 522](#) describes the messages that are displayed and the attribute to correct.

Table 90 Hierarchy errors in the Configure Login Authentication wizard

No entry in hierarchy ‘ou= <i>orgname</i> , dc= <i>domain_component1</i> , dc= <i>domain_component2</i> dc= <i>domain_component3</i> ...	This error message appears in the Configure Login Authentication wizard when the value defined...
...belongs to user object class ‘ <i>user_object_class</i> ’	...in the User Object Class attribute is not valid for the value defined in User Search Path attribute.
...has a group name attribute ‘ <i>groupname</i> ’	...in the Group Name Attribute field is not valid on the LDAP or AD server.
...has a user id attribute ‘ <i>user_id</i> ’	...in the User ID Attribute field is not valid on the LDAP or AD server.
...belongs to object class ‘ <i>group_object_class</i> ’	...in the Group Object Class field is not valid on the LDAP or AD server.
...belongs to object class ‘ <i>group_object_class</i> ’	...in the Group Object Class field is not valid on the LDAP or AD server
...has a group member attribute ‘ <i>group_member_attribute</i> ’	...in the Group Member Attribute field is not valid on the LDAP or AD server.

User Search Path hierarchy *ou=orgname,dc=domain_component1,dc=domain_component2,dc=domain_component3* does not exist or is empty.

Appears in the Configure Login Authentication wizard when the value defined in the User Search Path attribute is not valid on the LDAP or AD server.

No ldap search path for usernames

Appears in the Configure Login Authentication wizard when the value defined in the User Search Path attribute is not valid on the LDAP or AD server.

Group Search Path hierarchy *ou=orgname,dc=domain_component1,dc=domain_component2,dc=domain_component3* does not exist or is empty.

Appears in the Configure Login Authentication wizard when the value defined in the Group Search Path attribute is not valid on the LDAP or AD server.

Error querying for user groups

Appears in the Configure Login Authentication wizard when the value defined in the Group Search Path attribute is not valid on the LDAP or AD server.

LDAP bind failed because the server is down

Appears in the Configure Login Authentication wizard when:

- ◆ The Port Number defined for the LDAP, LDAPS, or AD server is incorrect.
- ◆ The hostname specified in the Provider Server Name field is incorrect or the hostname is not resolvable.
- ◆ When the LDAPS server requires a certificate but the Server certificate file or Client certificate file field is empty.

networker_server (Permission denied, user '*LDAP_user*' on '*NMC_server*' does not have 'Configure NetWorker' OR 'Change Application Settings' privilege to configure this resource) - NSR.

This error message appears in two scenarios:

- ◆ While distributing the authority configuration file to a new NetWorker server, the new NetWorker server cannot authenticate the LDAP user account.

To resolve this issue, configure the NMC server to use Native NetWorker Management Console authentication and then reconfigure the LDAP or AD authorities and distribute them to all the required servers.

For example:

1. In the **Distribute Authority Configuration File** window, click **Finish**.
2. Start the **Configure Login Authentication** wizard again.
3. In the **Select Authentication Method** window, click **Next**.
4. Record the values in each attribute field for the configured LDAP or AD authorities; click **Back**.
5. In the **Select Authentication Method** window, select **Native NetWorker Management Console**; click **Next**.
6. Select all servers with a status **Requires Update**; click **Distribute**.

7. Click **Finish**.
 8. Start the **Configure Login Authentication** wizard again and create the LDAP or AD authorities. [“Configuring the NMC server for LDAP, LDAPS, or AD authentication” on page 508](#) provides more information.
- ◆ When an LDAP or AD user tries to modify the Server resource (NSR) on a NetWorker server but the user is not a member of the **Application Administrators** or the **Security Administrators User Group**.

To resolve this issue:

1. Close the NetWorker server and NMC server browser windows.
2. Log in to the NMC server with an LDAP or AD account that is a member of the **Application Administrators** or the **Security Administrators User Group**.

Failed to retrieve authentication control attributes from NetWorker server [*NetWorker_server*]

Appears when an LDAP or AD user that is not a member of the **Security Administrators** User Group on the NetWorker server attempts to distribute the authority configuration file to the NetWorker server.

To resolve this issue:

1. In the **Distribute Authority Configuration File** window, click **Finish**.
2. Close the NMC server browser window.
3. Log in to the NMC server with an LDAP or AD user that is a member of the **Security Administrators** User Group on the NetWorker server. LDAP or AD users that have the **Console Security Administrator** role on the NMC server are a member of the **Security Administrators** User Group on the NetWorker server, by default.

NOTICE

Members of the **Security Administrators** User Group on a NetWorker server only have permissions to modify the **Audit Log** server and User Group resources. [“Modifying User Group privileges” on page 567](#) describes how to modify the **User Group** membership on a NetWorker server.

Could not validate external authority Failed to get status of file (clientCertificate) '*full_path_to_client_certificate*': No such file or directory. Provide valid path or copy the certificates/key to the specified path

This message appears when the Wizard attempts to distribute the authority configuration file to the NetWorker server but the paths that you specified for the certificate files are incorrect.

To resolve this issue:

1. In the **Distribute Authority Configuration File** window, click **Finish**.
2. Start the **Configure Login Authentication** wizard again.
3. In the **Select Authentication Method** window, click **Next**.
4. Correct the pathnames in the certificate fields and retry the distribution.

Note: For Windows paths, use double slashes in the path. For example, `c:\\my_ldap_server`.

NSR Could not validate external authority LDAP bind failed because the server is down

This message appears when there is an issue with the LDAPS certificate.

To troubleshoot LDAPS certificate issues, use the **openssl** command:

Note: A third-party provider provides an **openssl** binary for Windows. <http://www.openssl.org/related/binaries.html> provides more information.

1. Confirm that you can establish a SSL connection to the LDAPS server using the local copy of the certificate files:

```
openssl s_client -connect ldaps_server_name:ssl_port -CAfile
full_path_to_server_certificate -cert full_path_to_client_certificate -key
full_path_to_client_key_file
```

Where:

- *full_path_to_certificate* is the full path to the Server certificate file on the local host. If the environment has a hierarchy of CA authorities, then specify the root CA or the certificate file that contains all CA authority certificates.
- *full_path_to_client_certificate_file* specifies the full path to the Client Certificate file on the local host. This option is only required when LDAPS requires a client certificate.
- *full_path_to_client_key_file* specifies the full path to the Client Certificate file on the local host. This option is only required when LDAPS requires a client key.

Example 47 The LDAPS server, myldaps.emc.com requires a CA certificate only. The certificate file, ca.cert resides in the cst directory of a NMC server on Windows. In this example, type the following command:

```
openssl s_client -connect myldaps.emc.com:636 -CAfile
"C:\Program Files\EMC NetWorker\Management\GST\cst\ca.cert"
```

Note: When the connection succeeds, the command returns the message **Verify return code: 0 (ok)**

Example 48 The LDAPS server, myldaps.emc.com requires a Client Certificate and a Client Key. The certificate files and the key file resides in the cst directory of a NMC server on Windows. In this example, type the following command:

```
openssl s_client -connect myldaps.emc.com:636 -CAfile
"C:\Program Files\EMC NetWorker\Management\GST\cst\ca.cert" -cert
"C:\Program Files\EMC NetWorker\Management\GST\cst\client.cert"
-key "C:\Program Files\EMC
NetWorker\Management\GST\cst\client.key"
```

Note: When the connection succeeds, the command returns the message **Verify return code: 0 (ok)**

2. If the connection does not succeed, contact the LDAPS administrator to request a new copies of the certificate files. To manually copy the CA certificate file from the LDAP server, you can perform the following step:

- Connect to the LDAPS server to display the Server Certificate (ca.cert) file:

```
openssl s_client -showcerts -connect ldaps_server_name:ssl_port
```

The **openssl** command may display two certificates. The second certificate is usually the CA certificate.

Troubleshooting login errors

This section provides a list of possible causes and resolutions for Console login error messages:

- ◆ [“You do not have privileges to use NetWorker Management Console” on page 526](#)
- ◆ [“Could not authenticate this user name and password, try again!” on page 526](#)
- ◆ [“The specified user name is restricted and cannot be used to log into the system” on page 526](#)

You do not have privileges to use NetWorker Management Console

Appears when a valid LDAP or AD account tries to log into the NMC server but the account does not exist on the NMC server or is not assigned a Console role.

To resolve this issue, create the LDAP or AD account manually and try to log in again. [“Adding LDAP or AD users to the NMC server manually” on page 519](#) describes how to create LDAP and AD user accounts manually.

Could not authenticate this user name and password, try again!

Appears when you attempt to log into the NMC server with:

- ◆ A username that is not recognized or an incorrect password.
To resolve this issue, use the correct user name and password combination for the configured NMC server authentication method.
- ◆ An AD user that has the option **User must change password at next login** enabled.
To resolve this issue, change the password before attempting to log in to the NMC server.

The specified user name is restricted and cannot be used to log into the system

Appears when you use the username Administrator to log into the NMC server and the NMC server authentication is LDAP or AD. A NMC server that uses AD or LDAP authentication does not support the administrator username.

To resolve this issue, log into the NMC server with a different LDAP or AD username.

Restricting a user’s view of managed servers

By default the NMC server adds members of the **Console Security Administrators** to the **Security Administrators** user group on each NetWorker server that the NMC server manages.

To restrict the NetWorker servers that a user can view and manage, modify the privileges on the user object. [“Implications of restricting an NMC, LDAP, or AD user’s view of managed servers” on page 527](#) provides information about the implications of restricting user views. [“Restricting report views” on page 434](#) provides information on how restricting views affects reporting.

NMC, LDAP or AD users can only manage data on a NetWorker servers that the user can view. To perform operations on the NetWorker server, grant NMC, LDAP or AD users explicit privileges on the NetWorker server. [“Managing server access” on page 558](#) describes how to configure privileges for NMC, LDAP, or AD users on a NetWorker server.

To restrict the NetWorker server that an NMC, LDAP or AD user can view and manage:

1. Log in with an account that has the **Console Security Administrator** role.
2. From the **Console** window, click **Setup**.
3. In the left pane, click **Users**.
4. Right-click a user, then select **Permissions**. The **Edit User** window appears and the **Permissions** tab displays.
5. To grant the user privileges to view various hosts, use the arrow keys to select the allowed hosts.
6. Click **OK**.

Implications of restricting an NMC, LDAP, or AD user’s view of managed servers

The effects of restricting user views for various functions are as follows:

- ◆ In the **Events** window: The user sees only events from allowed NetWorker servers.
- ◆ In the **Enterprise** window: The user sees all the hierarchy folders, but only the allowed NetWorker servers appear in those folders.
- ◆ In the **Libraries** window: The user sees only the devices controlled by allowed NetWorker servers.
- ◆ In the **Reports** window: The user sees report data only from allowed NetWorker servers.
- ◆ In the **Setup** window:
 - The user sees properties for all users, and in addition to its own properties and privileges.
 - The user can modify its own properties, but not privileges. Only the **Console Security Administrator** can view and modify user privileges.

Because each user can view and manage different sets of NetWorker servers, the contents of the reports can vary among users. For example, a shared backup summary report entitled “Building C Backups” will display different data for different users (even if run you the report simultaneously) when the users’ privileges include different NetWorker servers. This applies to all report types, whether default or customized, private or shared.

If no data is available for a particular server, that server will not appear in any lists or reports, regardless of the user privileges.

Resetting the administrator password (native NMC authentication only)

Reset a lost or forgotten administrator password to the default value by using the `GST_RESET_PW` environment variable.

Microsoft Windows

To reset the native Console security administrator password on Windows:

1. From the **Control Panel** program, create the `GST_RESET_PW` environment variable with a value of `1`.
2. Restart the **EMC GST Service**. When the **EMC GST Service** starts, the NMC server administrator password resets.
3. Log into the NMC server specify **administrator** as the username and the password.
4. Return to the **Environment Variables** window and remove the `GST_RESET_PW` environment variable. This step prevents a password reset each time the **EMC GST Service** starts.

UNIX systems

To reset the native administrator password on a UNIX system:

1. Set `GST_RESET_PW` to a non-null value using the appropriate command for the shell, for example, in ksh:

```
export GST_RESET_PW "non_null_value"
```

2. Stop and restart the NMC server in the same shell. [Table 8 on page 53](#) provides more details.
3. Log into the NMC server and type **administrator** as the username and the password.
4. Set `GST_RESET_PW` back to `null` by using the appropriate command for the shell, for example, in ksh:

```
export GST_RESET_PW=
```

The next time the NMC server restarts, the password is not reset again.

Moving the NMC server

You can move a NMC server from one host to another only if both hosts use the same operating system. Reasons to move the NMC server include:

- ◆ The current host has insufficient processing capabilities. For example, if there is a need for more memory or a faster processor.
- ◆ The current host has insufficient space for the Console database.
- ◆ The current host was damaged beyond repair.

To move the NMC server component to another host:

1. Backup the existing Console database with the `savepsm` command at level **Full**. [“Performing a manual backup of the Console database” on page 111](#) provides more details.

2. Setup the new host with the same operating system as the host that runs the current software and connects to the network.
3. On the new host, install the NetWorker Client software and the NetWorker NMC server component.

NOTICE

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.

4. On the remote NetWorker server, setup the new target host as a client resource. [“Task 6: Create a backup Client resource” on page 64](#) describes how to create a Client resource.
5. For the Client resource of the source NMC server, add the appropriate users to the **Remote Access** attribute on the **Globals (2 of 2)** tab.
6. Stop the NMC server service on the source NMC server. [Table 8 on page 53](#) provides more information.
7. Stop the NMC server service on the target NMC server.
8. Ensure you stop the GST service on the original backup host or on the host used for the directed recover. Use the appropriate command or tools for the operating system to ensure the GST service is not running.
9. For UNIX systems, update the library path environment variable with the NMC server library path.

For example:

- On AIX, update the LIBPATH variable with *Console_install_dir/sybase/lib64* where *Console_install_dir* is */opt/LGTONMC* by default.
- On Linux and Solaris, update the LD_LIBRARY_PATH variable with *Console_install_dir/sybase/lib* where *Console_install_dir* is */opt/LGTONMC* by default.

NOTICE

If you did not install the NMC server software in the default location, add */Console_install_dir/bin* to the library path environment variable.

10. On the NetWorker server, you must have:
 - Client resources for the original backup host and the directed recover host.
 - The root or administrator account of the directed recover host listed in the **Remote Access** field of the client resource for the original backup host.
11. Run the **recoverpsm** command on the recovery host:

```
recoverpsm [-f] [-d recover_directory] -s NetWorker_server -c original_console_server -S gst_on_original_console_server -0
```

For example,

- To recover the NMC server database and **gstdb.conf** credential file to the original backup location on the directed recover host, type this command:

```
recoverpsm -f -c original_console_server
```

- To recover the NMC server database and **gstdb.conf** credential file, to a directory on the directed recover host that differs from the original backup location type this command:

```
recoverpsm -c original_console_server -f -d recover_directory
```

[Table 91 on page 530](#) provides a description of **recoverpsm** options. The *EMC NetWorker 8.0 Command Reference Guide* or the UNIX man pages provide a complete description of the **recoverpsm** command line options.

Table 91 The recoverpsm options

Option	Meaning
-f	Instructs the software to overwrite existing console database files.
-d <i>recover_directory</i>	Specifies the destination directory for the recovered Console database and if you did not specify the -O option, the destination directory of the <i>gstd_db.conf</i> configuration file. Include a full path for the directory. The recover process does not support partial paths. When <i>the recover_directory</i> is not the same as the console database directory, copy the database file and the <i>gstd_db.conf</i> file, to the directory where the console database resides after the recover completes.
-S <i>gst_on_source_Console_server</i>	Specifies the existing console database on the source NMC server. You must prepend gst_on_ to the short name of the source NMC server.
-O	Do not recover the database credential file, <i>gstd_db.conf</i> . Use this option to preserve preexisting database login credentials on recovery server. Use the recover program or the NetWorker User program to recover the <i>gstd_db.conf</i> configuration file to a new location.
-s <i>NetWorker_server</i>	Specifies the name of the NetWorker server where the console database backup resides.
-c <i>source_console_server</i>	Specifies the short name of the source NMC server, where the existing console database resides.

12. If you 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.
13. Start the NMC server program on the target NMC server. [Table 8 on page 53](#) provides detailed instructions.
14. If you are using the License Manager and the License Manager host has changed, obtain a Host Transfer Affidavit from EMC support. Use the host ID of the License Manager host for the new authorization codes. If the License Manager host has not changed, then new authorization codes are unnecessary.

Setting system options

The NMC server includes several options that affect performance. These options enable users to fine-tune the performance of the NMC server. Only Console administrators can set or change system options.

Setting a system option

To set a system option:

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. [Table 92 on page 531](#) provides a description of the available system options.
4. Click **OK**.

NOTICE

Do not adjust these system options without careful consideration. A mistake in setting system options can seriously degrade performance.

Table 92 NMC server system options (1 of 2)

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 gstd.raw file. Increase this value to troubleshoot only.
Polling interval for events and reporting (seconds)	Default value: 20 Range: 0-20 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: 0-20 Defines the frequency in which the NMC server contacts the managed NetWorker servers for activity updates.
Polling thread factor	Default value: 5 Range: 0-20 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.
Polling interval for NetWorker libraries (hours)	Default value: 12 Range: 0-20 Defines how frequently the Console GUI polls the Libraries defined for a NetWorker server to gather information. This information appears in Libraries task of the main Console GUI window.

Table 92 NMC server system options (2 of 2)

System option	Description
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: disabled 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. <ul style="list-style-type: none"> When enabled, the Console username determines the Console user access. “Individual User Authentication” on page 532 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: enabled Before the NMC server connects to a managed NetWorker server, the NMC server confirms that the NetWorker server daemons are running. <ul style="list-style-type: none"> 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.

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 is 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 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 servers to which 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. [Appendix B, “Firewall Support”](#) provides information about firewalls.

When you set the User Authentication for NetWorker system option, the NMC server software creates a separate network connection the NMC server 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 NMC server to the managed NetWorker server.

Setting environment variables

[Table 93 on page 534](#) describes the NMC server environment variables available in this NetWorker release. The environment variable values are incorporated only when the NMC server checks for environment variables at start, similar to how the `nsrd` daemon handles its corresponding variables.

Setting environment variables on UNIX

Table 93 Console environment variables

Environment variable	Description
GST_DEBUG	Range: 1–20 Default: 0 Use this environment variable to troubleshoot the product when the Console GUI is not accessible. This environment variable: <ul style="list-style-type: none"> • Defines the level of debug information written to the gstd.raw log file. • Overrides the Debug Level attribute defined in System Options. Setting the value higher increases: <ul style="list-style-type: none"> • The number of operation and status messages that the NMC server records. • The size of the gstd.raw file.
GST_RESET_PW	Resets the default administrator password for native NMC authentication. “Resetting the administrator password (native NMC authentication only)” on page 528 provides detailed information.
GST_RESET_DBPWD	Changes the database connection credentials. “Changing database connection credentials” on page 536 provides detailed information.

To set environment variables in the gstd file, a Bourne shell script on UNIX systems:

1. Change the file permissions. By default, the gstd file is a read-only file. The file location varies depending on the operating system:
 - Solaris and Linux: /etc/init.d/gstd
 - AIX: /etc/rc.gstd
2. Open the file in a text editor and add these lines to the beginning of the file:


```
variable_name=variable_value
export variable_name
```

where:

 - *variable_name* is the name of the environment variable.
 - *variable_value* is the value to assign to the environment variable.
3. Save the changes.
4. Stop and restart the NMC server (gstd) program. [Table 8 on page 53](#) provides detailed instructions.

The preceding example is for the Bourne shell. For other shells, refer to the shell-specific documentation or man pages that describe how to set an environment variable.

Setting environment variables on Windows systems

On Microsoft Windows, set environment variables by using the Control Panel System applet on the NetWorker server:

1. Browse to **Control Panel -> System and Security -> System -> Advanced System Settings**.
2. In the **General** tab click **Environment Variables...**
3. Click the **New** button.
4. Specify the name and value of the environment variable. Refer to [Table 93 on page 534](#) for a list of GST environment variables.
5. Stop and start the NetWorker Remote Exec and EMC gstd services in order for the environment variables to take effect.

Accessing the Console Configuration Wizard

NOTICE

Only Console Application Administrators can use the Console Configuration Wizard.

Perform the following tasks from the Console menu options or from the Console Configuration Wizard.

- ◆ Set the NetWorker Administrator password.
- ◆ Specify the NetWorker server to backup the Console database.
- ◆ Add NetWorker servers to the enterprise.

To access the Console Configuration Wizard:

1. From the **Console** window, click **Setup**.
2. From the **Setup** menu, select **Configuration Wizard**.

NetWorker NMC server maintenance tasks

This section lists maintenance tasks associated with the NetWorker NMC server.

Changing the service port used by the NetWorker Console database

By default, the NetWorker Console database uses port 2638 for TCP/IP communications. Other applications may also use this port or, as is the default installation of case with EMC AlphaStor or EMC DPA. Each application may have its own instance of the iAnywhere database installed. Having the NMC server installed on a host with one of these applications, creates a conflict.

To resolve this problem, change the service port as follows:

1. Stop the GST Service. [Table 8 on page 53](#) has information about stopping the GST service.
2. Stop the service of the other product using the iAnywhere database.
3. Ensure that no dbsrv12 processes are running.

- From a terminal or command prompt window, use the appropriate command (for example, **setenv** for csh, **export** for sh) to update the library path environment variable with the following location:

- Solaris: /opt/LGTONMC/bin:/opt/LGTONMC/sybase/lib
- Linux/AIX: /opt/lgtonmc/bin:/opt/lgtonmc/sybase/lib
- Microsoft Windows (assumes default installation location):
C:\Program Files\EMC NetWorker\Management\GST\sybase\bin

The environment variable to set varies by operating system, for example:

- Solaris/Linux: LD_LIBRARY_PATH
- AIX: LIBPATH

- Edit the `gstd.conf` file to add or change the following line:

```
db_svc_port=port_number
```

For example:

```
db_svc_port=2639
```

The following locations contain the `gstd.conf`:

- Solaris: /opt/LGTONMC/etc
 - Linux and AIX: /opt/lgtonmc/etc
 - Microsoft Windows: C:\Program Files\EMC NetWorker\Management\GST\etc
- Run the **gstconfig** command to update the port value in the NetWorker NMC server configuration file. The following locations contain the **gstconfig** command:
 - Solaris: /opt/LGTONMC/bin
 - Linux/AIX: /opt/lgtonmc/bin
 - Windows: C:\Program Files\EMC NetWorker\Management\GST\bin
 - Close the terminal or command prompt window.
 - Start the EMC GST Service.

Changing database connection credentials

When the NetWorker NMC server starts for the first time, it automatically generates the login credentials used to log into 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.

To force the server to change the credentials:

- Stop the GST Service. [Table 8 on page 53](#) has information about stopping the GST service.
- Set the environment variable `GST_RESET_DBPWD` to any value. On Microsoft Windows system, set this value as a System Variable. Reboot the system after setting the variable.

3. Restart the GST Service.
4. Delete the GST_RESET_DBPWD environment variable. On Microsoft Windows system, reboot the machine after setting the variable.

NMC server IP address/hostname updates

If you modify the IP address or hostname of the NMC server or if you add or remove protocols such as IPv6, perform the following:

1. Stop the GST Service. [Table 8 on page 53](#) has information about stopping the GST service.
2. Browse to the NetWorker bin directory then run the platform-specific commands:
 - On Windows, run `gstconfig` in the `C:\Program Files\EMC NetWorker\Management\GST\bin` directory.
 - On Solaris, as root:
 - a. Define and export the `LD_LIBRARY_PATH` variable:


```
LD_LIBRARY_PATH=/opt/LGTONMC/bin:/opt/LGTONMC/sybase/lib
export LD_LIBRARY_PATH
```
 - b. Update the GST configuration:


```
./gstconfig
```
 - On Linux, as root:
 - a. Define and export the `LD_LIBRARY_PATH` variable:


```
LD_LIBRARY_PATH=/opt/lgtonmc/bin:/opt/lgtonmc/sybase/lib
export LD_LIBRARY_PATH
```
 - b. Update the GST configuration:


```
./gstconfig
```
 - On AIX:
 - a. Define and export the `LIB_PATH` variable:


```
LIBPATH=/opt/lgtonmc/bin:/opt/lgtonmc/sybase/lib
export LIBPATH
```
 - b. Update the GST configuration:


```
./gstconfig
```
3. Restart the EMC GST daemon.
4. For NMC hostname changes only, delete the client resource created to perform NMC server database backups and create a new client. [“Scheduling backups for the Console server database” on page 110](#) describes how to create and configure a client for console database backups.

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 EMC Legato NetWorker Installation Guide provides more information about displaying international fonts on a NMC server that is operating 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 enterpriseNetWorker from a single host.

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 Installation and Administration Guide* provides more information on how to install and use the NetWorker License Manager.

Entering an enabler code

To add an enabler code when you use NetWorker License Manager:

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

To delete an enabler code if you are using LLM:

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

To enter an authorization code if you are using LLM:

1. From the **Console** window, click **Setup** and then click **Licensing**.
2. Right-click the license to be authorized, then select **Properties**. The **Properties** dialog box appears.
3. 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).
4. Click **OK**. The license is now permanently enabled.

Changing the License Manager server

The License Manager server that manages NetWorker Console licenses can be changed at any time.

To change which host runs the License Manager:

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**.

CHAPTER 18

NetWorker Server Management

This chapter covers these topics:

◆ Enterprise	542
◆ Configuring a NetWorker server	550
◆ Report home	551
◆ Parallelism and multiplexing	553
◆ Managing server access	558
◆ Working with the Multi-Tenancy Facility	569
◆ Server communication issues within Microsoft Windows	585
◆ Indexes	585
◆ Managing Client Push	594
◆ Monitoring Changes to NetWorker Server Resources	599
◆ Log file size management	600
◆ Internationalization	602

Enterprise

The Enterprise is a visual representation of the NetWorker Console control zone. Various servers in the enterprise such as NetWorker, Data Domain, and Avamar servers can be monitored for events, and various reports can be generated on events, backups, and user activity.

For more information, see:

- ◆ [“Enterprise components” on page 542](#)
- ◆ [“Organizing NetWorker servers” on page 542](#)
- ◆ [“Viewing the enterprise” on page 543](#)
- ◆ [“Managing various servers in the enterprise” on page 544](#)

Enterprise components

The Enterprise includes these components:

- ◆ Hosts

A host, also known as a *managed node*, is the NetWorker, Data Domain, or Avamar 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, thus putting all the hosts from the same city or country in the same folder. [Example 49 on page 542](#) provides information.
- ◆ By function of the computers backed up by the NetWorker servers, such as having the servers that back up web servers in one folder, and the servers that back up mail servers in another.
- ◆ By administrative divisions within the Enterprise, such as having separate folders for servers that back up Marketing, Sales, or Engineering computers.

Since there can be multiple copies of a host in the Enterprise, multiple folders can also be created and maintained. With each folder based on different organizational criteria, the organization can be viewed in different, yet parallel and complementary ways.

Example 49 An enterprise arranged by geographic location

[Figure 44 on page 543](#) shows an Enterprise arranged by geographic location. There are three folders, one for each country where the NetWorker servers are located: USA, France, and Australia. Each folder contains a number of hosts that correspond to NetWorker servers named for the city where they are located. The Australia folder, for instance, contains three host computers labeled *perth1*, *perth2*, and *sydney*.

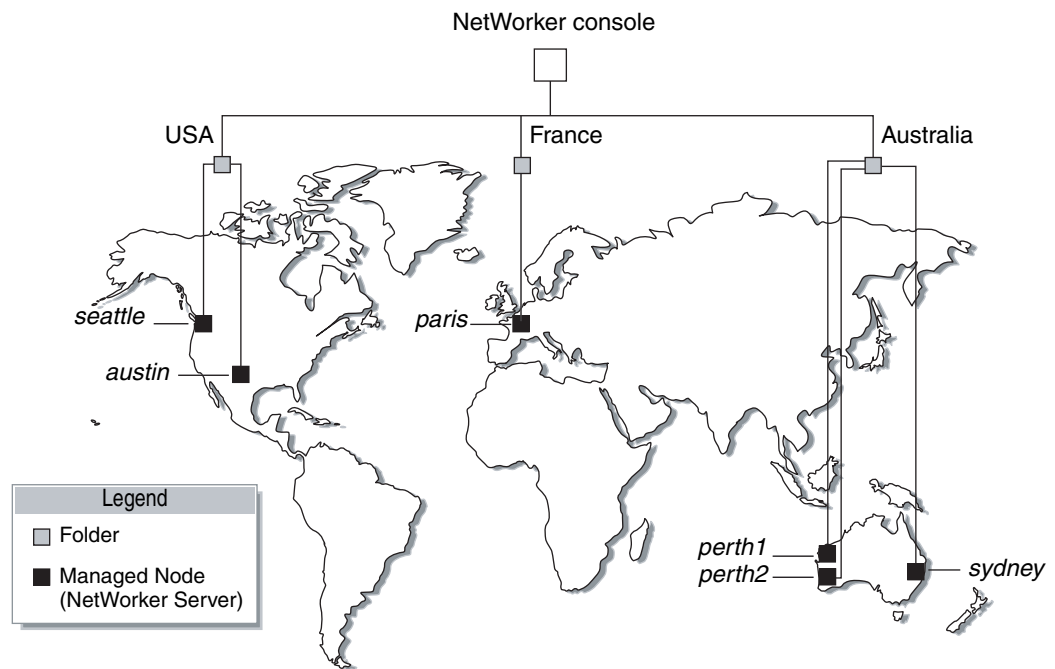


Figure 44 NetWorker servers worldwide

Viewing the enterprise

In the Console window, the organization of the NetWorker servers can be viewed in much the same way as the contents of a file system can be viewed by using a file manager program.

To view the Enterprise:

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 [Table 94 on page 543](#).

Table 94 Viewing the enterprise

To:	Do this:
Show or hide contents of the Enterprise.	Click Enterprise.
Show or hide contents of a folder.	Click the folder.
Show the managed applications installed on a host computer.	Click the host.

Managing various servers in the enterprise

NetWorker Console enables centralized management of NetWorker, Data Domain, and Avamar 3.7 and later servers within the Enterprise. You can add, delete, move, and copy servers as needed. All of these functions can be performed through the Console window. When using NetWorker software to manage a large number of NetWorker servers, it might be more efficient to add or delete multiple hosts by using a single command from the command prompt. [“Adding or deleting multiple servers by using a hostname file” on page 548](#) provides further information.

The server management activities include, but are not limited to, operations related to devices and libraries, and events that require user intervention.

Adding a managed host

The Console window can display server events, which allow server activity-reports to be generated.

NOTICE

Data Domain servers are added as managed hosts automatically when a Data Domain device is configured with the new New Device Wizard. The *EMC NetWorker Data Domain Devices Integration Guide* provides more information about Data Domain as a managed host.

NetWorker Console supports hosts that use NetWorker server software release 7.2.x or later and Avamar release 3.7.2 or later.

To add one of these hosts to the Enterprise:

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. Enter a hostname, IP address, DNS name, or WINS name in the **Host Name** attribute, then click **Next**.

NOTICE

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**.

A host can also be added by using the **Console Configuration Wizard**. [“Accessing the Console Configuration Wizard” on page 535](#) provides information.

Deleting a host

You can delete a single host or multiple hosts within a folder.

To delete a host or hosts:

1. From the **Console** window, click **Enterprise**.
2. Right-click the host to delete, 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

Multiple copies of a host can be created for a single NetWorker or Avamar server. For example, one copy of a host can be in its logical position in the Enterprise, while another copy of the host is in a Hosts-to-Watch folder where it can easily be monitored. This makes it possible to check the server without navigating through the Enterprise.

To copy a host:

1. From the **Console** window, click **Enterprise**.
2. Right-click the host to copy, then select **Copy**.
3. Right-click a new location, then select **Paste**.

NOTICE

Also, the drag-and-drop feature can be used while holding down the **Ctrl** key to copy hosts.

Moving a host

To move a host from one location to another in an Enterprise:

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**.

NOTICE

Also, the drag-and-drop feature can be used to move hosts.

Managing folders in the enterprise

NetWorker software allows you to manage folders within the Enterprise. This means that folders can be added, renamed, deleted, and moved as needed.

New folders can be added directly beneath the Enterprise node or beneath other folders.

Adding a folder

To add a folder:

1. From the **Console** window, click **Enterprise**.
2. Right-click the location within the Enterprise where the new folder is 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

NOTICE

If there are restrictions in place controlling which hosts a user is allowed to see, the folder might appear empty.

To delete a folder:

1. From the **Console** window, click **Enterprise**.
2. Right-click the folder to delete, then select **Delete**.
 - If hosts are present, 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 no hosts are present, the folder is deleted.

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**.

Non-unique hosts, and folders containing *only* non-unique hosts, are deleted without additional prompting.

Copying a folder

To copy a folder:

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** and a copy of the folder appears in its new location.

NOTICE

Also use the drag-and-drop feature can be used to copy folders while holding down the **Ctrl** key.

4. A folder cannot be copied within the same Enterprise level.

Moving a folder

To move a folder:

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

Also, the drag-and-drop feature can be used to move folders.

Renaming a folder

To rename a folder:

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, it may be more convenient to add or delete multiple NetWorker servers by using the **gstmodconf** command and a hostname file. With this method, hosts are added or deleted at the base level of the Enterprise. The hosts are added as NetWorker managed nodes with the features Capture Events and Gather Reporting Data enabled. [“Using the gstmodconf command” on page 549](#) has more information about the **gstmodconf** command.

Restrictions

Copies of hosts cannot be added with the **gstmodconf** command. If a host already exists anywhere in the Enterprise (either at the base or within a folder), copies of it cannot be added by this command.

It is not possible to use this command to add a host to a folder. It can only add a host to the base level. After the host has been added, the graphical user interface can be used to move the host to a selected folder. [“Moving a host” on page 545](#) provides information about this procedure.

When the **gstmodconf** command is used for deletion, it deletes hosts from the base level. It does *not* delete hosts that are within folders.

Creating the hostname file

To add or delete multiple hosts at the same time, specify their names in a hostname file. The hostname file is a simple text file.

To create a hostname file, use these guidelines:

- ◆ Only one hostname may be listed 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 [Example 50 on page 548](#):

Example 50 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 -p port -l login -P password
```

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide a complete description of the command and its options.

[Example 51 on page 549](#) shows how `gstmodconf` is used to add nodes from the file, `xyz_hostlist`. In this example, the Console server name is `myconsole` and the `xyz_hostlist` file contains:

Example 51 Adding multiple hosts with the `gstmodconf` command

```
apple
banana
grape
% gstmodconf -s myconsole -i xyz_hostlist
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
```

Error messages generated by the `gstmodconf` command

[Example 52 on page 549](#) shows the error that is generated if `gstmodconf` is used to add a host that already exists in the Enterprise:

Example 52 Attempting 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...
```

[Example 53 on page 549](#) shows the error that is generated if the `gstmodconf` command is entered *without* specifying the administrator password, when the password has been changed from its default value.

Example 53 Attempting to Use `gstmodconf` with an expired default password

```
% gstmodconf -s myconsole -i xyz_hostlist
Trying 111.22.3.444... auth failed.
gt_session_connect: c1nt_create: Remote system error-Connection
refused.
```

Configuring a NetWorker server

To use the Administration window to modify a NetWorker server, the server must meet the following requirements:

- ◆ The server must be included in the Enterprise.
 - “[Adding a managed host](#)” on page 544 provides information about adding hosts to the Enterprise.
- ◆ The server must not restrict users from viewing the NetWorker server.

The Administration window can be used to modify an existing NetWorker server by specifying attributes in the Properties dialog box of the NetWorker server. The configuration of these attributes, such as setting the parallelism level or designating administrator privileges, can affect backup performance and security.

Set up the server

NOTICE

When setting up the NetWorker server, be sure to enter the NetWorker product serial number. The product serial number is located on the Enabler Certificate that was sent when the NetWorker product was ordered.

To set up the NetWorker server:

1. From the **Administration** window, click **Configuration**.
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**.

Licensing the NetWorker server

The NetWorker *Licensing Guide* provides information on 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 a job will be declared inactive and will be terminated since it has last been heard from the NetWorker server. This timeout applies to all processes throughout the entire runtime operation.

NOTICE

The new Job inactivity timeout attribute applies to all processes throughout the entire runtime operation regardless of the job type and state. The savegrp inactivity timeout differs in that it only applies to save processes during a media session.

To set the Job inactivity timeout attribute:

1. In the **Administration** window, click **Configuration**.
2. In the left pane of the **Configuration** window, select the NetWorker server.
3. From the **File** menu, select **Properties**.
4. Select the **Configuration** tab.
5. For the **Job inactivity timeout** attribute, specify the maximum time in minutes that you want a job declared as inactive and terminated from the time it was last heard from.
6. Click **Ok**.

Report home

The report home feature is enabled by default during the installation of the NetWorker server and requires email capability on the NetWorker server. This connection enables the delivery of NetWorker configuration information to EMC Support when a default notification is triggered by an event in the NetWorker software. No other information or client data is sent to EMC Support.

By default, when a NetWorker event is triggered, an email is sent that includes the NSR RAP attribute data.

Enabling the report home feature

To enable the report home feature, use the **nsradmin** interface to specify the mail server:

1. From the NetWorker server, type the following to start the **nsradmin** interface:

```
nsradmin
```

2. At the **nsradmin** prompt, type:

```
print type: nsr report home
```

3. Edit the mail program attribute of the report home resource and type the name of the default mail server. For example, type:

```
update mail program: smtpmail -h mailserver
```

4. Type the following command to review the resource:

```
print
```

5. Quit **nsradmin**.

On windows, the default path for the report home output file is located in:

```
<NetWorker_install_path>\nsr\applogs\rh
```

Manually running a report home report

To configure a report home report to start immediately and run outside of the preselected scheduled time:

1. From the NetWorker server, type the following to start the **nsradmin** interface:

```
nsradmin
```

2. At the nsradmin prompt, type:


```
print type: NSR task; name: DefaultReportHomeTask
```
3. Edit the **autostart** attribute of the **DefaultReportHomeTask** resource and set it to **start now**. For example, type:


```
update autostart: start now
```
4. Type the following command to review the resource update:


```
print
```
5. Quit **nsradmin**.

Disabling the report home feature

To disable the report home feature:

1. From the NetWorker server, type the following to start the **nsradmin** interface:


```
nsradmin
```
2. At the nsradmin prompt, type:


```
print type: NSR task; name: DefaultReportHomeTask
```
3. Edit the **autostart** attribute of the **DefaultReportHomeTask** resource and set it to **Disabled**:


```
update autostart: Disabled
```
4. Type the following command to review the resource update:


```
print
```
5. Quit **nsradmin**.

Specifying additional email recipients

You can edit the **additional email address** attribute to include any internal company email address. You can use this feature to test that emails are correctly being sent.

To specify additional email recipients for the report home feature:

1. From the NetWorker server, type the following to start the **nsradmin** interface:


```
nsradmin
```
2. At the nsradmin prompt, type:


```
print type: nsr report home
```
3. Edit the **additional email recipients** attribute and type the additional email recipients. For example, type:


```
update additional email recipients: my_email@address.com
```


4. Type the following command to review the resource update:

```
print
```

5. Quit **nsradmin**.

Specifying the sender email address

You can edit the **sender email address** attribute to include any internal company email address. You can use this feature to test that emails are correctly being sent.

To specify additional sender email recipients for the report home feature:

1. From the NetWorker server, type the following to start the **nsradmin** interface:

```
nsradmin
```

2. At the **nsradmin** prompt, type:

```
print type: nsr report home
```

3. Edit the **sender email address** attribute and type the additional email recipients. For example, type:

```
update sender email address: my_email@address.com
```

4. Type the following command to review the resource update:

```
print
```

5. Quit **nsradmin**.

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 being processed by the system, thereby improving 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

Several attributes in various NetWorker resources are used to adjust the volume of data being processed by the system to improve overall performance.

The following attributes are related to parallelism:

- ◆ Client parallelism
- ◆ Server parallelism
- ◆ Savegrp parallelism
- ◆ Max active devices
- ◆ Media library parallelism

These attributes are described in detail in the following sections.

Client parallelism

Client parallelism is defined by using the Parallelism attribute of the Client resource, found in the NetWorker Console on the Globals (1 of 2) tab of the Client property dialog box.

Client parallelism defines the number of save streams that a client can send simultaneously during backup. Client parallelism has two default values, depending on whether the Client resource is a NetWorker server:

- ◆ For the Client resource for a NetWorker server, the default value for the client parallelism attribute is 12. This higher default value allows the server to complete a larger number of index backups even while the server's filesystem or other index backups are still running.
- ◆ Refer to [“Creating and configuring the NDMP client resource” on page 663](#) for recommended parallelism settings for NDMP clients.
- ◆ For all other clients (that is, clients that are not also NetWorker servers), the default value is 4. Many features within NetWorker, such as deduplication and VSS, will not work correctly with client parallelism set to a value that is higher than 4.

To avoid disk contention for all clients other than the Client resource of the NetWorker server, do not set the client parallelism higher than the number of physical disks involved in the client's backup, along with, on Microsoft Windows systems, the System State and System DB. Therefore, a typical Windows system configured with the ALL keyword for the client save set attribute will back up the C: and D: drives as well as the System State and System DB; in this case the default parallelism setting of 4 will be adequate. On the other hand, if you define multiple save sets on the same disk (C:\users, C:\system, C:\docs and so on), a higher client parallelism will result in multiple save streams attempting to access the disk at the same time.

Controlling client parallelism on 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 of the virtual clients that reside on a physical host are limited to the value specified for the physical host. For example, suppose you had ten virtual machines running on the same physical host. Each virtual machine is a NetWorker client and each client has a client parallelism of 4. This 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 4 save streams occurring for the backup of the ten virtual clients.

To base virtual client parallelism values on the physical host:

1. From the **Administration** window, click **Configuration**.
2. Click **Clients**.
3. Rick-click on a virtual NetWorker client resource and select **Properties**.
4. In the General tab, select the **Virtual Client** attribute and enter the name of the underlying physical host in the **Physical host** attribute.
5. In the Globals (1 of 2) tab, select the **Physical client parallelism** attribute.

Complete these steps for all virtual NetWorker clients that share the same physical host. Ensure that the value in the Physical host attribute is exactly the same for all virtual NetWorker client resources that share the same physical host.

Server parallelism

Server parallelism is defined by using the Parallelism attribute of the Server resource, found in the NetWorker Console on the Setup tab of the Server property dialog box.

Server parallelism defines the number of simultaneous save streams supported by the server. The default and the maximum allowed server parallelism varies depending on the edition of NetWorker software. Each enabled storage node connected to the NetWorker server can increase the parallelism maximum. The maximum parallelism for any NetWorker server and storage node combination can vary. The *EMC NetWorker 8.0 Release Notes* provides more information.

Optimally, your NetWorker server should be configured to process enough save streams to keep all backup devices in your datazone writing at their maximum speed. By tuning the server parallelism setting, along with other settings discussed in this section, you can maximize the speed that data is written to backup devices.

Savegroup parallelism

Savegroup parallelism is defined by using the Savegrp parallelism attribute of the Group resource, found in the NetWorker Console on the Advanced tab of the group property dialog box.

Savegroup parallelism determines the number of simultaneous save streams that will be allowed for a NetWorker group. The default value for this attribute is 0, which means that the attribute will have no effect on other parallelism settings. If the value is greater than zero, it will override any other parallelism considerations that savegrp would use.

Savegroup parallelism can help tune backups that involve multiple groups that overlap in their backup schedules, particularly if one group has a large number of clients. For example, if you have one group with 100 clients and another group with 4 clients, and you have the Savegrp parallelism attribute set to 0, savegrp may start as many clients in the first group as possible, up to the maximum number allowed by the server's server parallelism setting. This will mean that the smaller group will not be able to start any client backups, because the server parallelism is at its maximum, and the smaller group may experience timeouts and backup failures. To avoid this situation, you can set the savegroup parallelism for the larger group to a number below the server's server parallelism setting, guaranteeing that the clients in the smaller group will be able to begin their backups.

Max active devices

In a DDS environment, the maximum number of active devices for a storage node is defined by using the Max active devices attribute of the Storage Node resource. This attribute is found in the NetWorker Console in the Devices window, on the General tab of the Storage Node property dialog box.

This attribute sets the maximum number of devices that the NetWorker software may use from the storage node in DDS environments. In large environments with media libraries with large numbers of devices, it is possible that the storage node will not have the ability to optimize all drives in the library. The Max active devices attribute allows you to limit the

number of devices that the storage node will use at a given time, thereby allowing the storage node to have access to all devices in the library but limiting it to the number of devices it can fully optimize.

Media Library parallelism

Media library parallelism is defined by using the Max parallelism attribute of the Library resource, found in the NetWorker Console in the Devices window, on the Configuration tab of the Library property dialog box.

Media library parallelism allows you to designate the maximum number of available devices for the following autochanger operations:

- ◆ Inventorying volumes
- ◆ Labeling volumes

The Max parallelism attribute of the Library resource is normally set to one less than the number of devices within the library.

Library operations that operate on multiple volumes are more efficient if multiple devices can be used in parallel for these operations. However, because libraries may be performing multiple operations simultaneously, you may wish to restrict the number of devices available for inventorying and labeling operations, to leave some devices available for other library operations.

Multiplexing

Multiplexing is the ability to write multiple save 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 save streams to a particular device will improve performance of your NetWorker environment.

The NetWorker software has several attributes that can be used to increase or limit the number of save streams being written to a device.

The following attributes are related to multiplexing:

- ◆ [“Target sessions” on page 556](#)
- ◆ [“Max sessions” on page 557](#)
- ◆ [“Pool parallelism” on page 557](#)

These attributes are described in detail in the following sections.

Target sessions

The optimal number of target sessions per device is defined by using the Target sessions attribute of the Device resource, found in the NetWorker Console in the Devices window, on the Configuration tab of the Device property dialog box.

The Target sessions attribute allows you to set the optimal number of backup sessions accepted by an active device. This is not a hard limit; to set a hard limit for the number of sessions to a particular device, 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.

Consider the following when save sessions are started:

- ◆ 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 savestreams assignments with the same storage node, the device with the lowest target session value is used as a reference.

Max sessions

The maximum number of target sessions per device is defined by using the Max sessions attribute of the Device resource, found in the NetWorker Console in the Devices window, on the Configuration tab of the Device property dialog. This attribute defines the maximum number of save sessions for a device. Its value is never less than the target sessions value.

Pool parallelism

Pool parallelism is defined by using the Max parallelism attribute of the Pool resource, found in the NetWorker Console in the Media window, on the Configuration tab of the Pool property dialog box.

Pool parallelism determines the maximum number of simultaneous save streams for each device belonging to a NetWorker pool. The default value for this attribute is 0, which means that the attribute will have no effect on other parallelism settings.

Pool parallelism can be used to increase recovery times. For example, you can create a pool for backups of business critical data and use this attribute to restrict the number of save sets written in parallel to the media for this pool, which will increase the speed that data is recovered from that media. However, be aware that if the Max parallelism attribute of the pool is set to 1, there may be a prolonged delay between the backup of save sets. If this behavior occurs, try increasing the Max parallelism attribute for the pool resource.

Note: For AFTD and DD Boost devices, the pool's Max parallelism attribute is affected by the device's Max nsrmmmd count setting. 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 that AFTD_1 is selected by a back up pool whose Pool parallelism attribute is 1. Because each nsrmmmd process can initiate its own save sessions, the total number of save sessions that can be initiated 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 a tape device was used in the previous example, then the total number of save sessions would be 1.

Managing server access

This section describes how to restrict access to the NetWorker server and NetWorker operations.

- ◆ The privileges assigned to NMC, LDAP, and AD users and groups on a NetWorker server are configured in one of the following ways:
 - “[The administrator list](#)” on page 558
 - “[NetWorker User Groups](#)” on page 559
- ◆ “[Restrict backup and recover access to the NetWorker server](#)” on page 568 describes how to restrict server and client initiated backup and recover operations.

The administrator list

The NetWorker server software includes default administrator settings that provide the root user on a Unix NetWorker server and members of the Windows Administrators group on a Windows NetWorker server, all the NetWorker privileges required to change a NetWorker server’s configuration. These users are added to the **Administrator** list attribute in the NetWorker server resource, NSR

There are two way to update the administrators list:

- ◆ “[Using NMC](#)” on page 558
- ◆ “[Using nsraddadmin](#)” on page 559— If the NMC server and the NetWorker server are installed on separate hosts, the username of the Console daemon process owner must be added to the NetWorker Administrators list by using the **nsraddadmin** command.

Using NMC

To access the Administrators list attribute for a NetWorker server from NMC:

1. Connect to the NMC server with an account that has **Console Security Administrator** access.
2. Connect to the NetWorker server.
3. In the **Configuration** window, right-mouse click the NetWorker server in the left pane and select **Properties**. The **Administrator** attribute is under the **Setup** tab.

Consider the following:

- ◆ The **Administrator** attribute includes the following entries, by default:
 - On a Windows NetWorker server:


```
user=administrator, host=NetWorker_server_name
user=system, host=NetWorker_server_name
```
 - On a UNIX NetWorker server:


```
user=root, host=NetWorker_server_name
```

Using nsraddadmin

If the NMC server and the NetWorker server are installed on separate hosts, the Console administrator account must be granted access to administer and monitor all target NetWorker server.

To grant the Console administrator account access to the NetWorker server:

1. Connect to the NetWorker server with an administrator account on a Windows NetWorker server or the root account on UNIX NetWorker server.
2. From a command prompt, type:

```
nsraddadmin -u administrator@NMC_server
```

where *NMC_server* is the hostname of the NMC server.

NOTICE

The Console administrator account must be a member of the administrators list regardless of whether Native NMC authentication, LDAP authentication, or AD authentication is used.

NetWorker User Groups

To assign NMC, LDAP, and AD users and groups specific privileges on a NetWorker server, create and modify User Group resources.

By default, NMC, LDAP and AD users that have the NMC Console Security Administrator role are automatically added to a preconfigured Security Administrators user group on each NetWorker server that they have the right to manage. Members of the Security Administrators user group only have privileges to modify the Audit Log server and User Groups resources.

The following sections describe the NetWorker operations each privileges grants to a user and how to configure and manage the NetWorker User Groups:

- ◆ [“User Privileges” on page 560](#)
- ◆ [“User Groups” on page 563](#)

User Privileges

User privileges define the NetWorker operations and tasks that NMC, AD, and LDAP users are allowed to perform. Privileges are associated with user groups. With the exception of the Application Administrators user group and the Security Administrators user group, the privileges associated with a User Group can be modified. [Table 95 on page 560](#) provides a summary of the available privileges and the operations that are available to a user with each privilege.

Table 95 User Group privileges (1 of 3)

NetWorker Privilege	Operations Allowed
Change security settings	<p>The ability to modify:</p> <ul style="list-style-type: none"> • User groups • The Audit Log resource • The server resource <p>The Change security settings privilege requires that the following prerequisite privileges are also set: View Security Settings, Create Security Settings and Delete Security settings.</p>
View Security Settings	<p>The ability to view:</p> <ul style="list-style-type: none"> • User groups • The Audit Log resource • The server resource.
Create Security Settings	<p>The ability to create new user group resources.</p> <p>The Create Security Setting privilege requires that the following prerequisite privileges are also set: View Security Settings, Change Security Settings, and Delete Security settings.</p>
Delete Security Settings	<p>The ability to delete user created user groups. Preconfigured user groups cannot be deleted.</p> <p>The Delete Security Settings privilege requires that the following prerequisite privileges are also set: View Security Settings, Change Security Settings, and Delete Security settings.</p>
Remote access all clients	<p>The ability to:</p> <ul style="list-style-type: none"> • Remotely browse and recover data associated with any client • View configurations for all Client resources. This privilege is required to perform Directed Recovers. <p>This privilege supersedes the users defined in the Remote Access attribute of a Client resource.</p> <p>The Remote Access all clients privilege requires that the following prerequisite privileges are also set: Operate NetWorker, Monitor NetWorker, Operate Devices and Jukeboxes, Backup Local Data, and Recover Local Data.</p>
Configure NetWorker	<p>The ability to configure resources associated with the NetWorker server, storage nodes, and clients. This includes creating, editing, and deleting resources.</p> <p>Users with this privilege cannot configure User Group resources.</p> <p>The Configure NetWorker privilege requires that the following prerequisite privileges are also set: Operate NetWorker, Monitor NetWorker, Operate Devices and Jukeboxes, Backup Local Data, and Recover Local Data.</p>

Table 95 User Group privileges (2 of 3)

NetWorker Privilege	Operations Allowed
Operate NetWorker	<p>The ability to perform NetWorker operations. For example, members can:</p> <ul style="list-style-type: none"> • Reclaim space in a client file index. • Set a volume location or mode. • Start or stop a savegroup. • Disable a savegroup • Change the start time or the restart interval of the savegroup • Query the media database and client file indexes. <p>The Operate NetWorker privilege requires that the following prerequisite privileges are also set: Monitor NetWorker, Operate Devices and Jukeboxes, Backup Local Data, and Recover Local Data.</p>
Monitor NetWorker	<p>The ability to:</p> <ul style="list-style-type: none"> • Monitor NetWorker operations, including device status, save group status, and messages. • View media databases information. • View NetWorker configuration information (except the security settings described in the Change Security Settings privilege). <p>This privilege is not required to back up and recover local data, although it may be helpful for users to monitor messages and other information.</p>
Operate devices and jukeboxes	<p>The ability to perform device and autochanger operations, for example, mounting, unmounting, and labeling. Users with this privilege can also view device status and pending messages, as well as view information in the media database.</p> <p>The Operate devices and jukebox privilege requires that the Monitor NetWorker privilege is also set.</p>
Recover local data	<p>The ability to recover data from the NetWorker server to their local client, as well as view most attributes in the client's configuration. Members can also query the client's save sets and browse its client file index.</p> <p>This privilege does not provide permission to view information about other clients and does not override file-based privileges.</p> <p>Users can only recover files with the appropriate user privileges for that operating system. User with the privilege still must be logged in as root (UNIX) or administrator (Microsoft Windows) to perform save set or NDMP recovers.</p>
Backup local data	<p>The ability to:</p> <ul style="list-style-type: none"> • Manually back up data from their local client to the NetWorker server • View most attributes in the client's configuration. • Query the client save sets and browse the client file index. <p>This privilege does not provide permission to view information about other clients. Note that this privilege does not override file-based privileges.</p> <p>Users can only back up files with the appropriate user privileges for that operating system. Users with the privilege still must be logged in as root (UNIX) or administrator (Microsoft Windows) to run the savegrp command or perform NDMP backups. To allow scheduled backups to operate correctly, the root user (UNIX) or administrator (Microsoft Windows) on the client has this privilege automatically.</p>

Table 95 User Group privileges (3 of 3)

NetWorker Privilege	Operations Allowed
View Application Settings	<p>The ability to view NetWorker resources including: Archive Requests, Client resources, Device resources, Directives, Group, Jukebox, Label, License, Notification, Policies, Pool, Schedule, Staging, and Storage Node. It also allows user group members to view the status of operations.</p> <p>This privilege does not allow user group members to view the Server, user or user groups, or Security Audit Log resources.</p> <p>The View Application Settings privilege requires that the following prerequisite privileges are also set: Change Application Settings, Create Application Settings, and Delete Application settings.</p>
Change Application Settings	<p>The ability to change NetWorker resources including: Archive Requests, Client resources, Device resources, Directives, Group, Jukebox, Label, License, Notification, Policies, Pool, Schedule, Staging, and Storage Node. It also allows user group members to view the status of operations.</p> <p>This privilege does not allow user group members to change the Server, user groups, or Security Audit Log resources.</p> <p>The Change Application Settings privilege requires that the following prerequisite privileges are also set: Change Application Settings, Create Application Settings, and Delete Application settings.</p>
Create Application Settings	<p>The ability to create NetWorker resources including: Archive Requests, Client resources, Device resources, Directives, Group, Jukebox, Label, License, Notification, Policies, Pool, Schedule, Staging, and Storage Node. It also allows user group members to view the status of operations.</p> <p>This privilege does not allow user group members to change the Server, user groups, or Security Audit Log resources.</p> <p>The Create Application Settings privilege requires that the following prerequisite privileges are also set: Change Application Settings, Create Application Settings, and Delete Application settings.</p>
Delete Application Settings	<p>The ability to delete NetWorker resources including: Archive Requests, Client resources, Device resources, Directives, Group, Jukebox, Label, License, Notification, Policies, Pool, Schedule, Staging, and Storage Node. It also allows user group members to view the status of operations.</p> <p>This privilege does not allow user group members to delete the Server, or User Groups resources.</p> <p>The Delete Application Settings privilege requires that the following prerequisite privileges are also set: Change Application Settings, Create Application Settings, and Delete Application settings.</p>
Archive Data	<p>The ability to archive data. The NetWorker application administrator must have configured NetWorker for a user with this privilege to execute this operation. Only the Client resource that pertains to the client that issues the archive command is viewable.</p>
Backup Remote Data	<p>Allows users to remotely backup data.</p>
Recover Remote Data	<p>Allows users to recover remotely backed up data.</p>

User Groups

User Groups provide the ability to assign a group of NMC, LDAP, and AD users with a defined set of privileges to perform NetWorker operations.

By default, the NetWorker 8.0 software defines preconfigured user groups with specific privileges. These preconfigured user groups cannot be deleted. [Table 96 on page 563](#) provides a summary of the preconfigured users groups and the default privileges associated with each user group.

Prior to the NetWorker 8.0 software, a single Administrators user group was created. Any modifications to the users in the Administrator user group were automatically reflected in the Administrator attribute of the Server resource. In NetWorker 8.0 and later, the Administrator user group is replaced by three new Administrator user groups and user group membership changes are not reflected in the Administrator attribute of the Server resource.

The following sections describe how to manage NetWorker User Groups:

- ◆ [“Modifying LDAP or AD User Group membership” on page 564](#)
- ◆ [“Modifying native NMC users User Group membership” on page 565](#)
- ◆ [“Microsoft Windows groups NetWorker privileges” on page 566](#)
- ◆ [“Modifying User Group privileges” on page 567](#)
- ◆ [“Creating a User Group resource” on page 567](#)
- ◆ [“Copying a User Groups resource” on page 568](#)
- ◆ [“Deleting a User Groups resource” on page 568](#)

Table 96 User Groups and associated privileges (1 of 2)

NetWorker server-side User Groups	Associated privileges	
Security Administrators	View security settings Change security settings Create security settings Delete security settings	
Application Administrators	Remote Access All Clients Configure NetWorker Operate NetWorker Monitor NetWorker Operate Devices and Jukeboxes Recover Local Data Recover Remote Data	Backup Local Data Backup Remote Data Create Application Settings View Application Settings Change Application Settings Delete Application Settings Archive Data
Monitors	Monitor NetWorker Operate Devices and Jukeboxes Recover Local Data Recover Remote Data Backup Local Data	Backup Remote Data View Application Settings View Security Settings Archive Data

Table 96 User Groups and associated privileges (2 of 2)

NetWorker server-side User Groups	Associated privileges
Operators	Remote Access All Clients View Application Settings Operate NetWorker Monitor NetWorker Operate Devices and Jukeboxes
Auditors	View security settings
Users	Monitor NetWorker Recover Local Data Backup Local Data
Database Operators	Remote Access All Clients Operate NetWorker Monitor NetWorker Operate Devices and Jukeboxes
Database Administrators	Remote Access All Clients Configure NetWorker Operate NetWorker Monitor NetWorker Operate Devices and Jukeboxes

Modifying LDAP or AD User Group membership

To add, remove, or edit the LDAP or AD users and groups assigned to a User Group resource:

1. From the **Administration** window, click **Configuration**.
2. Click **User Groups**.
3. Right-click the user group to edit, then select **Properties**. The **Properties** dialog box appears.
4. In the **External Roles** field, specify the LDAP or AD object in one of the following formats:
 - *username*
 - *groupname*
 - *Group@LDAP_or_AD_hostname*
 - *user@LDAP_or_AD_hostname*
 - *host=LDAP_or_AD_hostname*
 - *role=role,host=LDAP_or_AD_hostname*
5. Click **Ok**. If the format of the object is invalid or the object is not found in the LDAP or AD authority, an error is displayed:

Cannot find group or user object in any configured authority.

Modifying native NMC users User Group membership

To add, remove, or edit the native NMC authentication users assigned to a User Group resource:

1. From the **Administration** window, click **Configuration**.
2. Click **User Groups**.
3. Right-click the user group to edit, then select **Properties**. The **Properties** dialog box appears.
4. In the **Users** field, specify the NMC user. Specify the username with the following syntax:

```
name=value[, name=value, ...]
```

where *name* can be one of the following:

- user
- group
- host
- domain
- domain_sid
- domaintype (either NIS or WINDOMAIN)

For example, to specify a user named *jdoe* on a host named *jupiter*, enter this line in the **Users** attribute:

```
user=jdoe,host=jupiter
```

NOTICE

The formats *user@host*, *host* and *user*, and similar formats are supported, but are ambiguous as to whether host or domain is intended. As a result, the *name=value* format is recommended.

If the value has spaces, it should be surrounded by quotation marks, for example:

```
domain="Domain Admins"
```

You can also enter just a username, which allows that user to administer NetWorker from any host (this is the same as an entry of *user=usemame*). Wildcards can also be used in place of a value. However, wildcards should be used with caution because they can compromise your enterprise security. Netgroup names can also be entered but must be preceded by an ampersand (&).

Example 54 Using the NetWorker Application Administrators group

This example shows what to enter to provide NetWorker administrative privileges to the following:

- The user *root* from any host
- The user *operator* from the hosts *mars* and *jupiter*
- Any users, valid hosts for the users, and valid domains for the users and host that are included in the netgroup *netadmins*:

```
user=root
user=operator,host=jupiter
user=operator,host=mars
&netadmins
```

Microsoft Windows groups NetWorker privileges

The NetWorker server recognizes domain names and Microsoft Windows groups, both local and global. For example:

- ◆ Administrators group
- ◆ Domain Admins group

If you are logged into a domain, only the *global* group is recognized. You can find out the name of your group by running the Windows utility **findgrp.exe**, which is available with the Windows Resource Kit.

If you are logged into an standalone Windows computer, only the *local* group is recognized, because there is no global group.

In cases where a user belongs to a domain that cannot be contacted by the server and therefore the username cannot be verified, you can use a more specific user description to guarantee that the appropriate user will have the appropriate privileges on the NetWorker server.

The syntax for this user description is as follows:

- ◆ For a single user:

```
user=user_name, domain=domain_name
```

For example:

```
user=joe,domain=NT-ENG
```

- ◆ For a group:

`group=group_name, domainsid=domain_id`

For example:

`group=Administrators,domainsid=S-1-5-32-323121-123`

Modifying User Group privileges

With the exception of the Application Administrators and Security Administrators user groups, the privileges associated with a user group can be changed.

To change the privileges assigned to a User Group resource:

1. From the **Administration** window, click **Configuration**.
2. Click **User Groups**.
3. Right-click the user group to edit, then select **Properties**. The **Properties** dialog box appears.
4. In the **Privileges** field, select or unselect the privileges as required.
5. Click **Ok**.

If a **Privilege** is selected but dependent privileges are not selected, an error is reported.

Creating a User Group resource

To create a User group resource:

1. From the **Administration** window, click **Configuration**.
2. Right-click **User Groups**, then select **New**. The **Create User Group** dialog box appears.
3. In the **Name** attribute, enter the name of the user group. The optional **Comment** attribute can be used to enter a description of the user group.
4. In the **Users** attribute, specify the users or groups to add to the user group. [“Modifying LDAP or AD User Group membership” on page 564](#) describes how to add LDAP and AD users. [“Modifying native NMC users User Group membership” on page 565](#) describes how to add native NMC authenticated users.
5. In the **Privileges** attribute, select the privileges to assign to the user group. [“User Privileges” on page 560](#) describes the NetWorker operations allowed with each privilege.
6. Click **OK**.

Copying a User Groups resource

To copy a User Group resource:

1. From the **Administration** window, click **Configuration**.
2. Click **User Groups**.
3. Right-click the user group to copy, then select **Copy**. The **Create User Group** dialog box appears, and contains the same information as the user group that was copied, *except* for **Name** attribute.
4. In the **Name** attribute, enter a name for the new user group.
5. Edit any other attributes as appropriate, then click **OK**.

Deleting a User Groups resource

To delete a User Group resource:

1. From the **Administration** window, click **Configuration**.
2. Click **User Groups**.
3. Right-click the user group to delete, then select **Delete**.
4. When prompted, click **Yes** to confirm the deletion.

NOTICE

You cannot delete a preconfigured user group.

Restrict backup and recover access to the NetWorker server

There are three ways to configure backup and restore session access to a NetWorker server:

- ◆ [“Allow manual saves” on page 568](#)
- ◆ [“Accept new save sessions” on page 569](#)
- ◆ [“Accept new recover sessions” on page 569](#)

Allow manual saves

This option determines if client initiated backups can be performed with the **save** command or the NetWorker User application (Windows clients only). This option is enabled by default.

To prevent manual saves from NetWorker clients:

1. From the **Administration** window, click **Configuration**.
2. In the left pane, right-click the NetWorker server and select **Properties**.
3. In the **Setup** tab, deselect **Manual saves**.

Accept new save sessions

In NetWorker 8.0 and later, the NetWorker server can be configured to not accept new save sessions from manual or schedule backups. For example, save sessions can be rejected to allow routine NetWorker Server maintenance such as server reboots, to occur without cancelling in progress backup or clone operations during the shutdown process. By default the NetWorker server is configured to accept new save sessions.

To prevent the NetWorker server from accepting new save sessions:

1. From the **Administration** window, click **Configuration**.
2. In the left pane, right-click the NetWorker server and click **Properties**.
3. In the **Miscellaneous** tab, unselect **Accept new sessions**.

Accept new recover sessions

In NetWorker 8.0 and later, the NetWorker server can be configured to not accept new recover and cloning sessions. For example, recover sessions can be rejected to allow routine NetWorker Server maintenance such as server reboots, to occur without cancelled in progress recover operations during the shutdown process. By default the NetWorker server is configured to accept new recover sessions.

To prevent the NetWorker server from accepting new recover sessions:

1. Connect to the NetWorker server in the NetWorker Management Console application.
2. In the **Administration** window, in the NetWorker Management Console, connect to the NetWorker server.
3. In the **Configuration** view, right-click the NetWorker server and click **Properties**.
4. In the **Miscellaneous** tab, deselect **Accept new recover sessions**.

Working with the Multi-Tenancy Facility

Use the Multi-Tenancy Facility option to create multiple restricted data zones. End users can access a single NetWorker server without being able to view data, backups, recoveries, or modify in other data zones. In addition, Tenant administrators within a restricted data zone can only see a very limited amount of the information (log and alerts) managed by the global administrator or other restricted data zones from the Console or the CLI.

The Multi-Tenancy Facility feature is enabled by configuring a Restricted Data Zone resource on the NetWorker server. A restricted data zone is associated with its own NetWorker resources. The *EMC NetWorker 8.0 Release Notes* provides information about NetWorker Module support within a restricted data zone.

This topic includes the following sections:

- ◆ [“Users within a restricted data zone” on page 570](#)
- ◆ [“Configurations for the Multi-Tenancy feature” on page 573](#)
- ◆ [“Configuring a restricted data zone” on page 574](#)
- ◆ [“Restricted Data Zone resource associations” on page 583](#)
- ◆ [“Viewing data within a Restricted Data Zone” on page 584](#)

Users within a restricted data zone

The following roles apply to the Multi-Tenancy Facility feature:

- ◆ global administrator
- ◆ tenant administrator

Global administrator

A global administrator oversees several restricted data zones to ensure proper setup and use of the restricted data zones, and to resolve issues that the tenant administrators cannot resolve by themselves.

A global administrator for a restricted data zone:

- ◆ Must be listed in the Administrator list in the NetWorker Server Properties window of the NetWorker server. [“Managing server access” on page 558](#) provides more information.
- ◆ Requires the following NetWorker software access rights:
 - The Console User role, if access to NMC is required.
 - Either of the following privileges on the NetWorker server:
 - Configure NetWorker, which includes the following privileges:
 - Operate NetWorker
 - Monitor NetWorker
 - Operate Devices and Jukeboxes
 - Backup Local Data
 - Recover Local Data
 - or
 - Create Application Settings, which includes the following privileges:
 - View Application Settings
 - Change Application Settings
 - Delete Application Settings

[“Managing server access” on page 558](#) provides information about how to configure and set access rights and privileges on the NetWorker server.

Tenant administrator

A tenant administrator is limited to managing backups and recoveries within a restricted data zone. Tenant administrators can create, view, operate, manage, and modify the NetWorker resources within their own restricted data zone; not the restricted data zones of other tenants.

NOTICE

A tenant administrator cannot create or edit a Restricted Data Zone resource. Although, they can view the resource from the Console or the CLI.

The procedures to configure a tenant administrator differ when using native authentication and LDAP authentication. The following topics provide detailed instructions about how to create a user account for a tenant administrator:

- ◆ [“How to create a tenant administrator when using native authentication” on page 571](#)
- ◆ [“How to create a tenant administrator when using LDAP authentication” on page 572](#)

Considerations

When configuring tenant administrators within the restricted data zone, consider the following:

- ◆ Non-default resources in the global data zone are not visible to tenant administrators who only have privileges within the restricted data zone.
- ◆ Tenant administrators cannot be a member of more than one unique restricted data zone. However, a tenant administrator can appear in multiple resource instances of the same restricted data zone. This allows for fine grain control over the tenant administrator and their privileges because all restricted data zones group together in a similar way that Client resources with the same name group together. [“Configuring multiple Restricted Data Zones resources that use the same name” on page 582](#) provides details.
- ◆ A restricted data zone can have multiple tenant administrators.
- ◆ The Privileges section of the Restricted Data Zone resource specifies the privileges for a tenant administrator.

For resources that have a restricted data zone associated with them, the Restricted Data Zone resource and the User Groups resource use the same Privileges attributes.

- ◆ Tenant administrators can only create devices and jukeboxes on storage nodes that are associated to them exclusively within the restricted data zone.

NOTICE

Do not use the wild card character * in the **Users** attribute of a Restricted Data Zone resource.

How to create a tenant administrator when using native authentication

The global administrator configures a restricted data zone and assigns tenant administrators. [“Configuring a restricted data zone” on page 574](#) describes how to configure a restricted data zone and add a tenant administrator to the restricted data zone.

To create a tenant administrator:

1. (Optional) Create the Console user create a Console user account if access to NMC is required.

To create a Console user account:

- a. Log in to the Console server as a Console Security Administrator.
- b. From the **Console** window, click **Setup**.
- c. In the left pane, right-click **Users** and select **New**. The **Create User** dialog box opens and displays the **General** tab.

- d. In the **User Name** field, type a username for the tenant administrator. The username must not:
 - Exceed 20 characters.
 - Use spaces or any of these characters:
 - :
 - < >
 - /
 - Use characters with an ASCII value less than or equal to 32.
 - Begin a username with an underscore (`_`) character.
- e. (Optional) In the **Full Name** field, type the full name of the user.
- f. (Optional) In the **Description** field, type the user description.
- g. In the **Role** field, select **Console User**.
- h. In the **Password** field, type the user password for the tenant administrator.

NOTICE

Passwords must be at least eight characters long and cannot be the same as the user name. The Console server enforces this requirement when you create or edit users or when you change a user password after you upgrade from a previous release.

- i. In the **Confirm Password** field, retype the password and click **OK**.

NOTICE

If records are enabled, all Console users can query records.

2. Add the Console user account that you created in [step 1 on page 571](#) or other user account to the **Users** field in the Restricted Data Zone resource. [“How to configure a restricted data zone” on page 574](#) provides more information.
3. Assign privileges to the tenant administrator within the restricted data zone. [“How to configure a restricted data zone” on page 574](#) provides more information.

How to create a tenant administrator when using LDAP authentication

The global administrator configures a restricted data zone and assigns tenant administrators. [“Configuring a restricted data zone” on page 574](#) describes how to configure a restricted data zone.

When an LDAP user logs in for the first time, a user object is automatically created on the Console server. You only need to map LDAP user roles or LDAP user names to Console user roles. If necessary, you can also create user objects before users log in for the first time. For example, you may want to restrict user access to managed servers before the user logs in for the first time.

[“Adding or removing LDAP or AD console users” on page 518](#) provides more information about how to configure NMC and the NetWorker software into LDAP mode.

To add a Console user account when using LDAP authentication:

1. (Optional) Create the Console user:
 - a. Log in to the Console server as a Console Security Administrator.
 - b. From the **Console** window, click **Setup**.
 - c. In the left pane, right-click **Users** and select **New**. The **Create User** dialog box opens and displays the **Identity** tab.
 - d. In the **User Name** attribute, type the LDAP user name.
 - e. (Optional) In the **Full Name** field, type the full name of the LDAP user.
 - f. (Optional) In the **Description** field, type the user description.
 - g. Click **OK**.
2. Map LDAP users to Console roles.
 - a. In the left pane, select **Roles**.
 - b. In the right pane, right-click **Console User** and select **Properties**.
 - c. In the **External Roles** attribute, add each LDAP user role or LDAP user name to be mapped. Type each entry on a separate line.
3. Click **OK**.
4. Assign privileges to the tenant administrator within the restricted data zone. [“How to configure a restricted data zone” on page 574](#) provides more information.

Configurations for the Multi-Tenancy feature

The following two configurations are supported for the Multi-Tenancy Facility feature:

- ◆ A global administrator assigns tenant administrators who can create, view, operate, manage, and modify the NetWorker resources associated with their own restricted data zone; not the restricted data zones of other tenants.

In this configuration, the global administrator sets up the initial privileges and restrictions such as the maximum numbers of clients, devices, jukeboxes, and storage nodes, and then lets the tenant administrators set up their own resources to meet their individual business requirements. Restrictions are set in place to ensure that tenant administrators do not overuse resources, while still enabling them the freedom to set up the restricted data zone. An error message displays if the set restrictions are exceeded.

- ◆ A global administrator can manage multiple Restricted Data Zone resources. In this configuration, the global administrator controls the layout of the restricted data zones, including the set up of clients, devices, jukeboxes, storage nodes, and schedules.

Configuring a restricted data zone

To configure a restricted data zone, you apply a set of privileges on a per resource basis. Note that you apply User Groups privileges on a per data zone level.

The association of a restricted data zone to resource is one-to-many, such that a restricted data zone can be associated with multiple resources. However, a single resource can not be associated to different restricted data zones. When a resource is associated to a restricted data zone, it is immediately made available for use by the restricted data zone.

Note that NetWorker clients and groups that are associated to a restricted data zone must both belong exclusively to that particular restricted data zone, and must not participate in other restricted data zones. The Multi-Tenancy Facility feature does not support a scenario where one group is used to back up clients that belong to different restricted data zones.

Error messages appear if there is an attempt to:

- ◆ Associate more of a particular type of resource than is specified in the **Restriction** attribute of the Restricted Data Zone resource.
- ◆ Configure a resource which cannot be associated to a Restricted Data Zone.

How to configure a restricted data zone

Tenant administrators cannot create or edit a Restricted Data Zone resource. However, they can view the resource from the Console or CLI.

To configure a restricted data zone, you must create or edit a Restricted Data Zone resource:

1. Ensure that you have the at least the minimum privileges to create or edit a Restricted Data Zone resource. [“Global administrator” on page 570](#) provides more information about required privileges.
2. To configure a restricted data zone where tenant administrators are assigned:
 - a. Ensure that the tenant administrators are not listed in Users attribute of any User Groups resources. Privileges set in a User Groups resource apply to all resources, not just the resources within a Restricted Data Zone.

NOTICE

Privileges that are set for the tenant administrator apply to all resources within the restricted data zone.

- b. Remove the *@* character from the Users attribute of each of the NetWorker User Groups resources. By default, *@* appears in the Users field of the User Groups resource named Users.

NOTICE

If the *@* characters are left in the access attribute of any NetWorker User Group resource or in the Remote Access attribute of any NetWorker resource such as a Client resource, then that resource will be visible from any restricted datazone.

3. Configure the following customized NetWorker resources for the restricted data zone, if required.

Default directives, labels, media pools, schedules, and policies are available to all restricted data zones at all times:

- Groups — A backup group specifies the time of day when a backup occurs. [Chapter 7, “Backup Groups and Schedules,”](#) provides information about groups.

NOTICE

NetWorker groups and clients that are associated to a restricted data zone must both belong exclusively to that particular restricted data zone, and must not participate in other restricted data zones. The Multi-Tenancy Facility feature does not support a scenario where one group is used to back up clients that belong to different restricted data zones.

- Directives — Optional instructions that control how files and directories are processed during backup and recovery. [“Preconfigured global directive resources” on page 294](#) provides more information about the preconfigured directive options.
 - Labels — 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. [“Label templates” on page 318](#) provides more information about the preconfigured label templates.
 - Media Pools — A collection of volumes to which backup data is written. Pools are used to sort backup volumes so that the volumes are easy to locate when they are required. [“Media pools” on page 304](#) provides more information about media pools.
 - Schedules — The NetWorker software ships with preconfigured schedules. If these schedules meet backup requirements, use them as is. Otherwise, create new schedules to accommodate any site-specific requirements. [“Schedules” on page 260](#) provides more information about schedules.
 - Policies — Specifies a time range that the Restricted Data Zone resource takes effect. You can use the Policies resource to create or modify the preconfigured policies. [“Preconfigured time policies” on page 283](#) provides more information about policies.
4. In the server **Administration** interface, click **Configuration**.
 5. Right-click **Restricted Data Zones** and select **New**. The **Create Restricted Data Zone** dialog box appears.
 6. Name the restricted data zone:
 - a. Click the **General** tab.
 - b. In the **Name** field, type the name of the of the restricted data zone.
 - c. (Optional) In the **Comment** field, add a descriptive comment for the restricted data zone.

7. Set restrictions to limit the number of clients, devices, storage nodes, and jukeboxes that the tenant administrator can create or associate within the restricted data zone:
 - a. Click the **General** tab.
 - b. In the **Number of clients** field, set restrictions to limit the number of clients that the tenant administrator can create or associate within the restricted data zone.
 - c. In the **Number of devices** field, set restrictions to limit the number of devices that the tenant administrator can create or use within the restricted data zone.
 - d. In the **Number of storage nodes** field, set restrictions to limit the number of storage nodes that the tenant administrator can create or associate to the restricted data zone.
 - e. In the **Number of jukeboxes** field, set restrictions to limit the number of jukeboxes that the tenant administrator can create or use within the restricted data zone.

NOTICE

You must use AlphaStor when two restricted data zones share a jukebox.

8. List the tenant administrators and assign their privileges within the restricted data zone:
 - a. Add the user accounts that are granted permission to create and manage the NetWorker resources within this particular restricted data zone.

The procedures differ when using native authentication or LDAP authentication:

- For LDAP authentication:
 - In the **External roles** field, add the user accounts that are granted permission to create and manage the NetWorker resources within this particular restricted data zone as a tenant administrator. [“How to create a tenant administrator when using LDAP authentication” on page 572](#) describes how to create a tenant administrator by using LDAP.
 - Leave the **Users** field blank.
- For native authentication, in the **Users** field, add the user accounts that are granted permission to create and manage the NetWorker resources within the particular restricted data zone as a tenant administrator. [“How to create a tenant administrator when using native authentication” on page 571](#) describes how to create a tenant administrator by using native authentication.

Use the following format:

sam@jupiter or **user=sam, host=jupiter** – User Sam on machine jupiter

host=jupiter – Any user on machine jupiter

NOTICE

Tenant administrators that are listed in the **Users** field can only create the NetWorker resources within their particular restricted data zone; not the restricted data zones of other third party users. [“Viewing data within a Restricted Data Zone” on page 584](#) provides more information.

- b. In the **Privileges** field, select the privileges that are required for the tenant administrators to create and manage resources within this particular restricted data zone. [“Customizing privileges” on page 579](#) provides more information.

NOTICE

Privileges you set for the tenant administrator apply to all resources within the restricted data zone.

- 9. In the **Group and Clients** tab:

NOTICE

NetWorker groups and clients that are associated to a restricted data zone must both belong exclusively to that particular restricted data zone, and must not participate in other restricted data zones. The Multi-Tenancy Facility feature does not support a scenario where one group is used to back up clients that belong to different restricted data zones.

- a. Select the groups that the specified tenant administrator can access within this particular restricted data zone. A backup group specifies the time of day when a backup occurs. [Chapter 7, “Backup Groups and Schedules,”](#) provides information about groups.

NOTICE

The group must only contain NetWorker clients that belong exclusively to that particular group.

If a group is associated or disassociated to a restricted data zone, then all of its related clients are associated or disassociated automatically when a change to the resource is committed. This setting ensures that resources are not partially associated to a restricted data zone.

- b. Select the clients that can be administered by the specified users within this particular restricted data zone.

NOTICE

A client that belongs to a group must be associated with a restricted data zone. If a client is not associated with a restricted data zone, then the client is considered to belong to the global data zone and is not considered to be part of the restricted data zone.

- 10. In the **Jukeboxes and Devices** tab:

- a. Select a jukebox that can be accessed by the specified tenant administrator within this particular restricted data zone.
- b. Select a device that can be accessed by the specified tenant administrator within this particular restricted data zone.

NOTICE

If a jukebox is associated or disassociated to a restricted data zone, then all of its related devices are associated or disassociated when a change to the resource is committed. This setting ensures that resources are not partially associated to a restricted data zone.

A global administrator can share storage nodes by assigning devices and jukeboxes on a storage node to multiple restricted data zones. However, the storage nodes are not marked as associated in the Restricted Data Zone resource. [“Associating storage nodes” on page 582](#) provides more information.

A tenant administrator can only create devices and jukeboxes on storage nodes that are associated to them exclusively in the restricted data zone by the global administrator.

11. If required, select the already configured customized NetWorker resources that you configured in [step 3 on page 575](#) for the restricted data zone:

NOTICE

Default directives, labels, media pools, schedules, and policies are available to all restricted data zones at all times

- In the **Directives** tab, select a directive that can be accessed and used by the specified tenant administrator within this particular restricted data zone. Directives are optional instructions that control how files and directories are processed during backup and recovery. [“Preconfigured global directive resources” on page 294](#) provides detailed information about the preconfigured directive options.
- In the **Labels** tab, select labels for the backup volumes. The tenant administrators within the restricted data zone can access these volumes. [“Label templates” on page 318](#) provides more information about the preconfigured label templates.
- In the **Media Pools** tab, select the media pools that can be accessed and used by the specified tenant administrators within this particular restricted data zone. Media pools are a collection of volumes to which backup data is written. Pools are used to sort backup volumes so that the volumes are easy to locate when they are required. [“Media pools” on page 304](#) provides detailed information.
- In the **Schedules** tab, select a schedule that can be accessed and used by the specified tenant administrators within this particular restricted data zone. [“Schedules” on page 260](#) provides detailed information.
- In the **Policies** tab, select a policy to specify a time range that the restricted data zone resource takes effect. You can use the Policies resource to create or modify the preconfigured policies. [“Preconfigured time policies” on page 283](#) provides more information.

12. In the **Storage Nodes** tab, select a storage node that the tenant administrator can access and use within the restricted data zone:
 - A global administrator can share storage nodes by assigning devices and jukeboxes to multiple restricted data zones. However, the storage nodes are not marked as associated in the Restricted Data Zone resource. [“Associating storage nodes” on page 582](#) provides more information.
 - A tenant administrator within a restricted data zone can only create devices and jukeboxes on storage nodes that are associated to their restricted data zone.
13. In the **Operation Status** tab, do not change the default setting.

The Operation Status is automatically set and cannot be manually associated to a restricted data zone.
14. Click **Ok** to complete the configuration. You can view the newly created restricted data zone in the **Restricted Data Zone** pane.

Customizing privileges

You can customize privileges associated with a tenant administrator within a restricted data zone to fit specific requirements. You can assign the tenant administrator privileges in the **Privileges** field of the **Restricted Data Zone** resource.

[Table 97 on page 579](#) lists preconfigured privileges and their associated privileges.

Table 97 Tenant administrator privileges (1 of 4)

Privilege	Privileges
Remote Access All Clients	<p>Allows a tenant administrator to:</p> <ul style="list-style-type: none"> • Remotely browse and recover data associated with any client within the restricted data zone. • View configurations for all Client resources that are within the restricted data zone. <p>This privilege is required to perform Directed Recovers. This privilege supersedes the Remote Access attribute in the Client resource.</p> <p>You must also enable the following five privileges:</p> <ul style="list-style-type: none"> • Operate NetWorker • Monitor NetWorker • Operate Devices and Jukeboxes • Backup Local Data • Recover Local Data
Operate NetWorker	<p>Allows a tenant administrator to perform NetWorker operations.</p> <p>For example, members can perform the following tasks with their restricted data zone:</p> <ul style="list-style-type: none"> • Reclaim space in a client file index. • Set a volume location or mode. • Start or stop a savegroup. • Query the media database and client file indexes. <p>You must also enable the following five privileges:</p> <ul style="list-style-type: none"> • Monitor NetWorker • Operate Devices and Jukeboxes • Backup Local Data • Recover Local Data

Table 97 Tenant administrator privileges (2 of 4)

Privilege	Privileges
Configure NetWorker	<p>Allows a tenant administrator to configure the following resources associated with the NetWorker server, its storage nodes, and clients. This includes creating, editing, and deleting resources within the restricted data zone:</p> <ul style="list-style-type: none"> • Clients • Devices • Directives • Groups • Jukeboxes • Labels • Media Pools • Policies • Schedules • Storage Nodes <p>You must also enable the following five privileges:</p> <ul style="list-style-type: none"> • Operate NetWorker • Monitor NetWorker • Operate Devices and Jukeboxes • Backup Local Data • Recover Local Data <p>Tenant administrators can only edit resources within their own particular restricted data zone. However, they can create new resources if they are not blocked by the quantity restrictions in the Restricted Data Zones resource.</p> <p>Tenant administrators with this privilege can not configure NetWorker resources that have not been assigned to them within the restricted data zone.</p> <p>Tenant administrators can not configure the following resources:</p> <ul style="list-style-type: none"> • User Group resource • Restricted Data Zone resource <p>Tenant administrators with this privilege can not enable the Remote Access attribute in the Client resource. Permission to change the Remote Access attribute in the Client resource is granted only through the Change Security Settings privilege.</p>
Monitor NetWorker	<p>Allows a tenant administrator to:</p> <ul style="list-style-type: none"> • Monitor NetWorker operations, including device status, save group status, and messages. • View media databases information. • View NetWorker configuration information within the restricted data zone. <p>This privilege is not required to back up and recover local data. However, this privilege might be helpful for a tenant administrator to monitor messages and other information.</p> <p>While most information is limited to what is within the restricted data zone. A tenant administrator with this privilege can view all of the current log messages and alerts that are present in the NetWorker software by using NMC or the nsrwatch command.</p>
Operate Devices and Jukeboxes	<p>Allows a tenant administrator to perform device and autochanger operations within the restricted data zone. For example, operations such as mounting, unmounting, and labeling.</p> <p>You must also enable the Monitor NetWorker privilege.</p> <p>A tenant administrator with this privilege can also view device status, and pending messages, and view information in the media database.</p>
Recover Local Data	<p>Allows a tenant administrator to recover data from the NetWorker server to the local client and to view most attributes in the client's configuration.</p> <hr/> <p>Note: This privilege does not provide permission to view information about other clients. This privilege does not override file-based privileges. Tenant administrators can only recover files with the appropriate user privileges for the operating system.</p> <p>A tenant administrator with this privilege still must log in as root (UNIX) or administrator (Microsoft Windows) to perform save set or NDMP recovers.</p> <hr/>

Table 97 Tenant administrator privileges (3 of 4)

Privilege	Privileges
Backup Local Data	<p>Allows a tenant administrator to manually back up data from local clients in the restricted data zone to the NetWorker server. Use the nsradmin command to query information about the current client backup.</p> <p>Note: This privilege does not override file-based privileges. A tenant administrator can only back up files with the appropriate user privileges for the operating system. A tenant administrator with this privilege still must log in as root (UNIX) or administrator (Microsoft Windows) to run the savegrp command or perform NDMP backups. To allow scheduled backups to operate correctly, the root user (UNIX) or administrator (Microsoft Windows) on the client has this privilege by default.</p>
Archive Data	This privilege has no effect in a restricted data zone.
Backup Remote Data	This privilege has no effect in a restricted data zone.
Recover Remote Data	This privilege has no effect in a restricted data zone.
Delete Application Settings	<p>Allows a tenant administrator to delete application settings that were set in the NetWorker software for the particular data zone.</p> <p>You must also enable the following three privileges:</p> <ul style="list-style-type: none"> • Create Application Settings • View Application Settings • Change Application Settings
Change Application Settings	<p>Allows a tenant administrator to change application settings that were set in the NetWorker software for the particular data zone.</p> <p>You must also enable the following three privileges:</p> <ul style="list-style-type: none"> • Create Application Settings • View Application Settings • Delete Application Settings
View Application Settings	Allows a tenant administrator to view application settings that were set in the NetWorker software for their particular data zone.
Create Application Settings	<p>Allows a tenant administrator to create application settings for their particular data zone.</p> <p>You must also enable the following three privileges:</p> <ul style="list-style-type: none"> • View Application Settings • Change Application Settings • Delete Application Settings
Change Security Settings	<p>Allows a tenant administrator to edit the Remote Access attribute in the Client resources that belongs to the particular restricted data zone.</p> <p>You must also enable the following three privileges:</p> <ul style="list-style-type: none"> • Delete Security Settings • Create Security Settings • View Security Settings
Delete Security Settings	<p>Prohibits a tenant administrator from deleting a security setting.</p> <p>Enable this privilege if you selected the Change Security Settings privilege. This privilege has no impact on the restricted data zone.</p> <p>You must also enable the following three privileges:</p> <ul style="list-style-type: none"> • View Security Settings • Change Security Settings • Create Security Settings

Table 97 Tenant administrator privileges (4 of 4)

Privilege	Privileges
Create Security Settings	Prohibits a tenant administrator from deleting a security setting. Enable this privilege if you selected the Change Security Settings privilege. This privilege has no impact on the restricted data zone. You must also enable the following three privileges: <ul style="list-style-type: none"> • View Security Settings • Change Security Settings • Create Security Settings
View Security Settings	Allows a tenant administrator to view the security settings that were set in the NetWorker software for the particular data zone.

Associating storage nodes

A global administrator can share storage nodes by assigning devices and jukeboxes to multiple restricted data zones. However, the storage nodes are not marked as associated in the Restricted Data Zone resource.

A tenant administrator within a restricted data zone can:

- ◆ Create a storage node if the global administrator has set the correct privileges and restrictions. [“Customizing privileges” on page 579](#) provides more information.
- ◆ Create devices and jukeboxes only on storage nodes that are associated within the restricted data zone.

You can associate storage nodes to a restricted data zone in two ways:

- ◆ Exclusive use
For exclusive use, associate the storage node within the restricted data zone. This best practice allows tenant administrators to create their own devices and modify and delete the storage node. When a tenant administrator with the proper privilege creates a storage node, the storage node and restricted data zone automatically associate during the create process.
- ◆ Group use
For group use, do not explicitly associate the storage node to any restricted data zone. The global administrator must create and associate devices and jukeboxes to the restricted data zone. This allows the restricted data zones and the global data zone to all share the same storage nodes.

Configuring multiple Restricted Data Zones resources that use the same name

You can configure multiple Restricted Data Zone resources with the same name. Using one name for multiple Restricted Data Zone resources allows fine grain control over a restricted data zone's administrators and their privileges. The use of one name to group multiple Restricted Data Zone resources is similar to the way you group Client resources with the same name.

To provide optimal flexibility, you can define multiple Restricted Data Zone resources for the same data zone, with the same name.

For example, a list of restricted data zones might include two instances of a restricted data zone named Company_One. Each instance is assigned a unique tenant administrator, each with different privileges. In this scenario, one tenant administrator is designated for

configuration and the other for operation. Each tenant administrator has a different privilege set. All information in a restricted data zone, except the tenant administrators of the data zone and their privileges, are propagated to all other Restricted Data Zone resources with the same name.

You can use the **Comment** attribute in the Restricted Data Zone resource to help distinguish between multiple Restricted Data Zone resources with the same name.

NOTICE

A tenant administrator cannot be included in more than one restricted data zone. However, a tenant administrator can be included in more than one Restricted Data Zone resource within the same restricted data zone.

Limiting the tenant administrator from creating too many resources

You can set limits on the following Restricted Data Zone resources:

- ◆ Number of clients
- ◆ Number devices
- ◆ Number of storage nodes
- ◆ Number of jukeboxes

By setting resource limits you prevent tenant administrators from overusing resources and licenses:

- ◆ If tenant administrators exceed the restriction limitations when creating or updating a resource, an error message appears notifying them that the update or creation will not be performed.
- ◆ If tenant administrators review their own restricted data zone, the list of selectable resource is limited to the resources that were specified for them in the Restricted Data Zone resource by the global administrator.

Restricted Data Zone resource associations

When tenant administrators create a new resource, providing that they have permission to do so, that new resource is automatically associated to their restricted data zone. When the resource is created, the restricted data zone's restriction counts are checked. An error message appears if the restriction count is exceeded.

When a global administrator or a tenant administrator deletes a resource from the restricted zone, the resource is automatically disassociated from its original restricted data zone.

For resources that exist in multiple instances such as clients, when you delete the last copy of the resource, the Restricted Data Zone resource disassociates. You can manually disassociate resources by deleting the original resource.

Special case associations

Associations between the following resources have a special built-in inheritance policy:

- ◆ Groups and clients
- ◆ Jukeboxes and devices

If a group or jukebox is associated or disassociated to a restricted data zone, then all of its related clients or devices are associated or disassociated when a change to the resource is committed. This function ensures that resources are not partially associated to a restricted data zone.

Viewing data within a Restricted Data Zone

To aid in the segregation of restricted data zones, some resources have limited visibility. Tenant administrators within a restricted data zone cannot view or modify the resources of other third-party restricted data zones. Resources and save set information do not appear to a tenant administrator, who is outside of a restricted data zone.

Tenant administrators can view some resources, but cannot associate them with a restricted data zone. This allows some of the basic functionality of the NetWorker software to be performed.

Only the default pre-sets of the following NetWorker resources are always visible to the tenant administrator within the restricted data zone:

- ◆ Directives
- ◆ Label Templates
- ◆ Media Pools
- ◆ Schedules
- ◆ Policies

For example, the Encryption directive is always visible to all restricted data zones.

Tenant administrators see a limited view of NetWorker resources. When tenant administrators review their own restricted data zone, the list of selectable resources is limited to the resources that the global administrator specified for them in the Restricted Data Zone resource. Tenant administrators can perform report queries to view basic information about clients and save sets. File data is not available.

Server communication issues within Microsoft Windows

This section addresses various client/server communication issues that occur when running NetWorker software in a Microsoft Windows environment.

Name resolution

If the network consists of only Microsoft Windows computers, you may find WINS or LMHOSTS is adequate for using NetWorker software. However, when using the software with clients running on other platforms, such as UNIX, you must use a local host file or DNS name resolution.

You must add the NetWorker server name to either the local hosts file (located in %SystemRoot%\system32\drivers\etc) or the Domain Name Server that contains the names of all servers on your network.

Backup Operators group

The Microsoft 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 Backup Operators group on the problematic computer to ensure that you are properly assigned.

Dynamic Host Configuration Protocol

A NetWorker server requires a static (fixed) hostname address. Typically, addresses for Dynamic Host Configuration Protocol (DHCP) clients change because they use dynamic addressing. If the address changes, the authorization code for that NetWorker server becomes invalid. If the NetWorker server is a DHCP client, a static TCP/IP address for the server must be preserved. NetWorker clients can still use dynamic addressing.

Backup and Recover Server service

In Microsoft Windows operating systems, the NetWorker Backup and Recover Server service is normally started by the Windows System account. This allows backup and recovery services to run even if no one is logged onto that computer.

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 591](#) 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 it 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 saved on that volume. The NetWorker server selects which volume to mount for recovering a file by mapping the saved files to their volumes.

Each entry in the client file index includes this information for a backed-up file:

- ◆ Filename
- ◆ Number of blocks
- ◆ Access privileges
- ◆ Number of links
- ◆ Owner
- ◆ Group
- ◆ Size
- ◆ Last modified time
- ◆ Backup time

The online indexes grow with each backup, as entries are added for each newly backed-up file and save set. As long as an index entry for a file remains in the client file index, you can recover 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 is unable to access the media database and is thus unable to access and recover data. Unless you configure the server to control the size of the online indexes by using browse and retention policies, they continue to grow until they exceed the capacity of the file system.

NetWorker uses browse and retention policies to manage the lifecycle of data, and to automatically control the size of the client file index. [“About browse and retention policies” on page 276](#) provides information on 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 592](#) provides more information.

Checking online indexes

Each time the NetWorker server starts, it 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 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 outlined in [“Recovering expired save sets” on page 398](#).

It is also good maintenance practice to periodically run the **nsrck -F** and **nsrim -X** commands to check the integrity of the client and media indexes. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information about these commands.

Viewing information about the indexes

To view information about the indexes:

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.

[Table 98 on page 587](#) identifies the index information displayed for each client.

Table 98 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.

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. It also includes an option to remove old save set cycles.

Viewing client save set information

To view the information about client save sets:

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.

[Table 99 on page 588](#) identifies the information in the **Save Sets** dialog box for each save set.

Table 99 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)

[“Reducing client file index size” on page 591](#) provides information about reducing the size of the client file indexes by using the Remove Oldest Cycle button.

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.

To query the media database:

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 indicating that no save sets were found that matched the specified query. Click **OK** to close the dialog box.

You can also query the media database by using the **mminfo -av** command. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide a complete description of the **mminfo** command and its options.

Cross-checking client file indexes

By cross-checking, you can 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 crashes before the NetWorker server has completely updated the indexes. Once the server is running again, cross-check to accurately update the online indexes.

To cross-check a client file index:

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 NetWorker server displays this prompt:

```
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.

To refresh the index information:

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

To designate a nondefault client file index location when creating a new client:

1. From the **Administration** window, click **Configuration**.
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, enter the full path of the directory where the client file index will reside.
5. For the remaining tabs, enter information as necessary to create the new client. [“Setting up a scheduled backup” on page 59](#) provides instructions.
6. Click **Ok**.

Changing the client file index location for an existing client

NOTICE

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” on page 590](#) provides more information.

To change the client file index location to a nondefault location for an existing client:

1. From the **Administration** window, click **Configuration**.
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, enter 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
```

The resulting output will be similar to:

```
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.

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.

To move an existing client file index:

1. Ensure that no backup is currently running on the NetWorker server.
2. Copy the client file index from its current location to the new location. For example:
 - For Windows, copy `<NetWorker_install_dir>\index\jupiter` to `<new_location>\indexes\jupiter`.
 - For UNIX, type this line at a command prompt:


```
cp -rp /nsr/index/jupiter /new_location/indexes
```
3. Update the **Index Path** attribute of the **Client** resource to point to the new location of the index. [“Changing the client file index location for an existing client” on page 590](#) provides instructions.

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.

Reducing client file index size

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 592](#) provides details.
- ◆ Delete volume-based entries from the client file index. [“Deleting volume-based online index entries” on page 593](#) 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 browse policy associated with a particular save set by using the `nsrmm -w` 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. [“Modifying the browse and retention policy on a save set” on page 286](#) 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 590](#) provides details.

Reducing media database size

Reduce the size of media database on the NetWorker server by using one or more of these methods:

- ◆ Remove volumes that contain recyclable save sets from the NetWorker inventory. [“Removing volume-based entries from the online indexes” on page 592](#) provides details.

When a volume is removed from the media database, the entries associated with that volume are removed from the media database and the online file index on the client. If you select this option, you will still be able to recover the data on the volume by using the **scanner** program.

NOTICE

Very little disk space is gained by 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 a volume’s mode” on page 222](#) provides details.

When a volume is recycled, the NetWorker server performs these procedures:

- Relabeling
- Deletion from the media database
- Reinitialization

Once a volume is recycled, its content cannot be recovered.

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. [“Modifying the browse and retention policy on a save set” on page 286](#) provides details.

- ◆ Compress the media database. [“Compressing the media database” on page 594](#) provides details.

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. Removing the oldest cycle frees up disk space.

To remove the oldest save set cycles:

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.

After the Remove Oldest Cycle operation has finished, the statistics in the Index Save Sets dialog box are updated to reflect the current state of the client file index.

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

You can remove just the entries contained in the client file index by using the **nsrmm** command. This changes the status of the browsable save sets to recoverable. At the command prompt, enter:

```
nsrmm -d -P -S ssid
where ssid is the save set ID for the save set.
```

Use **mminfo** to determine the save set ID. At the command prompt, type:

```
mminfo -v -c client_name
```

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide information on the **nsrmm** and **mminfo** entries.

When save sets are marked recoverable, users may no longer browse to recover these files. Use the save set recover procedure to recover data when the client file index entries are removed. [“Recovering the data” on page 373](#) provides details.

Removing client file index and media database entries

You can remove both the client file index and media database entries. 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, and the volume is undamaged, the data remains recoverable by using the **scanner** program. [“Recovering expired save sets” on page 398](#) 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 the 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. The entry of the volume in the media database is never actually purged. 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

To delete volume-based entries from the online indexes:

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.** [“Removing client file index entries” on page 593](#) provides details about this option.
 - **File Index Entries Only.** [“Removing client file index and media database entries” on page 593](#) provides details about this option.
5. Click **OK**.

The NetWorker server first cross-checks the indexes before it purges a volume. As a result, the volume might still appear in the Volumes window for a brief period of time.

Deleting volume based index entries from the command line

You can also remove online index entries with the `nsrmm` program. To remove both client file index and media database entries for a particular save set, enter this line at the command prompt:

```
nsrmm -d -S ssid
```

To remove all information related to a particular volume, enter this line at the command prompt:

```
nsrmm -d volume_name
```

Compressing the media database

You can free up more space on the server by compressing the media database.

To compress the media database:

1. Delete the appropriate file:

- On Windows:

```
NetWorker_install_dir\mm\.cmprssd
```

- On UNIX:

```
/nsr/mm/.cmprssd
```

2. Type at the command prompt:

```
nsrim
```

Managing Client Push

The software distribution feature, Client Push distributes software and performs software updates to one or more NetWorker hosts from the NetWorker server.

The *NetWorker Installation Guide* and the **nsrpush** man page describes how to use Client Push to update NetWorker products.

This section covers these maintenance tasks:

- ◆ [“Changing the location of the software repository” on page 594](#)
- ◆ [“Removing software package information from the software repository” on page 595](#)
- ◆ [“Transferring files and folders by using nsrpush” on page 597](#)

Changing the location of the software repository

Use the **nsradmin** command to change the location of the software repository.

1. Log in to the NetWorker server as root on UNIX or administrator on Windows.

2. Connect to the nsrccd database:

- UNIX: **nsradmin -d /nsr/res/cpdb**
- Windows: **nsradmin -d “*NetWorker_install_path*\nsr\res\cpdb”**

3. Set the current query type to **NSR Client Push Master**:

```
nsradmin> . type: NSR Client Push Master

Current query set
```

4. Review the current attribute settings:

```
nsradmin> print

                                type: NSR Client Push Master;
                                name: Client Push Master;
    actual repository location: /nsr/repository;
    default repository location: /nsr/repository;
                                exclude clients: ;
```

5. Update the **actual repository location** value:

```
nsradmin> update actual repository location: /new_repository

    actual repository location: /new_repository;
Update? y
updated resource id 0.2.15.116.0.0.0.43.78.222.34.14.10.5.172.45(7)
```

6. Confirm the value of the **actual repository location** attribute:

```
nsradmin> print

                                type: NSR Client Push Master;
                                name: Client Push Master;
    actual repository location: /new_repository;
    default repository location: /nsr/repository;
                                exclude clients: ;
```

7. Quit **nsradmin**.
8. Kill the **nsrccd** process on the NetWorker server.

Removing software package information from the software repository

After you add new packages to the repository, you can remove old package information from the command line or from a GUI.

- ◆ [“Removing information from the repository by using the Software Administration Wizard” on page 595](#)
- ◆ [“Removing information from the repository by using the nsrpush command” on page 596](#)

Removing information from the repository by using the Software Administration Wizard

Use the Software Administration Wizard to remove information about software packages in the repository, from a GUI.

1. Connect to the NetWorker server from the **NetWorker Management Console**.
2. In the **Configuration** menu, select **Software Administration Wizard...**
3. Click **Next** in the **Welcome to the Software Administration Wizard** page.
4. In the **Select Operation** window, accept the default **Add or remove products from my software repository**, then click **Next**.

5. In the **Software Repository Operations** window, select **Remove software products from the repository**, then click **Next**.
6. In the **Select Products to Remove** window, select the products you want to remove, then click **Next**.
7. If the remove operation completes successfully, then click **Ok** when the pop up window appears. If the remove operation fails, review the **nsrccd.raw** file located in the `/nsr/logs` (UNIX) or `NetWorker_install_path\nsr\logs` (Windows) directory for further details.

Removing information from the repository by using the nsrpush command

Use **nsrpush** to remove information about software packages in the repository, from the command line.

Perform these steps from a command prompt on the NetWorker server as the administrator on Windows and the root account on UNIX:

1. Display a list of products that are in the software repository:

For example:

```
nsrpush -l

Products in the repository
=====

NetWorker 8.1

win_x64
Storage Node
Server
License Manager
Language Packs
English Language Pack
French Language Pack
Japanese Language Pack
Korean Language Pack
Chinese Language Pack
Client
Management Console
```

2. Remove the packages:

```
nsrpush -r -p Product -v Version -P platform
```

For example, to remove the NetWorker 8.1 win_x64 package:

```
nsrpush -r -p NetWorker -v 8.1 -P win_64
```

```
Remove from repository status: succeeded
```

If the remove operation fails, then review the **nsrccd.raw** file located in `/nsr/logs` on UNIX or `NetWorker_install_path\nsr\logs` on Windows for further details.

Transferring files and folders by using nsrpush

You can use **nsrpush** to transfer files and folders from a central location to the /nsr directory on a UNIX host or the *NetWorker_install_dir*\nsr directory on a Windows host.

- ◆ [“Requirements for file and folder transfers” on page 597](#)
- ◆ [“Transferring files and folders” on page 597](#)
- ◆ [“Troubleshooting file and folder transfers” on page 598](#)

Requirements for file and folder transfers

Before you transfer files and folders, ensure that the target host is in the Client Push inventory and there is sufficient free disk space in the tmp folder on the target host. Client push uses the c:\windows\temp folder on Windows and /tmp on UNIX.

When the operating system of the target host differs from the NetWorker server, for example, when the NetWorker server is on Windows and the target host is a UNIX, you must configure a proxy host to store the cross platform files and folders.

When selecting a proxy host, ensure that the host:

- ◆ Is the same platform as the cross platform packages.
For example, if the NetWorker server is a Linux host, use a Windows proxy host to transfer files to Windows x86, Windows x64, and Windows ia64 hosts.
- ◆ Has the NetWorker 7.6 or later client software installed.
- ◆ Is a client of the NetWorker server.

When choosing a directory on the proxy host to store the source files, ensure that the directory:

- ◆ Resides on a local file system.
- ◆ Uses a path that does not contain spaces or special characters.

Transferring files and folders

You can transfer files and folders from a central location to all NetWorker hosts in a datazone or selected hosts. You can specify hosts in the **nsrpush** command or by using an input file.

From a command prompt on the NetWorker server, use **nsrpush** to transfer files and folders.

Performing cross platform file and folder transfers

To transfer files from the NetWorker server to NetWorker hosts on a cross platform operating system, the syntax of the **nsrpush** command is as follows:

```
nsrpush -Tx -c proxy_host -C proxy_source_path -U|-W -all|-If  
input_file|hostname...
```

where:

- ◆ *proxy_host* is the hostname of the host that contains the source files and folders.
- ◆ *proxy_source_path* is the folder on the Proxy host that contains the source files and folders.

- ◆ **-U** specifies a UNIX cross platform host and **-W** specifies a Windows cross platform. Use the appropriate option for the target host.
- ◆ **-all** transfers the source files and folders to all inventoried NetWorker hosts that are not in the exclude list.
- ◆ **-IF *input_file*** transfers the source files and folder to all inventoried NetWorker hosts that are listed, one per line, in the input file. When specifying *input_file*, include the name of the file and the path to the file on the NetWorker server.
- ◆ *hostname* is the name of the target host. Separate multiple hostnames with spaces.

Example 55 Transferring files from a proxy host

A NetWorker datazone uses a NetWorker 8.1 server on Windows and has two NetWorker UNIX clients, `pwd.emc.com` and `lad.emc.com` that require new DD Boost libraries in the `/usr/bin` directory. The directory `/usr/ddlib/bin` on UNIX host `mnd.emc.com` contains the files.

To transfer the files, type:

```
nsrpush -Tx -c mnd.emc.com -C /usr/ddlib -U pwd.emc.com lad.emc.com
```

Performing same platform file and folder transfers

To transfer files from the NetWorker server to same platform NetWorker hosts, the syntax of the **nsrpush** command is as follows:

```
nsrpush -Tx -m source_path -all|-If input_file|hostname...
```

where:

- ◆ *source_path* specifies the path on the NetWorker server that contains the source files.
- ◆ **-all** transfers the source files and folders to all inventoried NetWorker hosts that are not in the exclude list.
- ◆ **-IF *input_file*** transfers the source files and folder to all inventoried NetWorker hosts that are listed, one per line, in the input file. When specifying *input_file*, include the name of the file and the path to the file on the NetWorker server.
- ◆ *hostname* is the name of the target host. Separate multiple hostnames with spaces.

Example 56 Transferring files from the NetWorker server

A NetWorker datazone uses a NetWorker 8.1 server on Windows and has two Windows clients, `dmd.emc.com` and `jad.emc.com` that require new DD Boost libraries. The directory `c:\ddlib` on the NetWorker server contains the files.

To transfer the files, type:

```
nsrpush -Tx -m c:\ddlib dmd.emc.com jad.emc.com
```

Troubleshooting file and folder transfers

This section describes how to troubleshoot file and folder transfer issues.

- ◆ [“Transfer media path doesn't exist: pathname” on page 599](#)

Transfer media path doesn't exist: pathname

This error message appears when the *source_path* or *proxy_source_path* specified in the `nsrpush` command does not exist on source or proxy host. To resolve this issue, ensure that you specify a valid path.

Monitoring Changes to NetWorker Server Resources

The Monitor RAP (resource allocation protocol) attribute in the NetWorker Server resource tracks both before and after information related to additions, deletions, or modifications to NetWorker server resources and their attributes. NetWorker records these changes in the `rap.log` file, located in the `NetWorker_install_dir\logs` directory. The `rap.log` file lists the username, the source computer, and the time of the change. NetWorker logs sufficient information in the `rap.log` file to enable an administrator to undo any changes.

In NetWorker 8.0 and later, the Security Audit Log feature provides the NetWorker server and the NMC Console server with the ability to log specific security audit events related to their operations. [“NetWorker Accountability” on page 934](#) provides more information about the Security Audit Log feature and how to configure it.

How to disable/enable the Monitor RAP Attribute

The Monitor RAP attribute is enabled by default, to change the attribute's setting:

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**.

Log file size management

NetWorker stores messages generated by the NetWorker daemons or the Console server daemons in raw files. [“Viewing log files” on page 803](#) describes how to view and render raw log files.

To manage the size of raw log files, modify attributes stored in the NSRLA database. [Table 100 on page 600](#) summarizes the attributes in the NSRLA database that manage log file sizes.

Table 100 Attributes used to manage raw 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. When you configure the runtime rendered log attribute, NetWorker trims the runtime rendered log file and the associated raw file simultaneously. Default: disabled
runtime rollover by time	When set, this attribute invokes 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 raw file simultaneously. Default: undefined

How the trimming mechanism trims the log files differs depending on the how you define the log file size management attributes. [Table 101 on page 601](#) summarizes the behavior of the trimming mechanism.

Table 101 Trimming behavior

Attribute configuration	Trimming behavior
runtime rollover by time or runtime rollover by size is configured	<ol style="list-style-type: none"> 1. NetWorker copies the contents of the existing log file to a new daemon <code>date_time.raw</code> or <code>rapdate_time.log</code> file. 2. NetWorker truncates the existing <code>daemon.raw</code> or <code>rap.log</code> file to 0 MB. <p>Note: When this mechanism starts on a NetWorker server that is under a heavy load, this process may take some time to complete.</p>
runtime rollover by time or runtime rollover by size is not configured	<ol style="list-style-type: none"> 1. NetWorker checks the log file size when the <code>nsrexecd</code> process starts on the computer. 2. When the log file size exceeds the size defined by the maximum size MB attribute, NetWorker renames the existing log file to <code>daemon_date_time.raw</code> or <code>rapdate_time.log</code>. NetWorker creates a new empty <code>daemon.raw</code> or <code>rap.log</code> file. <p>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 be much larger than the value defined by maximum size MB.</p>

Example 57 Enabling runtime rollover by time

To enable **runtime rollover by time** at 12:34 AM for the `gstd.raw` file on the console server:

1. Set the current query to the **gstd.raw** log file:

```
nsradmin> . type:NSR log;name:gstd.raw
```

2. Update the **runtime rollover by time** attribute:

```
nsradmin> update runtime rollover by time: "00:34"
runtime rollover by time: "00:34";
Update? y
updated resource id 0.14.5.28.0.0.0.31.78.125.11.176.10.5.172.45(3)
```

3. Verify the attribute value:

```
nsradmin> print

type: NSR log;
administrator: root, "user=root,host=bu-t3-7.lss.emc.com";
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: /opt/lgtonmc/logs/gstd.raw;
```

4. Type **Quit** to exit the `nsradmin` prompt.

Internationalization

NetWorker releases 7.4 and later have been internationalized. As a result, the NetWorker software now supports language packs, which can be installed as part of the NetWorker installation, or can be installed separately after the NetWorker software has been installed. The *NetWorker Installation Guide* provides more information.

Internationalization support in the NetWorker software is dependent on the underlying operating system's internationalization support. If you are planning on using non-English data in the NetWorker software, make sure the appropriate support for that language has been installed and configured on the operating system.

There are a number of issues and limitations related to the use of NetWorker software in a multi-language environment, which are addressed in the following sections.

Log file viewer

NetWorker log files must be viewed using the `nsr_render_log` program. [“Viewing log files” on page 803](#) provide more information.

Interoperability with previous releases of NetWorker

Multiple locales within the same datazone are fully supported only if all NetWorker installations in the datazone are at release level 7.4. The following limitations apply:

- ◆ In datazones with a NetWorker release 7.3 server and NetWorker 7.4 clients, support for scheduled backups of path or filenames containing non-ASCII characters is limited to the support provided by NetWorker release 7.3.
- ◆ Binaries from releases earlier than 7.4 may not correctly display UNIX save sets containing non-ASCII characters.

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 not supported by the user's current locale will display incorrectly. 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 be displayed correctly.

Character display in graphical user interfaces

Character display from within NetWorker graphical user interfaces varies depending on the platform on which the graphical user interface is running:

- ◆ Microsoft Windows:
 - Any Unicode encoded data will be displayed correctly.
 - When viewing UNIX path and filenames, path and filenames created by using a character set supported by the current locale or UTF-8 will be displayed correctly. Paths created by using another character set may not display correctly. Because

Microsoft Windows does not have native support for many 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, it will not be displayed correctly on Microsoft Windows.

- ◆ Unix:
 - Characters not supported by current locale may not be displayed correctly.
- ◆ Mac OS X
 - Because of differences in Unicode support, non-ASCII paths and filenames on Mac OS X machines may not display correctly when browsing the filesystem from a non-Mac platform.

Maximum path and save set length

For the NetWorker software, the maximum supported length for a pathname is twelve kilobytes, and the maximum length for a save set name is 1024 bytes. The number of characters allowed by each of these limits will vary depending on the locale.

All operating systems have an internal limit for path and filenames. The limit varies depending on the operating system and file system being used. 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 filename that is greater than the limit supported by the operating system, but an attempt to access this path will result in a failure.

Locale-specific configuration issues on UNIX/Linux

This section describes certain configuration issues related to Client and Archive Request resources for clients that run in UNIX/Linux environments.

Configuring the Save Set attribute for Client and Archive Request resources

For clients that run on non-ASCII locales on Unix platforms, or for clients on Microsoft Windows that are being configured from UNIX hosts that use non-ASCII locales, special considerations apply when typing a path or filename in the Save Set attribute for Client or Archive Request resources. The path or filename must be entered in the locale that was used when the path or file was created. If you enter a path or filename by using a locale other than the one that was used when the path or file was created, subsequent backups will fail with the following error message:

```
No such file or directory
```

To configure a Client resource in this situation, either:

- ◆ Use the **All** keyword for the **Save Set** attribute.
- ◆ Log into a client host by using the correct locale, and configure the client from this machine.

Client resources with multiple locales on UNIX/Linux

To back up a UNIX or Linux machine that contains path or filenames with multiple locales, you must create a separate Client resource for each locale.

For example, to configure a multi-locale UNIX machine with data in both Japanese and French, you must create two different Client resources, one defining the save sets for the Japanese data, another defining the save sets for the French data.

[“Multiple clients from the same computer” on page 622](#) has information about creating multiple clients.

Locale settings with NDMP

When running NDMP backups, the locale setting has to be consistent in your environment. All UNIX flavored locale settings on the filer (including UTF-8) must be the same and the NMC client can be run only on a UNIX client set to the exact same locale setting as the filer.

Backup and recovery operations can be run on any locale, but if you try to browse on a locale that is different from the original locale the filenames appear as random characters.

CHAPTER 19

NetWorker Client Management

This chapter covers these topics:

- ◆ NetWorker client overview 606
- ◆ Client configuration..... 606
- ◆ Creating a client probe 610
- ◆ Associating a probe with a Client resource 611
- ◆ Creating a lockbox to store and retrieve pass phrases securely 611
- ◆ NetWorker authentication 612
- ◆ Multiple clients from the same computer..... 622
- ◆ Scheduled backups of non-ASCII files or directories 623
- ◆ Controlling access to a NetWorker client 624
- ◆ Client priority 625
- ◆ Dedicated client/server interface for backup and recover operations..... 626

NetWorker client overview

A NetWorker client is both a physical computer with NetWorker client software installed on it, and a NetWorker server resource that specifies a set of files and directories to be included in a scheduled backup. As such, a single NetWorker client computer can have several Client resources specified that can back up to the same or even different NetWorker servers.

[“Defining a client and save set combination” on page 623](#) further explains concept of a client computer with multiple NetWorker Client resources.

NetWorker client software is available for a variety of platforms. No matter which platform the client resides on, it can be backed up to any NetWorker server. For example, you can back up a NetWorker client on a Microsoft Windows computer to a NetWorker server on a Solaris computer.

Client configuration

Before a client can be backed up by a NetWorker server, the client computer must have the appropriate NetWorker client software installed. The NetWorker Installation Guide provides more information.

Creating a client

NetWorker client creation is part of the process of creating a scheduled backup. [“Setting up a scheduled backup” on page 59](#) provides information on how to create a client.

Editing a client

Use this procedure to change Client resource attributes. You cannot change the name of a client with this procedure. [“Changing a client name” on page 607](#) provides information on how to change a client name.

To edit a client:

1. From the **Administration** window, click **Configuration**.
2. Click **Clients**.
3. Right-click a client, then select **Properties**. The **Properties** dialog box appears, with the **General** tab displayed.
4. Edit the attributes of the client, then click **OK**.

Copying a client

Use this procedure to copy Client resource attributes.

To copy a client:

1. From the **Administration** window, click **Configuration**.
2. Click **Clients**.

3. Right-click a client, then select **Copy**. The **Create Client** dialog box appears. By default, the new client retains the information from the client that was copied, except the Name attribute, which is blank.
4. Enter the name for the new client and edit other attributes as appropriate, then click **OK**.

Changing a client name

The only way to change the name of a client is to re-create the client.

1. From the **Administration** window, click **Configuration**.
2. Click **Clients**.
3. Right-click the client to be renamed, then select **Properties**. The **Properties** dialog box appears.
4. Click the **Globals (1 of 2)** tab.
5. Record the **Client ID** attribute listed for the client, then click **Cancel** to close the **Properties** dialog box.
6. Delete the original client from the **Administration** window. See [“Deleting a client” on page 607](#).
7. Stop all NetWorker services.
8. On the NetWorker server that backs up this client, rename the directory containing the client file index for this client from *old_client_name.domain.com* to *new_client_name.domain.com*. The default location for the client file index is:
 - For UNIX/Linux:


```
/nsr/index/client_name.domain.com
```
 - For Microsoft Windows:


```
<NetWorker_install_path>\index\client_name.domain.com
```
9. Restart the NetWorker services.
10. Create a new client, making sure that you enter the client ID that you recorded in [step 5](#) in the **Client ID** attribute of the **Globals (1 of 2)** tab of the **Create Client** dialog box. [“Setting up a scheduled backup” on page 59](#) provides instructions on how to create a client.

Deleting a client

When a client is deleted, the NetWorker server can no longer back up or recover files from the client computer. The backup history for the client remains in the client file index and media database until the entries are removed.

To delete a client:

1. From the **Administration** window, click **Configuration**.
2. Click **Clients**.

3. Right-click the client to be deleted, then select **Delete**.
4. When prompted, click **Yes** to confirm the deletion.

Note: Even if you delete a client from the NetWorker server, the previously backed-up data from the client is still accessible and recoverable directly from the volume that contains the data by using the **scanner** command.

Recovering a deleted client

To recover a deleted client, create a new client, making sure that you enter the name of the deleted client in the **Name** attribute on the **General** tab of the **Create Client** dialog box. [“Task 6: Create a backup Client resource” on page 64](#) provides instructions on how to create a client.

The NetWorker server recalls the client ID for this name and inserts it into the **Client ID** attribute on the **Globals (1 of 2)** tab of the **Create Client** dialog box.

Editing a client NSRLA database

The NSRLA database contains a NetWorker resource, called the NSRLA resource, which has attributes that apply to the client, such as the Disable Directed Recover attribute. In some cases, it may be necessary to edit the NSRLA resource. The NSRLA resource can be edited by using the character-based **nsradmin** program.

Note: [“Directed recoveries” on page 369](#) provides information about permissions for directed recoveries and the Disable Directed Recover attribute in the NSRLA resource.

To edit the NSRLA database:

1. Log in as root or as Windows Administrator on the NetWorker client.
2. Type this 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
```


- To change the value of attributes in the **NSRLA** resource, type this line at the **nsradmin** prompt:

```
update attribute:value;
```

For example, to update the **Disable Directed Recover** attribute, type:

```
update disable directed recover:Yes
```

- Type **Yes** when prompted to confirm the change.

NOTICE

When modifying an attribute with the **nsradmin** program, the attribute name and value must be specified correctly. If the attribute name and value are not specified correctly, the attribute is not updated. No error message is provided.

Restricting access to view the client NSRLA database

By default, a UNIX root user on any host can view the NSRLA database on any UNIX NetWorker client. In the same way, any Windows administrator can view the NSRLA database settings on any Windows NetWorker client. You can restrict access so that only the root (UNIX) or Administrator (Windows) on the local host can view the NSRLA database.

To restrict the ability to view the NSRLA database:

- Log in as root or as Windows Administrator on the NetWorker client.
- Type the following at the command prompt:

```
nsradmin -p nsrexec
```

The **nsradmin** prompt appears.

- Determine the current settings for the attributes in the NSRLA resource:

```
print type:NSRLA
```

By default the Administrator attribute will look similar to the following:

Windows

```
administrator: Administrators, "group=Administrators,host=saturn";
```

UNIX

```
administrator: root, "user=root,host=saturn";
```

- Change the value of the Administrator attribute to remove the single *Administrators* or *root* value:

Windows

```
update administrator: "group=Administrators,host=this_host"
```

UNIX

update administrator: "user=root,host=*this_host*"

where *this_host* is the name of the local host.

5. Type **Yes** when prompted to confirm the change.

Now, only the Windows Administrator or root on the local host can view the NSRLA database.

Creating a client probe

NOTICE

Users are responsible for creating and supporting user defined probes.

To create a client probe:

1. From the **NetWorker Administration** window, click **Configuration**.
2. Right-click probes, and select **New**. The **Create NSR probe** window opens.
3. Type the name of the probe in the **Name:** field.
4. (Optional) Include details for the probe script in the **Comment:** field.
5. Type the name, and path of the probe script in the **Command:** field. The probe resource script must be placed in the same directory as the nsr binaries for each client referencing the probe. For example, /usr/sbin on Solaris and Linux. The *NetWorker Installation Guide* provides details on default installation paths for all operating systems.

A user defined probe is any program which passes a return code. Return codes are interpreted by NetWorker as follows:

- Return code 0: backup is required
- Return code 1: backup is not required
- All other return codes are interpreted as errors during probe execution, and backups are not taken

Note: The probe script name must begin with save, or nsr.

6. Associate the probe with a Client resource as described in [“Associating a probe with a Client resource” on page 611](#).

Note: The **Command Options:** field is applies to NetWorker Module probes only.

Associating a probe with a Client resource

To associate a probe resource with a Client resource:

1. Click **Clients**, and right-click the client in the **Configuration** screen of the **NetWorker Administration** window.
2. Select **Properties**, and the **Properties** window opens.
3. In the **Apps & Modules** tab, select the probe from the **Probe resource name:** list. All defined probe resources appear in the list. If the list is empty, then probe resources were not defined. [“Creating a client probe” on page 610](#) provides details on creating a probe resource.

Note: Each client in a group can have a probe associated with it, but a probe is not required. However, a probe-based backup group must have at least one probe-enabled client associated with it. If a single probe is referenced by multiple clients, make sure that the probe resource script is in the same directory as the NetWorker binaries for each of the clients referencing the probe.

Creating a lockbox to store and retrieve pass phrases securely

The NetWorker server provides NetWorker modules, and certain NetWorker features, such as AES encryption, with the ability to securely store and retrieve passwords over a network by using a NetWorker lockbox.

To use the lockbox pass phrase feature for a NetWorker module:

1. Log in as root, or as Windows administrator on the NetWorker server.
2. In the client **NetWorker Administration** window, click the **Configuration** tab.
3. Right-click **Lockboxes**, and select **New**. The **Create NSR Lockbox** window opens.
4. Type a name in the **Name:** field.
5. Type the users that will have permission to store, retrieve, and delete passwords in the **Users:** field. Use the `user@destination_client` or `user@hostname` format when listing users.
6. Verify that the lockbox is listed in the **NetWorker Administration** window.

Set the Datazone pass phrase for a NetWorker server

[“Set the Datazone pass phrase for a NetWorker server” on page 108](#) provides information on how to set the Datazone pass phrase on a NetWorker server.

Set the Remote access attribute for NetWorker client resource when a lockbox is created for a NetWorker server on a cluster

If a lockbox is created for a NetWorker server configured on a cluster, the following step must be performed for the lockbox to work on the cluster node that the NetWorker server is running on:

- ◆ The **Host** section for the NetWorker client resource's **remote access attribute** must list the names of all the cluster nodes. The values of the names must be in the following form:
 - For UNIX cluster nodes, use the hostname command output
 - For Windows cluster nodes, use the full computer name (go to **Control Panel > System > Computer name** to get the full name)

Note: The lockbox can be created prior to or after performing this step.

This configuration must be performed during the initial startup of the virtual NetWorker server (that is, before the virtual server fails over to the other cluster). Otherwise, the lockbox will have to be deleted and then re-created.

Error messages and error handling for the Datazone pass phrase

The following messages are written to the daemon.log file if an error occurs during the encryption, or decryption of the Datazone pass phrase:

- ◆ Encryption error:

```
Error encrypting key. Erasing datazone pass phrase.
```

The Datazone pass phrase is saved as a blank phrase.

- ◆ Decryption error:

```
Error decrypting key. Erasing datazone pass phrase for current use.
```

A blank pass phrase is used in place, allowing any backup requiring a key to continue by using the default key.

NetWorker authentication

This section describes how to configure authentication between NetWorker hosts. This section also includes special considerations for authentication.

Two types of authentication are supported for NetWorker hosts:

- ◆ nsrauth
- ◆ oldauth

Strong authentication (nsrauth)

The **nsrauth** authentication mechanism (enabled by default) is strong authentication that is based on the Secure Sockets Layer (SSL) protocol, which is provided by the OpenSSL library. NetWorker hosts and NetWorker user permissions are authenticated by using **nsrauth**. The **nsrauth** authentication mechanism is available for hosts that run NetWorker software release 7.3 or later.

Each NetWorker host has a **nsrexecd** service, which provides authentication services. Each **nsrexecd** has its own private key and self-signed certificate for authentication. The private key is generated by **nsrexecd** when it starts up or one can be loaded from a file. The corresponding self-signed certificate is generated by the private key. The private key is RSA and is 1024 bits in length. The encryption method that is used once an SSL session is set up is AES-128. The session information sent over the SSL connection includes:

- ◆ Session keys
- ◆ Session ID
- ◆ User's information
- ◆ User's NetWorker permissions

This product should only be used where network-layer authentication prevents rogue hosts from connecting or accessing network traffic.

Authentication for backwards compatibility (oldauth)

For compatibility with earlier NetWorker releases, **oldauth** authentication is supported. If two hosts cannot authenticate by using strong authentication (**nsrauth**), you can enable authentication by using **oldauth**. One can specify the minimum authentication strength that is allowed for any host relationship. [“Specifying the minimum authentication strength between hosts” on page 614](#) provides more information.

Access privileges for authentication configuration

This section describes how to set up access privileges to maintain **nsrauth** strong authentication configuration settings.

The Console server must have permission to update resources on each NetWorker host whose authentication information will be updated:

To grant authentication update privileges to the Console server:

1. Log in as root or as Windows administrator on the NetWorker host whose authentication information will be updated.
2. Type the following at the command prompt:

```
nsradmin -p nsrexec
```

The **nsradmin** prompt appears.

3. For each of the following resources:

- NSRLA
- NSR Peer information
- NSR system port ranges

Complete the following steps:

- a. Determine the current settings for the attributes in the resource:

```
print type:resource_name
```

- b. Change the value of the Administrator attribute in the resource:

```
update administrator:current_values,"user=Console_user,host=
Console_host"
where:
```

- *current_values* are the values that are currently listed for the resource’s Administrator attribute.
- *Console_user* is the user ID of the Console user.
- *Console_host* is the name of the Console host.

- c. Type **Yes** when prompted to confirm each change.

Specifying the minimum authentication strength between hosts

You can specify that only certain authentication methods be allowed between specific hosts. For example, one could specify:

- ◆ That a NetWorker release 7.3 or later server be allowed to authenticate with legacy NetWorker clients by using **oldauth**.
- ◆ That all other NetWorker release 7.3 or later clients must use only **nsrauth** Strong authentication (enabled by default).

The Console server user must have permission to update authentication resources. [“Access privileges for authentication configuration” on page 613](#) provides more information.

To specify the authentication strength between NetWorker hosts:

1. From the **Administration** window, click **Configuration**.
2. Open the **Local Hosts** folder.
3. From the **Local Hosts** list, right-click the host with the authentication relationships to be configured and select **Configure Local Agent**.
4. Select the **Advanced** tab.
5. In the **Auth Methods** attribute, at the top of the list enter the minimum allowable authentication strength for the NetWorker hosts that will connect to this host. Type values in this format:

clientgroup_or_host, auth_strength

where:

- *clientgroup_or_host* is either an IP address that represents a group of hosts or an explicit hostname. For example:
 - *pluto.company.com* - An explicit NetWorker client name.
 - *10.102.0.0/255.255.0.0* - A subnet IP address representing all NetWorker clients on the subnet. Alternatively, you could type this value as *10.102.0.0/16*.
 - *0.0.0.0/0* - This value represents all clients in the domain,
- *auth_strength* is the authentication strength. The allowable values are:
 - **nsrauth** - Use strong authentication only.
 - **oldauth** - Use oldauth authentication only.
 - **nsrauth/oldauth** - Attempt to use strong authentication. If strong authentication fails, use **oldauth** authentication.

For example, if all hosts must use only strong authentication, enter:

0.0.0.0/0,nsrauth

NOTICE

The order in which values appear in the Auth Methods list is important. The first client match found starting from the top of the list is the authentication value used. For example, to specify strong authentication for all clients except for one, ensure that the explicit entry

for the single client appears at the top of the list before the more general entry that represents all clients in the domain. If no matches are made, authentication defaults to 0.0.0.0/0, nsrauth, which means that all clients can only authenticate by using nsrauth. [“Enforcing strong authentication on selected hosts” on page 615](#) provides details.

6. Click **OK**.
7. Restart the NetWorker services for the host with the authentication relationships to be configured. [“Stopping and starting a NetWorker server, client, or storage node” on page 55](#) provides information about restarting NetWorker services.

Example 58 Enforcing strong authentication on selected hosts

In this example, a NetWorker release 7.3 server has both release 7.3 clients and legacy clients. The requirement is to ensure that the release 7.3 clients authenticate by using only **nsrauth** strong authentication, while the legacy clients authenticate by using **oldauth** authentication.

1. From the **Administration** window, click **Configuration**.
2. Open the **Local Hosts** folder.
3. From the **Local Hosts** list, right-click the NetWorker server and select **Configure Local Agent**.
4. Select the **Advanced** tab.
5. In the **Auth Methods** attribute, type the two legacy hostnames, for example:

jupiter.company.com, nsrauth/oldauth

pluto.company.com, nsrauth/oldauth

In this example, two legacy hosts identified as *jupiter.company.com* and *pluto.company.com* are allowed to use **oldauth** authentication to ensure backwards compatibility. All other hosts use only strong authentication (**nsrauth**).

6. To specify that all other hosts must use strong authentication, type this value at the bottom of the list in the **Auth Methods** attribute:

0.0.0.0/0, nsrauth

The entries in the list should now be ordered as follows:

pluto.company.com, nsrauth/oldauth

jupiter.company.com, nsrauth/oldauth

0.0.0.0/0, nsrauth

The more explicit entries, those that identify a single client, are at the top of the list. The first matching entry, starting from the top of the list, is the entry used to specify a host's authentication strength.

7. Click **OK**.
8. Restart the NetWorker services for the NetWorker server. [“Stopping and starting a NetWorker server, client, or storage node” on page 55](#) provides information about restarting NetWorker services.

Maintaining NetWorker local host authentication credentials

Each NetWorker host that uses **nsrauth** strong authentication has unique credentials that are used to identify itself to other NetWorker hosts during the **nsrauth** strong authentication process. These credentials are known as local host authentication credentials.

Additionally, each NetWorker host maintains a copy of the local host credentials that belong to every NetWorker host to which it has authenticated. These credentials are maintained as part of the local host's Peer resources.

You can export, import, or create new credentials for a local host. You can also delete and import host credentials in Peer resources. In most cases, you do not need to perform any of these maintenance tasks. However, there are some cases in which you may need to perform these tasks. These examples provide some scenarios in which these maintenance tasks might be performed.

Note: To complete the procedures in this section, the Console server user must have permission to update authentication resources. [“Access privileges for authentication configuration” on page 613](#) provides details.

Example 59 Securing the initial peer host authentication process

When two NetWorker hosts authenticate with each other, no user intervention is required. Each host keeps a copy of the other host's authentication credentials in a Peer resource. The copy is created the first time the hosts authenticate with each other. Subsequent authentication attempts are verified by matching the host with the authentication credentials stored in the Peer resource.

To eliminate the possibility that an attacker could compromise this process, manually update the Peer resource with authentication credentials rather than have this occur automatically.

For example, suppose that the NetWorker server *jupiter* will authenticate with the NetWorker client *pluto*.

To manually update the **Peer** resources:

1. Export authentication credentials.
 - a. Export authentication credentials for *jupiter* to a file.
 - b. Export authentication credentials for *pluto* to a file.

[“Exporting local host credentials” on page 618](#) provides information about exporting authentication credentials.
2. Create a **Peer** resource for *pluto* on *jupiter*:
 - a. Open the credentials file for *pluto*. Refer to this file in these steps.
 - b. From the **View** menu in the **Administration** window, select **Diagnostic Mode**.
 - c. Open the **Local Hosts** folder.
 - d. From the **Local Hosts** list, click *jupiter*.
 - e. Right-click and select **New**.

- f. In the **Name** attribute, enter the Name value from the credential file.
 - g. In the **Instance ID** attribute, enter the **NW Instance ID** value from the credential file.
 - h. In the **Peer Hostname** attribute, enter the **My Hostname** value from the credential file. For example:


```
certificate:\
"-----BEGIN CERTIFICATE-----
-----END CERTIFICATE-----";
my hostname: pluto.company.com;
name: pluto.company.com;
NW instance ID: d9b61002-0004-fe5c3a37-42d5a842-00010000-8945657f;
private key:\
"-----BEGIN RSA PRIVATE KEY-----
-----END RSA PRIVATE KEY-----";
```
 - i. Click **OK**. A **Peer** resource for pluto is created.
3. Create a local host certificate file for pluto. [“Creating a local host Peer resource certificate” on page 620](#) provides details.
 4. Load the certificate for pluto into the **Peer** resource:
 - a. Right-click the **Peer** resource for pluto and select **Properties**.
 - b. From the **Change Certificate** attribute list, select **Load** certificate from file.
 - c. In the **Certificate file to load** attribute, enter the path and name of the certificate file and the click **OK**.
 5. Repeat steps 2 through 4 to create a **Peer** resource for jupiter on pluto.

Example 60 Local host credentials are exported

In most cases, local host credentials are retained. For example, when the NetWorker software is uninstalled and reinstalled. However, in some cases, such as with unexpected data loss or corruption, new local host credentials may need to be re-created.

Creating new local host credentials can be time-consuming, because Peer resources on all NetWorker hosts that authenticate with the local host must be updated with the new local host credentials. To avoid these updates and to save time, export the local host credentials to a file for safekeeping. Store this file in a protected location, such as on a USB stick in a safe and not on a machine attached to the network. If necessary, re-import these local host credentials.

In this example, the user can either create new local host authentication credentials or save the existing credentials to a file and re-import the credentials at a later time. To save time and simplify the procedure, the user chooses to export the local host credentials to a file so that the credentials can be re-imported, if necessary.

To complete this scenario:

1. Export the local host credentials to a file. [“Exporting local host credentials” on page 618](#) provides details.
2. An event occurs in which the local host credentials are lost.
3. Re-import the host credentials to the local host mentioned in [step 1](#) . [“Importing local host credentials” on page 619](#) provides details.

There is no need to update Peer resources because the local host authentication information was recovered from the file that was exported in [step 1](#) .

Example 61 Local host credentials are not exported

This scenario is similar to [“Local host credentials are exported” on page 617](#), except that the local host credentials were not exported before an event such as the corruption of local host credential data. In this scenario, you must create new local host credentials and then delete all existing Peer resources for the local host.

To complete this scenario:

1. Create new local host credentials for the NetWorker host. [“Creating new local host certificate keys” on page 619](#) provides details.
2. Delete all peer certificates that are set up for the NetWorker host whose credentials were created in [step 1](#) . [“Deleting NetWorker local host Peer resources” on page 620](#) provides details.

Exporting local host credentials

To export local host credentials:

1. From the **Administration** window, click **Configuration**.
2. Open the **Local Hosts** folder.
3. From the **Local Hosts** list, right-click the NetWorker server and select **Configure Local Agent**.
4. Select the **Advanced** tab.
5. From the **NW Instance Info Operations** attribute list, select **Export**.
6. In the **NW Instance Info File** attribute, enter a directory and name for the credential file.
7. Click **OK**. A credential file is saved to the location specified.

Importing local host credentials

Note: On UNIX platforms, the credential file that is imported in this procedure must be set to read and write permission for root only. For example:

```
chmod 600 certificate_name
```

To import local host credentials:

1. From the **Administration** window, click **Configuration**.
2. Open the **Local Hosts** folder.
3. From the **Local Hosts** list, right-click the NetWorker server and select **Configure Local Agent**.
4. Select the **Advanced** tab.
5. From the **NW Instance Info Operations** attribute list, select **Import**.
6. In the **NW Instance Info File** attribute, enter the directory and name for the credential file to be imported.
7. Click **OK**.

The credential file will be imported from the location specified.

Creating new local host certificate keys

To create new local host certificate keys:

1. From the **Administration** window, click **Configuration**.
2. Open the **Local Hosts** folder.
3. From the **Local Hosts** list, right-click the NetWorker host and select **Configure Local Agent**.
4. Select the **Advanced** tab.
5. From the **NW Instance Info Operations** attribute list, select **New Keys**.
6. Click **OK**. A new certificate will be created for the NetWorker host.
7. On all other local host resources, delete the Peer resource that was set up for the NetWorker host whose certificate was changed in this procedure. [“Deleting NetWorker local host Peer resources” on page 620](#) provides details.

Note: Consider exporting the new certificate information to a credentials file. [“Local host credentials are exported” on page 617](#) provides information about the benefits of doing so.

Maintaining local host Peer resources

Each NetWorker host has unique authentication credentials that are used during the **nsrauth** strong authentication process. These credentials include certificate information.

Each NetWorker host retains a record of the certificate of all NetWorker hosts to which it has successfully authenticated. These records are maintained as part of each NetWorker host’s Peer resource database.

Note: To complete the steps in this section, the Console server user must have permission to update authentication resources. [“Access privileges for authentication configuration” on page 613](#) provides details.

Deleting NetWorker local host Peer resources

When the credentials for a NetWorker host change, so does its certificate. All hosts that have previously authenticated with the changed NetWorker host must be configured to accept the changed certificate.

Otherwise, authentication will fail because the changed certificate is no longer recognized.

For example, host A authenticates with host B. Host B will have a Peer resource for host A, which contains certificate information. Now the credentials for host A change. Host B will no longer recognize host A and authentication will fail. To solve this problem, delete host B's Peer resource for host A. The next time host A attempts to authenticate with host B, a new Peer resource will be created on host B and authentication will succeed.

To delete a local host Peer resource:

1. From the **Administration** window, click **Configuration**.
2. Open the **Local Hosts** folder.
3. From the **Local Hosts** list, click the NetWorker host whose Peer resource must be deleted.
4. Right-click the Peer resource that corresponds to the NetWorker host whose credentials were changed, then select **Delete**.

Creating a local host Peer resource certificate

To update certificate information for a local host Peer resource, delete the existing Peer resource or update the Peer resource with a certificate from a file. Before you load the certificate from a file, create the certificate as described in this section.

To create a Peer local host certificate:

1. Export the local host credentials file for the NetWorker host whose credentials have changed, if this has not already been done. [“Exporting local host credentials” on page 618](#) provides details.
2. With a text editor that is compatible with UNIX text files, open the credential file and copy and save the certificate information to a new file. Ensure that the new file is saved in a UNIX text-file format.

In this sample credential file, copy the information between the *begin certificate* and *end certificate* comments:

```
certificate:\
"-----BEGIN CERTIFICATE-----
MIIB9...
-----END CERTIFICATE-----";
my hostname: pluto.company.com;
```

```

name: pluto.company.com;
NW instance ID: d9b61002-0004-fe5c3a37-42d5a842-00010000-8945657f;
private key:\
“-----BEGIN RSA PRIVATE KEY-----
MIICW...
-----END RSA PRIVATE KEY-----”;
type: NW instance Information;
resource identifier: 1.0.21.1.139.45.65.23.121.56.111.101(1)

```

This file can now be loaded into all of the Peer resources that exist for the NetWorker host whose credentials have changed. [“Securing the initial peer host authentication process” on page 616](#) provides details.

Creating a custom certificate and private key for a host

NetWorker automatically creates certificate and private keys for each NetWorker host. However, it is possible to create certificate and private key information for a host manually. One may want to do this in special cases such as when a company mandated policy stipulates that certificate and private key information must be generated on a single host that has a trusted random number generation utility. The certificate and key information can then be transferred from the trusted host to other hosts within the enterprise.

To create custom certificate and private key information for a host:

1. On the host that is being used to create the custom certificate and private file, type this command:

```
nwinstcreate -ix
```

Complete the remaining screen prompts as appropriate. The *EMC NetWorker Command Reference Guide* or the UNIX man page provide more information about **nwinstcreate**.

2. On the host for which the custom certificate and private key file was created, place the file in the following directory:

```
NetWorker_install_path\nsrres
```

- a. On Windows hosts, you must give read, write, and modify privileges for the custom certificate and private key file to the Windows Local System Account (SYSTEM).

On Windows, NetWorker services such as **nsrexecd**, are run with SYSTEM privileges. By default, NetWorker services do not have adequate privileges to read and write the custom certificate and private key file.

- b. Ensure that the NetWorker client service, **nsrexecd**, is started on the host. [“Stopping and starting a NetWorker server, client, or storage node” on page 55](#) provides information about verifying and starting NetWorker services.

- c. Start the **nsradmin** program:

```
nsradmin -p nsrexec
```

The **nsradmin** prompt appears.

d. Type this command:

```
. type: nsrla
```

e. Import the custom certificate and private key file:

```
update nw instance info operations: import;
```

```
nw instance info file:
```

```
NetWorker_install_path\nsrrescertificate_file
```

f. Enter **Yes** when prompted to confirm the update.

The custom certificate and private key information will be used for the host.

Multiple clients from the same computer

The NetWorker server identifies each of its clients by the client computer name.

To provide optimal flexibility, the server lets you define multiple Client resources for the same computer, with the same computer name, provided that one of these is true:

- ◆ Each client save set is unique.
- ◆ Clients are included in different backup groups.
- ◆ Clients are associated with different schedules.
- ◆ Clients are associated with different browse and/or retention policies.

For example, looking at a list of configured NetWorker clients, you might see several instances of a client named *mars*. But each instance would contain a unique collection of save sets or would be configured differently regarding groups, schedules, or policies.

Defining multiple clients from the same computer or file system can be useful for backing up specialized files, such as databases. You can use the Comment attribute in the Client resource to help distinguish between multiple Client resources with the same name.

Redefining a file system into multiple client and save set instances

If a client has a large volume of data, you can schedule the client computer for several, separate client and save set backups. By redefining a large file system into multiple client and save set instances, you are able to:

- ◆ Automatically back up a large client file system.
- ◆ Balance the load by avoiding a full backup of the entire file system at one time.

You can then associate each client and save set instance with a different backup group and/or a different schedule. Associating different client instances with different backup groups varies the start time of the backups. Staggering the start times in this way may achieve the load balancing needed.

If different backup start times do not reduce the load adequately, you can associate the different client and save set instances with different backup schedules. Recall that a client's schedule determines the level of backup (if any) that is run on a particular day. By

using different schedules, you can specify that each client and save set instance run its full backup on a different day of the week. [“Schedules” on page 260](#) provide more information on schedules.

Defining a client and save set combination

The same save set can appear in the Save Set attribute of the Client resource for multiple client instances. This characteristic permits you to associate the same save set with more than one group or schedule for backup.

The save sets associated with a specific client instance are visible as a scrollable list in the Save Set attribute of the Create Client and Properties dialog boxes.

If the default value All appears in the Save Set attribute, all local data for the client computer is backed up according to the group and schedule in the Create Client and Properties dialog boxes.

To define a client and save set combination:

1. Create a new or edit an existing NetWorker client.
2. Click the **General** tab of the **Properties** dialog box for the client.
3. In the **Save Set** attribute, delete the default value **All**.
4. Complete one of these steps as appropriate:
 - To configure the client so that a specific file system is backed up, enter the file system pathname in the **Save Set** attribute.
 - To define multiple save sets on a client, enter each save set (partition, file system, or file) on a separate line in the **Save Set** attribute.

Scheduled backups of non-ASCII files or directories

When the Saveset attribute for a Client resource contains non-ASCII characters, the Save Operation attribute must set as follows:

- ◆ UNIX/Linux:
 - For NetWorker clients at release 7.4 or later, the value of the Save Operation attribute should be set to:


```
I18N:mode=nativepath
```
 - For clients at a release level prior to release 7.4, the value of the Save Operation attribute should be set to:


```
I18N:mode=utf8path
```
- ◆ For all Microsoft Windows clients, the Save Operation attribute should be set to:


```
I18N:mode=utf8path
```

If the Client Backup Configuration Wizard is used, it is not necessary to change the Save Operation attribute in the Client resource created by the wizard.

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 (the right to request a program to be executed on another client). [Table 102 on page 624](#) provides a list of tasks that require an update to the servers file.

Table 102 When to modify the servers file

Operations	Update required on the NetWorker client's servers file
Archive request	Addition of the NetWorker server (long and shortname)
Scheduled backup	Addition of the NetWorker server (long and shortname) Addition of the virtual NetWorker and all physical nodes if the NetWorker servers is clustered (long and shortname for all computers)
Remote Directed Restore	The administering client (long and shortname) is added to the destination client's servers file.
NDMP DSA backups	Addition of the NetWorker client (long and shortname) initiating the backup

- ◆ Consider the following:
 - If the servers file is empty, then any NetWorker host has tasking rights. This is a potential security concern.
 - Entries in the servers file restrict client-tasking access to only the machine names listed in the file.
 - The option to modify the servers file during the installation process is only available for Windows and Solaris. For all other operating systems, the servers file must be manually edited with a text editor after the installation has completed.
 - On UNIX computers, client-tasking rights can also be defined by starting the `nsrexecd` daemon with `-s servename`. Starting `nsrexecd` with the `-s` option supersedes the use of the servers files to restrict client-tasking rights.

Editing the servers file

To give clients client-tasking rights to a NetWorker client:

1. Stop the nsrexecd process on the client machine:
 - On Windows, stop the NetWorker Remote Exec service
 - On Unix run the command **nsr_shutdown**.
2. If necessary, remove the **-s** option from the **nsrexecd** command that is invoked by the boot-time startup file. [Table 103 on page 625](#) provides details on the location of the boot-time startup files for each Unix operating system.

Table 103 Boot-time startup file locations

Operating system	Boot-time startup file
Solaris	/etc/init.d/networker
Linux	/etc/init.d/networker
AIX	/etc/rc.nsr
HP-UX	/sbin/init.d/networker

If the **-s** option exists in the boot-time startup file, remove all occurrences of this in the startup file:

```
-s server_name
```

3. Open the servers file in a text editor.
The default installation location for this file is:
 - /nsr/res/servers (UNIX)
 - *NetWorker_install_path*\res\servers (Windows)
4. Enter one hostname per line.
5. Save the changes and exit the text editor.
6. Start the NetWorker Remote Exec server (Windows) or nsrexecd process (UNIX). [Table 9 on page 56](#) provides details on how to start the NetWorker daemons on the various UNIX operating systems.

Client priority

The Priority attribute in the Client resource specifies the order in which clients are contacted for backup. The attribute can contain a value between 1 and 1,000. The lower the value, the higher the priority.

The client with the lowest value for the Priority attribute is placed at the top of the list to be contacted by the NetWorker server. If a value is not specified in the Priority attribute, 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, including these scenarios:

- ◆ 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, the client is put at the end of the backup list.

To increase the number of times each client in a group is retried before the backup attempt is considered unsuccessful, change the value in the Client Retries attribute in the **Group** resource. [“Task 2: Set up a group for backup clients” on page 61](#) provides more information.

Note: The only way to guarantee that ClientA backs up before ClientB is to assign ClientA to a scheduled backup group that starts earlier than the group containing ClientB.

Dedicated client/server interface for backup and recover operations

If a NetWorker client must access a unique network interface on the server that is used for backup and recover operations, enter the value in the NetWorker client resource’s Server network interface attribute. The value you enter should correspond to the unique hostname of the network interface on the server.

The Server network interface attribute is located in the Globals (1 of 2) tab in the client resource. This attribute does not specify the NetWorker server hostname that is used for other NetWorker processes, such as monitoring or authentication operations. All other processes use the hostname of the NetWorker server.

The server network interface attribute is used only for the NetWorker Server. The hostname of the LAN interface which is dedicated for backup on the NetWorker server must be entered here. In order to set a dedicated LAN interface for backup on storage node, use the hostname of the LAN interface which is dedicated for backup as storage node attribute and enter the hostname of the LAN interface which is dedicated for backup on NetWorker server.

For example,

- ◆ To backup a client via NetWorker Server:
 - Storage Node Attribute: **nrsserverhost**
 - Server Network Interface Attribute: A hostname of the LAN interface dedicated for backup on the NetWorker Server.
- ◆ To backup a client via the storage node:
 - Storage Node Attribute: A hostname of the LAN interface dedicated for backup on storage node.
 - Server Network Interface Attribute: A hostname of the LAN interface dedicated for backup on NetWorker server.

CHAPTER 20

Block Based Backup and Recovery

This chapter includes the following sections:

- ◆ Overview of NetWorker block based backup and recovery..... 628
- ◆ Preparing for block based backups..... 630
- ◆ Performing block based backups..... 633
- ◆ Verifying block based backups 636
- ◆ Cloning block based backups..... 636
- ◆ Preparing for block based recovery..... 636
- ◆ Performing block based recovery..... 637
- ◆ Troubleshooting block based backup and recovery issues 642

Overview of NetWorker block based backup and recovery

NetWorker block based backups are high-performance backups that support NTFS and ReFS file systems.

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:

- ◆ Use the Volume Shadow Copy Service (VSS) snapshot capability to create a consistent copy of the source volume for backup.
- ◆ Use the Virtual Hard Disk (VHDx), which is sparse, to back up data to the target device.

Block based backups support only the following devices as target devices, where these devices are Client Direct enabled:

- ◆ Advanced File Type Devices (AFTDs)
- ◆ Data Domain devices

Block based incremental backups use the Change Block Tracking (CBT) driver to identify the changed blocks, and back up only the blocks that are changed.

NOTICE

You must install the latest recommended service packs and VSS patches.

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. A block based backup enables you to mount the backup by using the same file system that you used to back up the data. For example, if the data that you backed up is NTFS, you can mount the block based backup by using NTFS.

A block based backup provides the following capabilities:

- ◆ Mounting of the backup as a file system
- ◆ Maximum of 38 incremental backups after a full backup
- ◆ Mounting of an incremental backup
- ◆ Sparse backup support
- ◆ Backups to disk-like devices
- ◆ Backups to Data Domain
- ◆ Backups of operating system-deduped file systems as source volumes
- ◆ Virtual full backups to Data Domain
- ◆ Synthetic full backups to AFTD
- ◆ Incremental synthetic full backups to AFTD
- ◆ Backups of volumes up to 63 TB each
- ◆ Recoveries from Data Domain without using CIFS share

[Table 104 on page 629](#) lists the backup scenarios and the recovery scenarios that block based backups support.

Table 104 Supported backup and recovery scenarios

Backup scenarios	Recovery scenarios
<ul style="list-style-type: none"> • AFTD backups • Backups to Data Domain by using DD Boost • Full backups • Virtual full backups • Synthetic full backups • Incremental backups • Incremental synthetic full backups • Full backups and incremental backups intermixed with built-in provisions to anchor the incremental backups with an appropriate backup type 	<ul style="list-style-type: none"> • 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 for block based backups and recoveries

Block based backups support the backup and recovery of the following operating systems and configurations:

- ◆ Operating systems on x64:
 - Windows client 8.1
 - Windows client 8
 - Windows Server 2012 R2
 - Windows Server 2012
 - Windows Server 2008 R2
- ◆ Operating systems on x86:
 - Windows client 8.1
 - Windows client 8
- ◆ 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
- ◆ Full and incremental backups of New Technology File System (NTFS) and Resilient File System (ReFS)

Limitations of block based backups

Block based backup does not support the following capabilities and configurations:

- ◆ Non-Windows storage nodes
- ◆ FAT file system
- ◆ Backup levels 1 through 9
- ◆ Synthetic full backups of mixed mode save sets
- ◆ Cloning of incremental backups
- ◆ Granular save sets at either the folder level or the file level, for example, D:\data
- ◆ Checkpoint restart
- ◆ Staging
- ◆ Standard NetWorker directives
- ◆ The **scanner** command with the **-i** option for rebuilding indexes for block based backups
- ◆ The **nsrclone** command with the **-m** option for migrating block based backup save sets to other volumes
- ◆ Image recovery to a system volume.

Preparing for block based backups

You must complete the following tasks before you perform a block based backup and recovery:

1. [“Creating a backup device” on page 630](#)
2. [“Configuring block based backups” on page 632](#)
3. [“\[Optional\] Creating a CIFS share for block based recoveries” on page 632](#)

Creating a backup device

Depending on the backup requirements, you can create the following types of devices:

- ◆ AFTD
- ◆ Data Domain CIFS
- ◆ DD Boost

Creating an AFTD or a Data Domain CIFS device

You can create an AFTD by using one of the following methods:

- ◆ [“Create an AFTD by using the Device Wizard” on page 168](#) describes how to create an AFTD by using the NMC wizard.
- ◆ [“Create an AFTD by using the Properties window \(Windows\)” on page 173](#) describes how to create an AFTD by using the New Device Properties dialog box.

NOTICE

To make a local AFTD Client Direct enabled, you must specify the Uniform Naming Convention (UNC) path in the **Device access information** field of the **Create device properties** dialog box.

Creating a DD Boost device

Complete the following steps to create a DD Boost device:

1. In NMC, click **Devices**.
2. In the left pane, right-click **Devices** and click **New Device Wizard**.
The **Select the Device** page appears.
3. Select **Data Domain**, and click **Next**.
The **Data Domain Preconfiguration Checklist** page appears.
4. Click **Next**.
5. Complete the following steps 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, specify the IP address of the Data Domain system.
 - b. In the **Data Domain DDBoost Username** field, specify the username of the Data Domain user.
 - c. In the **Data Domain DDBoost Password** field, specify the password of the Data Domain user.
 - d. Specify the required values in the other fields.
 - e. Click **Next**.
6. Complete the following steps 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. Complete the following steps 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 specify 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**.

The **Select Storage Nodes and Fibre Channel Options** page appears.

8. Select the storage node, specify the required values in the other fields, and click **Next**.

The **Select SNMP Monitoring Options** page appears.

9. Specify the required field values, and click **Next**.

The **Review the Device Configuration Settings** page appears.

10. Review the configuration settings, and click **Configure**.

The **Device Configuration Results** page appears.

11. Click **Finish**.

Configuring block based backups

You can enable the block based backup feature as part of the client configuration. You must select the following NetWorker client configuration attributes:

- ◆ Client direct (selected by default)
- ◆ Block based backup

You can select these attributes by using one of the following methods:

- ◆ NetWorker Client Configuration wizard
- ◆ Client Properties window
- ◆ The nsradmin program

[Optional] Creating a CIFS share for block based recoveries

To recover data from either a Data Domain device or an AFTD, you must enable a CIFS share to access save sets on the device.

Creating a CIFS share on a Data Domain device

Complete the following steps to create a CIFS share on a Data Domain device and access the share from a client:

1. Open Data Domain Enterprise Manager.
2. In the left pane, under **DD Network**, select the Data Domain device.
3. Click **Data Management**.
4. Click **CIFS**.
5. Click **Shares**.
6. Click **Create**.
7. Complete the following steps in the **Create Share** dialog box:
 - a. In the **Share Name** field, specify the NetWorker server short name.
 - b. In the **Directory Path** field, specify the full pathname of the share folder that is available on the Data Domain device, in the following format:

/data/col1/NetWorker server short name

- c. In the **Clients** field, add the IP addresses or the hostnames of the clients from which you want to access the CIFS share.
 - d. Click **OK**.
8. Access the CIFS share from the clients.
 9. If you want to add more clients after creating a share, click **Modify**.

Creating an AFTD CIFS share

Complete the following steps to create a CIFS share on an AFTD on Windows:

1. Right-click the folder that you want to share, and click **Share with > Specific people....**
The **File Sharing** dialog box appears.
2. Select or add the people with whom you want to share the folder, and click **Share**.
The folder becomes shared.
The access credentials are the same as the administrator's credentials on the host.

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 57](#) provides more information about backing up data by using NetWorker.

You can perform a block based backup as any of the following types of backup:

- ◆ [“Scheduled backups” on page 633](#)
- ◆ [“Incremental backups” on page 634](#)
- ◆ [“Virtual full backups” on page 634](#)
- ◆ [“Synthetic full and incremental synthetic full backups” on page 634](#)
- ◆ [“Manual backups or client-initiated backups” on page 635](#)
- ◆ [“Save set backups” on page 635](#)
- ◆ [“Exclude list backups” on page 635](#)
- ◆ [“Windows Server 2012 R2 and 2012 deduplication volume backups” on page 635](#)
- ◆ [“CSV backups” on page 635](#)
- ◆ [“Windows BMR backups” on page 636](#)

Scheduled backups

The NetWorker **savegrp** program 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

The NetWorker **savegrp** program initiates a scheduled backup as either a full backup or an incremental backup.

You must perform the incremental backup of a volume only to the same device to which the full backup of the volume was performed.

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.
- ◆ You already performed 38 incremental backups.

NOTICE

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

The virtual full backups are applicable only to the Data Domain devices. When you perform an incremental backup to a Data Domain device, the backup is performed 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: Selecting any backup level apart from **full** results in performing a virtual full backup. Block Based Backup enables you to perform 38 virtual full backups only, after which the subsequent backup shifts to the **full** backup.

Synthetic full and incremental synthetic full backups

The synthetic full and incremental synthetic full backups are applicable only to AFTDs. A synthetic full backup consolidates data from all the existing full and incremental backups. An incremental synthetic full backup first performs an incremental backup from the immediate previous backup and then performs a synthetic full backup.

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 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 are defined only at the volume level.

Save set backups

You can use a block based backup to back up the following save sets:

- ◆ 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)

Exclude list backups

The exclude list backups exclude the files that the operating system specifies as unwanted files in the

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup Registry key. Also, the backup excludes unsupported writers and application files such as SQL, Exchange, and so on. However, the backup does not exclude the files present in the system volume information.

Windows Server 2012 R2 and 2012 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, and blocks are merged out of these 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 (CSV) across all nodes. Block based backups support only full backups of CSV volumes, 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.

For non-CSV volumes, if a failover happens after an incremental backup, the backup shifts to a full backup.

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 by using the NetWorker Client Configuration wizard, the Client Properties window, or the **nsradmin** program.

[“Windows Bare Metal Recovery” on page 719](#) provides more information about BMR backups.

Verifying block based backups

Perform the following steps to verify block based backups:

1. Run the following command:

```
mminfo -S -q ssid=<Save_set_ID>
```

2. From the output, see whether the following information is present:

```
*BlockBasedBackup: Yes;
```

For virtual full backups and synthetic full backups, see whether the following additional information is present respectively:

- *BlockBased Virtual Full: Yes;
- *Synthetic full: Yes

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 backups only.

[“Cloning” on page 339](#) provides more information about cloning the NetWorker backups.

Run the following command to clone the block based backup save sets:

```
nsrclone -b target_pool_name -S save_set_ID/clone_ID
```

Preparing for block based recovery

You need to be familiar with the recovery operations, workflows, and interfaces associated with block based backup recovery. You can perform a block based backup recovery by using either NMC or the NetWorker command-line interface (CLI).

You typically complete the following tasks to perform a recovery by using NMC:

1. Select the save set.
2. Perform either file level recovery or image/destructive recovery.

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 program mounts all the save sets on a device that supports the Client Direct functionality.

To recover data from either an AFTD or a Data Domain device by using the CIFS share, you must enable the CIFS share to access save sets on the device.

[“Preparing for block based backups” on page 630](#) provides information about how to create the CIFS share and devices.

Performing block based recovery

You can perform a block based recovery by using either NMC or the NetWorker CLI.

Using NMC for block based recovery

Complete the following steps to perform a block based recovery by using NMC:

1. Open NMC.
2. Click **Recover**.
3. From the menu bar, select **Recover > Recover > New Recover**.
4. Complete the following steps on the **Select the Client to Recover** page:
 - a. Under **Source client**, in the **Name** field, specify the name of the client on which the backed-up data exists.
 - b. Under **Destination client**, specify the client to which you want to recover the backed-up data.
 - c. For the type of backup that you want to recover, select **Block Based Backup**.
 - d. Click **Next**.
5. Complete the following steps 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. Depending on the type of recovery you selected in [step a](#), perform one of the following tasks:
 - For a file level recovery, select the save set to recover from the left pane and select the files to recover from the right pane.
 - For an image level recovery, select the save set that you want to recover from the left pane.
 - d. Click **Next**.

The **Select the Recovery Options** page appears.

6. Depending on the type of recovery you selected in [step a](#) of [step 5](#), perform one of the following tasks:
 - 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**.

The **Obtain the Volume Information** page appears.

7. Click **Next**.
8. Complete the following steps on the **Perform the Recovery** page:
 - a. Under **Identity**, specify a name for the recovery in the **Recover name** field.
 - 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 your choice.
 - c. If you want to stop the recovery at a certain time, specify the time in the **Specify a hard stop time** field.
 - d. Select the **Recover Resource persistence** option according to your choice.
 - e. Click **Run Recovery**.
 - f. The **Check the Recovery Results** page appears.

The recovery log appears when the recovery progresses.

After the recovery succeeds, a successful completion message appears at the bottom of the recovery log.

9. To export the log file, click **Export Log File**.
10. Click **Finish**.

Using the CLI for block based recovery

You can perform a block based recovery by using the **recover.exe** command, which is applicable only to local clients. However, you cannot perform a remote or redirected recovery by using this command.

Commands for performing file level recovery

- ◆ The following command mounts the backup and opens the command prompt at the mount point:

```
recover.exe -S ssid
```

You can 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.

- ◆ The following command mounts the backup and copies specific files from the input file to the destination:

```
recover.exe -S ssid -I input file -d destination
```

NOTICE

To avoid using the Windows version of the **recover.exe** command on Windows operating systems, perform one of the following tasks:

- ◆ Include *NetWorker_install_path\bin\recover.exe* at the command prompt.
- ◆ Ensure that the \$PATH environment variable lists *NetWorker_install_path\bin* before %SystemRoot%\System32.

Command for performing image and destructive recoveries

The following command mounts the source volume, unmounts the target volume, and copies the used blocks on the source volume to the target volume:

```
recover.exe -S ssid -r target_volume
```

Note: Ensure that the target volume is not a system volume.

You must ensure that the following requirements pertaining to size are met for performing the recovery:

- ◆ The size of the target volume is either the same or bigger than the size of the source volume.
- ◆ The cluster size of the source volume is the same or bigger than the cluster size of the target volume.

Command-line options for recover.exe

[Table 105 on page 639](#) describes the key options that you can use with the **recover.exe** command for a block based recovery.

Table 105 Key options for the block based recover.exe command

Option	Description
-r [volume GUID/mount point]	Specifies the supported destinations for save set recovery: <ul style="list-style-type: none"> • Volume name • Raw pathname • Volume GUID • Existing mount point
-k [mount/unmount]	Specifies to mount or unmount a save set.
-S [save set ID/clone ID]	Specifies the save set ID or the clone ID that you want to recover.
-I [input file]	Specifies a file that contains a list of files that you want to recover. This is useful for performing disaster recovery and remote recovery.

Performing Windows BMR recovery

The procedure to recover a block based backup through a Windows BMR recovery is the same as the procedure to perform a NetWorker Windows BMR recovery. 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 clone recovery

Complete the following tasks to perform a clone recovery:

- ◆ [“Recovering data from Client Direct enabled devices” on page 640](#)
- ◆ [“Recovering data from Client Direct disabled devices” on page 640](#)

Recovering data from Client Direct enabled devices

Client Direct enabled devices include AFTD, DD Boost, and Data Domain CIFS devices.

You can recover the data by using either of the following methods:

- ◆ Use NMC by performing [step 1](#) through [step 10](#) in [“Using NMC for block based recovery” on page 637](#).
- ◆ Use the NetWorker CLI by running either of the following commands:

```
recover.exe -S save_set_ID/clone_ID for file level recovery
recover.exe -S save_set_ID/clone_ID -r target_volume for image
recovery
```

Recovering data from Client Direct disabled devices

Client Direct disabled devices typically include tape devices. You must have a Client Direct enabled device to perform a recovery. When you perform a tape recovery, the 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 even before the retention period lapses.

You can recover the data by using either NMC or the NetWorker command **recover.exe**

Using NMC for the clone recovery

You must complete the following steps to use NMC to recover the data from Client Direct disabled devices.

1. Open NMC.
2. Click **Recover**.
3. From the menu bar, select **Recover > Recover > New Recover**.
4. Complete the following steps on the **Select the Client to Recover** page:
 - a. Under **Source client**, in the **Name** field, specify 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. Complete the following steps 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 **Subsets of save sets** and appropriate save sets that belong to the selected save set group.
 - d. Under **Pool**, select the pool that has the Client Direct enabled device to which you want to copy the cloned data.
 - e. Click **Next**.
The **Copying the Backup to Disk** page appears.
 - f. After the cloning succeeds, click **Next**.
6. Complete the following steps 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. Depending on the type of recovery you selected in [step a](#), perform one of the following tasks:
 - For a file level recovery, select the save set to recover from the left pane and select the files to recover from the right pane.
 - For an image level recovery, select the save set that you want to recover from the left pane.
 - d. Click **Next**.
The **Select the Recovery Options** page appears.
7. Depending on the type of recovery you selected in [step a](#) of [step 6](#), perform one of the following tasks:
 - 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**.
The **Obtain the Volume Information** page appears.
8. Click **Next**.
9. Complete the following steps on the **Perform the Recovery** page:
 - a. Under **Identity**, specify a name for the recovery in the **Recover name** field.
 - 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 your choice.

- c. If you want to stop the recovery at a certain time, specify the time in the **Specify a hard stop time** field.
- d. Select the **Recover Resource persistence** option according to your choice.
- e. Click **Run Recovery**.

The **Check the Recovery Results** page appears.

The recovery log appears when the recovery progresses.

After the recovery succeeds, a successful completion message appears at the bottom of the recovery log.

10. To export the log file, click **Export Log File**.

11. Click **Finish**.

Using the CLI for the clone recovery

Run either of the following commands to recover the data from Client Direct disabled devices:

```
recover -S save_set_ID/clone_ID -l pool_name for file level recovery
recover -S save_set_ID/clone_ID -l pool_name -r target_volume for
image recovery
```

The pool that you select must have a Client Direct enabled device. The pool must also be a backup clone pool.

Troubleshooting block based backup and recovery issues

The following table lists the issues that you might encounter during block based backups and recoveries and the troubleshooting techniques to resolve the issues.

Table 106 Troubleshooting block based backup and recovery issues

Error Message	Resolution
Block based backups are only supported with Client Direct.	Enable the Client Direct option in the Client Properties window.
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 Inputfile.	Perform an image recovery if applicable. Otherwise, select all the files except the system files such as System Volume Information and Recycle Bin for 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 DFA-enabled.	Enable the Client Direct option in the Client Properties window.

CHAPTER 21

NetWorker support for NDMP

This chapter covers these topics:

◆ Overview of NDMP	644
◆ Components in a NetWorker NDMP environment	644
◆ Configurations in a NetWorker NDMP environment	645
◆ Pre-configuration requirements for NDMP data operations	650
◆ Configuring Devices for NDMP operations.....	654
◆ Creating resources to support NDMP clients	662
◆ Creating and configuring the NDMP client resource	663
◆ Performing NDMP backups.....	671
◆ Troubleshooting NDMP configuration and backup failures	674
◆ Cloning NDMP save sets.....	678
◆ Reporting NDMP Data.....	679
◆ Performing NDMP recoveries	680
◆ Troubleshooting NDMP recover	691

Overview of NDMP

The “[network data management protocol \(NDMP\)](#)” is a TCP/IP-based protocol that specifies how network components talk to each other for the purpose of moving data across the network for backup and recovery.

The NDMP protocol addresses the problems associated with backing up data in heterogeneous environments when you use different operating system vendors, backup developers, and Network Attached Storage (NAS) devices.

NDMP enables disparate vendors to use a common NDMP protocol for the backup architecture. With the NetWorker NDMP interface, you can connect to hosts that have an active “[NDMP service](#)” or an NDMP data module installed. The NDMP host does not have the NetWorker software installed. NDMP allows a NAS device to back up data to other NDMP-controlled tape or disk devices that are on the network. NDMP passes control of the data and the file metadata to and from the NetWorker software.

The NetWorker server by default attempts to establish communications with a NAS filer by using NDMP version 4. If the NAS does not support or use NDMP v4, communications are autonegotiated to use the highest version that the NAS filer supports. NetWorker supports NDMP v3 and later however, there are some NetWorker features that require a specific version of NDMP on the NAS.

The *NetWorker Hardware Compatibility Guide* on the EMC Online Support web site provides a list of NAS filers that the NetWorker software supports.

Components in a NetWorker NDMP environment

Three main components support NDMP data operations with the NetWorker software:

- ◆ NDMP Data Server—The system that contains the NDMP data. The NDMP Data Server is also known as the datamover, the NDMP client, or the NAS filer. The NAS transfers the data to the primary storage devices through a data connection. You configure the NAS as a client of the NetWorker server; however, you do not install the NetWorker client software on the NAS.
- ◆ NDMP Tape Server—The host with the backup device to which NetWorker writes the NDMP data.
- ◆ Data Management Agent(DMA)—The NetWorker server is the DMA.

The DMA:

- Initiates the NDMP backup.
- Monitors the NDMP backup and recover operations.
- Maintains the media database and the client file index entries for NDMP backups.
- Maintains the resource database information for the NAS and NDMP Tape Server.

Configurations in a NetWorker NDMP environment

You can use three methods to configure the NDMP Data Server and the NDMP Tape Server to perform backups and recoveries. You can customize the NetWorker environment to support NDMP data operations in each scenario:

- ◆ “[NDMP local backup](#)” on page 645—The NDMP Data Server and the NDMP Tape Server reside on the same physical host.
- ◆ “[NDMP backups to non-NDMP devices \(NDMP-DSA\)](#)” on page 646—An NDMP Data Server Agent (NDMP-DSA) sends the NDMP data from the NDMP Data Server to the non-NDMP devices.
- ◆ “[Three-party backup with NDMP devices](#)” on page 649—The NDMP Data Server and the NDMP Tape Server reside on different physical hosts. The NAS passes NDMP metadata to the NetWorker server. The NetWorker server sends the data to an NDMP tape device attached to another NDMP Server.

[Table 107 on page 645](#) summarizes the differences between the NDMP Tape Server and NDMP-DSA.

Table 107 Distinctions between NDMP Tape Server and NDMP-DSA

NDMP Tape Server	NDMP-DSA
Supports only the NDMP type of tape device.	Supports any type of device that the NetWorker software supports.
Does not support backup to disk.	Supports backup to disk.
Does not support multiplexing.	Supports multiplexing.

NDMP local backup

In an NDMP local backup (Direct-NDMP), the NDMP Data Server (NAS) sends data to a locally attached tape device or library. The `nsrndmp_save` program runs on the NetWorker server, and only the metadata and the NDMP control information traverse the network between the NetWorker server and the NDMP host.

Advantages:

The NDMP data does not traverse the network, this prevents network congestion.

Disadvantages:

- ◆ The `nsrndmp_save` program queries the NDMP Tape Server at consistent intervals to determine the status of a backup. These queries have an impact on backup performance.
- ◆ The NetWorker software does not multiplex NDMP save sets and writes NDMP data serially to the local device. As a result, backups are faster, but recoveries are slower.
- ◆ NDMP local backups are not suitable when there are many large file systems to back up on a NAS filer.
- ◆ You cannot archive NDMP save sets.

The NetWorker server, or data management application (DMA), performs these tasks:

- ◆ Initiates the backup or the recovery request through the NDMP connection.
- ◆ Receives the file history information from the data server.

During a backup, the NAS filer is the NDMP Tape Server and the NDMP Data Server. The NAS filer performs these tasks:

- ◆ Receives the backup requests.
- ◆ Reads the backup data on the NAS disks.
- ◆ Produces a data stream for the backup.
- ◆ Writes the data stream to the tape or file device.

Figure 45 on page 646 illustrates a local backup configuration.

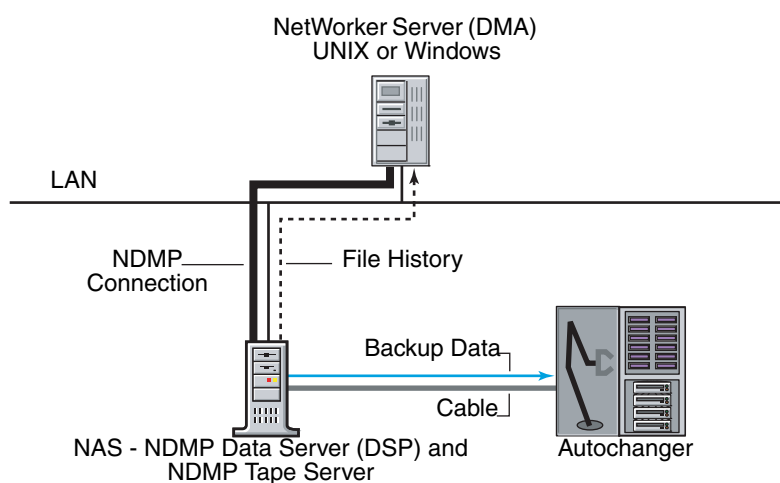


Figure 45 NDMP local backup configuration

NDMP backups to non-NDMP devices (NDMP-DSA)

In this scenario, the NetWorker software writes NDMP data to non-NDMP devices, including tape, disk, optical, and Data Domain devices. Use NDMP-DSA backups when there are many small file systems to backup, and network throughput is not a concern. An NDMP data backup to disk is faster than an NDMP backup to tape. Directing NDMP staged and cloned data to a non-NDMP device is faster than sending the data to an NDMP device.

The NetWorker software uses the NDMP Data Server Agent (DSA) and the `nsrndmp_save` command to send NDMP data to a non-NDMP device. The process associated with DSA is `nsrdsa_save`.

The benefits of using NDMP-DSA include the ability to:

- ◆ Write NDMP data to devices that also contain non-NDMP data.
- ◆ Multiplex NDMP save sets to improve backup speeds. Recovery speeds are slower
- ◆ Stage save sets from the disk to tape.
- ◆ Archive NDMP save sets.

You can back up NDMP data to a non-NDMP device in one of two ways:

- ◆ “Sending NDMP data to non-NDMP devices that are local to the NetWorker server” on page 647
- ◆ “Sending NDMP data to non-NDMP devices that reside on a NetWorker storage node” on page 647

Sending NDMP data to non-NDMP devices that are local to the NetWorker server

When you send NDMP data to non-NDMP devices that are local to the NetWorker server:

- ◆ The backup data traverses the network between the NetWorker server and the NDMP Data Server.
- ◆ The metadata, the NDMP control information, and the FH remain local to the NetWorker server and do not traverse the network.

Figure 46 on page 647 illustrates a NetWorker storage device directly attached to the NetWorker server. The NetWorker server initiates an NDMP-DSA backup. The `nsrnmdd` process on the NetWorker server processes the data and metadata. The `nsrindmp_2fh` and `nsrdmpix` processes on the NetWorker server process the file history (FH) data and then pass the FH data to the `nsrindexd` process.

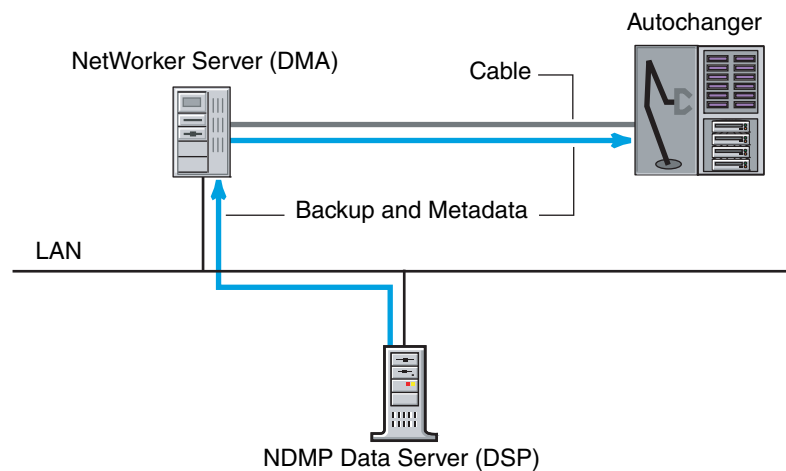


Figure 46 Backup initiated from a NetWorker server with an attached storage device

Sending NDMP data to non-NDMP devices that reside on a NetWorker storage node

You can configure NDMP backups to a NetWorker storage node in one of three ways:

- ◆ “Immediate save” on page 648
- ◆ “Non-immediate save” on page 648
- ◆ “Client Direct file access” on page 649

Immediate save

When you configure an NDMP backup with immediate save:

1. The **nsrdsa_save** backup command runs on the NetWorker storage node.
2. The NetWorker software uses TCP/IP and shared memory to communicate between the **nsrdsa_save** and **nsrmmd** processes.
3. The NetWorker server processes the backup data and sends the data to the non-NDMP device directly through the **nsrmmd** process on the storage node.

When the NetWorker software uses immediate save to send the NDMP data:

- ◆ The **nsrindexd** process on the NetWorker server processes the file history.
- ◆ After the data backup completes and the sessions with the NDMP Data Server and the NetWorker server close, the NetWorker software commits the FH to the client file index associated with the NDMP client.

Figure 47 on page 648 illustrates a NetWorker configuration that uses immediate save.

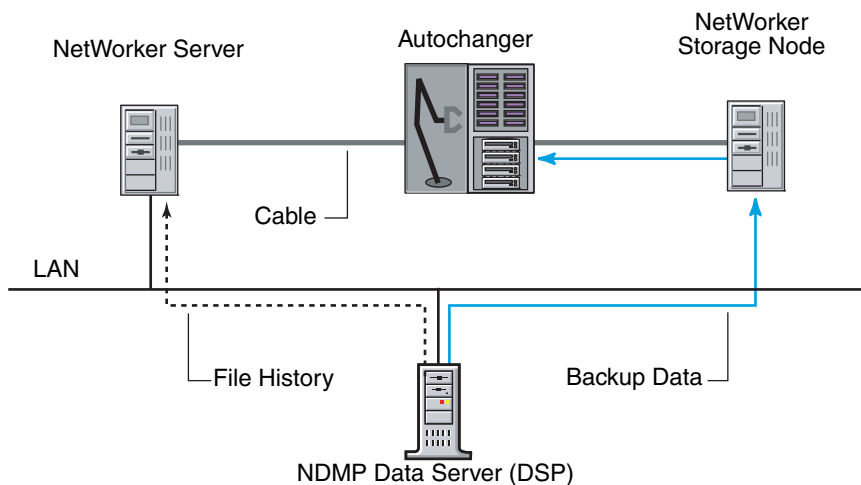


Figure 47 NDMP backup that uses immediate save

Non-immediate save

By default, NDMP backups to a non-NDMP device uses non-immediate save. When you configure an NDMP backup to use non-immediate save:

1. The **nsrdsa_save** backup command runs on the NetWorker server.
2. The **nsrdsa_save** process uses TCP/IP to read the data in a local buffer.
3. The **nsrdsa_save** process transmits the data to the **nsrmmd** process on the storage node.
4. The **nsrmmd** process writes the data to the storage device.

This approach is inefficient and has slow performance for the following reasons:

- ◆ Backup data traverses the network between the NetWorker server, the NDMP host, and the NetWorker storage node.
- ◆ Metadata and the NDMP control information traverse the network between the NetWorker server and the storage node.
- ◆ FH traverses the network between the NetWorker server and the NDMP Data Server.

Client Direct file access

Use Client Direct file access (DFA) technology only when you perform an NDMP backup to a disk such as an advanced file type device (AFTD). DFA writes the data directly to disk, and bypasses the nsrmmmd process on the storage node. The storage node only plays a role in loading the volume. This is a highly efficient, high-performance approach.

Three-party backup with NDMP devices

A three-party backup, or three-way backup, sends NDMP data to an NDMP Tape Server, but the NDMP Data Server and the NDMP Tape Server are not the same physical host.

There are two types of three-party backups:

- ◆ The NetWorker software sends the NDMP data to non-NDMP devices. [“NDMP backups to non-NDMP devices \(NDMP-DSA\)” on page 646](#) provides more information.
- ◆ The NetWorker software sends NDMP data to NDMP devices. In this scenario, the data flows from the NDMP Data Server to the NDMP Tape Server, and then to a library locally attached to the NDMP Tape Server. In this configuration, you cannot archive the NDMP save sets.

In addition to using a NetWorker server or storage node as the NDMP Tape Server, you can use these third party NDMP Tape Servers:

- ◆ NetWorker SnapImage Module 2.5 or later.
- ◆ A DinoStor TapeServer. This hardware connects one or more libraries to the network and allows you to back up any NDMP host to one location instead of requiring a local backup device for each server.

[Figure 48 on page 650](#) demonstrates a three-party configuration, which enables backup and recovery to a NDMP device attached to another NDMP server.

In this example:

- ◆ One server is the data server.
- ◆ The second server is the tape server.
- ◆ The third party is the NetWorker server (DMA).

This configuration is similar to the flow of data between a NetWorker client and a NetWorker server or storage node, except that it is not necessary to install the NetWorker software on either of the NDMP hosts. Data flows from the NDMP Data Server over the network to the NDMP Tape Server and then to tape. The NDMP Data Server sends the metadata to the NetWorker server.




Figure 48 Three-party NDMP backup to NDMP devices

Pre-configuration requirements for NDMP data operations

This section provides requirements to review before you configure the NetWorker software for NDMP data operations.

Note: [“Firewall Support” on page 853](#) describes how to determine port requirements for NDMP backups and recoveries when a firewall exists in the NetWorker datazone.

- ◆ [“NDMP feature requirements” on page 651](#)
- ◆ [“Locale requirements with NDMP” on page 652](#)
- ◆ [“Memory and space requirements for NDMP FH updates” on page 653](#)
- ◆ [“Performance Considerations” on page 653](#)
- ◆ [“NDMP licensing requirements” on page 654](#)

NDMP feature requirements

[Table 108 on page 651](#) provides information to review before implementing the Checkpoint restart, SMTape, iSCSI, vbb, and DAR/DDAR features in NetWorker.

Table 108 NDMP features

Feature	Information (page 1 of 2)
Checkpoint restart	<ul style="list-style-type: none"> • Supports NetApp and Isilon filers only. • Requires NDMP v4 restartable backup extension. • Enabling checkpoint restart support for the NDMP client results in slower backups because NetWorker writes the checkpoint files at defined intervals. The more frequently NetWorker writes checkpoint files, the slower the backup. • You cannot use NetWorker 7.6.x and earlier to restore a checkpoint restarted NDMP backup. <p>Note: “Checkpoint restart backups” on page 95 provides more information about checkpoint restarts.</p>
Snapmirror to tape (SMTape)	<p>Performs block-level backup of SnapMirror volumes on NetApp filers. Reduces the backup window when millions of files reside on the NetApp filer. Use the SMTape feature in instances where NDMP full backups become impractical. SMTape:</p> <ul style="list-style-type: none"> • Copies large NetApp file systems to secondary storage instead of using the standard NDMP full or differential backups • Supports a 240 KB block size. • Allows mirroring of backups to disk and tape devices. • Supports full volume backups and recoveries only. You cannot use SMTape for file indexes or file restores. • Supports save set recoveries, only. • Does not support incremental and differential backup levels. <p>The NAS generates a snapshot of the file system at the beginning of the SMTape operation. Use Environment variables to control the conditional call to retain or delete the snapshot. “Creating and configuring the NDMP client resource” on page 663 describes how to configure an NDMP client by using SMTape.</p>
iSCSI	<p>NetWorker supports iSCSI LUNS on EMC Celerra and NetApp filers. EMC Celerra filers do not support NDMP based backups and recoveries of iSCSI LUNS.</p> <p>NetApp filers support NDMP based backups and recoveries of iSCSI LUNS but you cannot perform an index-based recovery. To perform a full save set recover to another volume, the destination volume must be at least two and a half times as large as the source volume.</p> <p>Notice: NetApp recommends using SnapMirror to safeguard iSCSI LUNS instead of backups.</p>

Table 108 NDMP features

Feature	Information (page 2 of 2)
vbb	<p>vbb supports:</p> <ul style="list-style-type: none"> • EMC DART version 5.5 and later. • Index-based recoveries of a Celerra or VNX block level to the same volume or another location. <p>Use:</p> <ul style="list-style-type: none"> • Checkpoint configuration utility to configure checkpoint file systems on the EMC Celerra or VNX before you perform a backup. • Full Destructive Restore (FDR) to perform a full save set recovery of a raw volume of equal or greater size than the backup. <p>NetWorker performs a file-by-file recovery when you:</p> <ul style="list-style-type: none"> • Recover data from a Celerra or VNX block-level backup. This recover requires disk space in the root directory of the target filesystem to store temporary recovery files. • Perform save set recoveries and NDMP Directory Restores to an existing file system. <p>When you backup a volume that uses native Celerra deduplication, you cannot perform an index-based or NDMP Directory Restore of the backup You can only perform an FDR restore from a level FULL save set.</p> <p>Note: <i>Configuring NDMP backups on Celerra and Using Celerra Data Deduplication</i> on EMC Support Online provides detailed information about how to prepare the filer before you perform FDR.</p>
DAR and DDAR	<p>DAR and DDAR sends file information from the NAS filer to NetWorker. This allows a single file recovery or a directory recovery to position to the exact location of the data on the tape media. NetWorker does not read the file and record numbers sequentially to locate the data.</p> <p>You cannot use DDAR when you enable vbb.</p> <p>Supports EMC DART version 5.5 or later and NetApp with OnTap version 6.4 and later.</p> <p>Requires NDMP version 3 or later. Recoveries on earlier NDMP versions fail.</p>

Locale requirements with NDMP

When running NDMP backups, ensure that you use consistent locale settings in the environment.

- ◆ NetWorker supports the UTF-8 format with CIFS clients. NetWorker only supports NFS clients of a NetApp filer when the NFS clients can generate UTF-8 encoded data. If you set the UTF8=Y variable during an NDMP backup and the backup contains path names with non-ASCII characters, then index-based recoveries of the backup fail with an error message similar to the following:

```
"RESTORE: could not create path pathname"
```

- ◆ If you use **UTF8=Y** variable and perform a backup, then you must recover path names that contain non-ASCII characters by using either a save set recovery from the command line or an NDMP Directory restore.
- ◆ All UNIX locale settings on the NAS filer, including UTF-8 must be the same.

- ◆ Configure the NAS filer to use UTF-8 character sets. Contact the NAS vendor for configuration requirements.
- ◆ Use only UNIX Console clients that have the same locale setting as the NAS filer.
- ◆ You can perform backup and recovery operations on any locale. However, if you try to browse on a locale that is different from the original locale, then the file names appear as random characters.
- ◆ A single save set supports data that belongs to only one code set. If you have data in multiple code sets, you must create multiple save groups. [“Creating resources to support NDMP clients” on page 662](#) describes how to create groups for NDMP clients.
- ◆ A save set can contain filenames that belong to different languages if all characters in those languages belong to the same code set. For example ISO 8859-1 and ISO 8859-15 include most Western European languages, such as French, Spanish, and Portuguese. NetWorker can back up filenames from these languages in a single save set.

Memory and space requirements for NDMP FH updates

During an NDMP backup, the NDMP Data Server sends the FH metadata information to the NetWorker server. The NetWorker software does not verify or modify FH metadata received from the NAS. The NetWorker software uses the file history information to maintain appropriate indexes and media database entries for the NDMP client backups.

The `nsrmdmp_2fh` and `nsrdmpix` binaries interact with the raw database, instead of virtual memory, to process the FH metadata. As a result, memory requirements for this process are minimal. The NetWorker server stores metadata updates in the `\nsr\tmp` directory and commits the metadata to the client file index after the NDMP client backup completes.

Use the following formula to determine the required physical space for the default `\nsr\tmp` directory:

$$2 * (144 + \text{average file name length}) * \text{number of entries in the file system}$$

For example:

For one million file entries with an average file name length of 128, use this formula to compute the required temporary swap space:

$$2 * (144 + 128) * 1,000,000 = 544 \text{ MB approximately}$$

Performance Considerations

Volume loading and positioning operations do not occur during a volume selection process because of an information exchange between the `nsrmmd` process and the `nsrmdmp_save` or `nsrmdmp_recover` command. To avoid the overhead associated with the exchange of information, back up the data to a storage node device.

On NetApp filers with Data OnTap 6.4 and later, NetWorker reads all metadata from tape before recovering the files. For large save sets with 20 million files or more, the recovery time for a file can exceed three hours. This also applies to backups because NetWorker records the metadata for the whole volume onto the tape during a single file backup.

NDMP licensing requirements

NetWorker with NDMP requires an additional license, separate from the NetWorker base product according to a tiered or capacity-based licensing structure. The NetWorker *Licensing Guide* provides more information.

Configuring Devices for NDMP operations

Review this section for information about how to configure the NetWorker environment for NDMP data operations:

- ◆ [“NDMP device limitations” on page 654](#)
- ◆ [“DinoStor-managed jukeboxes” on page 655](#)
- ◆ [“Configuring NDMP on Isilon filer” on page 655](#)
- ◆ [“Determining NDMP device path names” on page 655](#)
- ◆ [“Configuring NDMP devices” on page 657](#)
- ◆ [“Configuring NDMP-DSA devices” on page 661](#)
- ◆ [“Configuring the Clone Storage Node” on page 662](#)

The *NetWorker Hardware Compatibility Guide* on the EMC Online Support web site provides a list of NDMP devices that the NetWorker software supports.

NDMP device limitations

Review these limitations before you configure NDMP devices.

- ◆ The following NetWorker server resource attributes do not apply to NDMP devices but do apply to storage nodes devices:
 - nsrmmmd polling interval
 - nsrmmmd restart interval
 - nsrmmmd control timeout
- ◆ You cannot use the **jbexercise** utility with an NDMP autochanger.
- ◆ You cannot configure NDMP devices on a dedicated storage node.
- ◆ NDMP media device handle must be a non-rewind device handle.
- ◆ You cannot configure Advanced file type devices and file type devices as an NDMP device.
- ◆ You cannot configure an NDMP Autochanger when the NDMP protocol is before version 3. You must determine the NDMP device handles, then use the **jbconfig** command to configure the autochanger. [“Determining NDMP device path names” on page 655](#) provides more information.

DinoStor-managed jukeboxes

The DinoStor software provides a web-based interface for administering and controlling the TapeServer settings.

Review this information before you use a DinoStor-managed jukebox with NetWorker:

- ◆ When you configure the DinoStor TapeServer, set the port number to **10000** on the NDMP page of the **Configure** tab.
- ◆ NetWorker supports:
 - SCSI tape devices.
 - GigE and 10/100 Base-T networks.
- ◆ You cannot use DDS because the DinoStor TapeServer is not fibre-equipped.

Configuring NDMP on Isilon filer

Before you can backup NDMP data, you must configure OneFS NDMP.

1. Use **ssh** to connect to a node in the cluster.
2. Use the **isi** command to create the NDMP username and password:

```
isi ndmp user create username password
```

3. Use the **isi** command to enable ndmp:

```
isi ndmp settings set --name dma --value emc
```

Determining NDMP device path names

To configure an NDMP standalone device or NDMP jukebox, you must first determine the path names of the media devices. If the NAS filer does not support the NDMP_CONFIG interface or uses NDMP version 3, you must also determine the library device handle.

There are two methods to determine the NDMP device path names and the library handle:

- ◆ [“Determining the NDMP device path names by using inquire” on page 655](#)
- ◆ [“Determining the NDMP path names with vendor-specific commands” on page 656](#)

Determining the NDMP device path names by using inquire

Use the **inquire** command to determine the path names and library handle.

1. From a command prompt on the NetWorker server, type:

```
inquire -N NAS_hostname -T
```

2. When prompted, specify the NAS username and password.

NOTICE

Use the **inquire** command with caution. Running **inquire** sends the **SCSI inquire** command to all devices detected on the SCSI bus. If you use the **inquire** during normal operations, unforeseen errors can occur, resulting in possible data loss.

Determining the NDMP path names with vendor-specific commands

Before you configure an NDMP autochanger you must determine the device path names of NDMP devices and the robotic arm. [Table 109 on page 656](#) provides vendor-specific information that you can use to determine the device path names.

Table 109 Determining NDMP path names

NAS	Vendor specific information to determine autochanger and device paths (page 1 of 2)
EMC Celerra and VNX	<p>Use the Celerra or VNX Administrator program or manually query the scsidevs file. To manually query the scsidevs file, log in to the filer with the ndmp account and type:</p> <pre>server_devconfig data_mover_name -p -s -n</pre> <p>The host responds with a list of media device names, for example:</p> <pre>server_2 : Scsi device table name addr type info jbox1 c1t010 jbox ATL P1000 62200501.21 tape2 c1t410 tape QUANTUM DLT7000 245Fq_ tape3 c1t510 tape QUANTUM DLT7000 245Fq_</pre> <p>To help avoid tape drive issues, set the ntape parameter for every tape drive discovered on a particular Data Mover. For example, if a Data Mover has five tape drives configured on it, set the parameter to NDMP <code>ntape=5</code>. To modify the NDMP ntape parameter, edit the <code>/nas/server/slot_#/param</code> file, where <code>slot_#</code> correlates directly to the server number and reboot the filer.</p> <p>Note: You cannot specify a value greater than 8 for <code>ntape</code>.</p> <p><i>Configuring NDMP on EMC Celerra</i> on the EMC Online Support web site provides detailed information about configuring an EMC Celerra filer.</p>
Isilon	<p>For an NDMP local backup only, configure Backup Accelerator: Use the isi fc list command to ensure that the state of each fibre channel ports is enabled. Use the isi tape rescan --reconcile command to scan for tape devices.</p> <p>Note: <code>--reconcile</code> deletes the device entries for devices that a Backup Accelerator node no longer manages.</p> <p>Use the isi tape ls -v to display a list of current devices.</p>
DinoStor-managed	<ol style="list-style-type: none"> 1. Access the DinoStor TapeServer interface. 2. Click the Configure page. 3. Click the SCSI tab. 4. Make note of the device names and device handles.
MiraPoint	<p>On the MiraPoint filer, type:</p> <pre>diag tape inquiry ** ** **</pre> <p>Note: The device pathname is in the format <code>/dev/nrstn</code>, where <code>n</code> starts at 0 and increases one number for each tape drive. This value is constant.</p> <p>When the filer uses NDMP v2, or does not support the NDMP_CONFIG interface, you must specify the autochanger handle, <code>/dev/ch0</code>, when running the jbconfig command.</p> <p>To determine the autochanger handle, type:</p> <pre>diag changer inquiry ** ** **</pre>

Table 109 Determining NDMP path names

NAS	Vendor specific information to determine autochanger and device paths (page 2 of 2)
NetApp	<p>Log in to the appliance as root or as a Windows Administrator and type:</p> <pre>sysconfig -t</pre> <p>The host responds with a list of media device names, for example:</p> <pre>Tape drive (DS_300B:3.126L10) Hewlett-Packard LTO-4 nrst01 - no rewind device, format is: LTO-2(ro)/3 2/400GB nrst0m - no rewind device, format is: LTO-2(ro)/3 4/800GB cmp nrst0h - no rewind device, format is: LTO-4 800GB nrst0a - no rewind device, format is: LTO-4 1600GB cmp</pre> <p>where:</p> <ul style="list-style-type: none"> • (DS_300B:3.126L10) indicates the switch (DS_3--6), the port number (3) and the LUN number(10). This information must match the output in the sysconfig -v command. • nrst01 is the media device name. <p>When the filer uses NDMP v2, or does not support the NDMP_CONFIG interface, to determine the autochanger handle, type:</p> <pre>sysconfig -m</pre> <p>The host responds with the devices on the host, for example:</p> <pre>Medium changer (DS_300B:3.126L9) ADIC Scalar i2000 mc5 - medium changer device</pre> <p>where mc5 is the autochanger handle.</p>
Procom NetFORCE	<p>Log in as root and type:</p> <pre>status dm</pre> <p>If the filer uses NDMP v2 or does not support the NDMP_CONFIG interface, you must determine the autochanger handle.</p> <p>On a Procom NetFORCE filer, the SCSI device name format is <i>isp1tSSL[L]</i>, where <i>isp1</i> is the autochanger handle. The Fibre Channel device format is <i>ffx1tSSL[L]</i>, where <i>ffx1</i> is the autochanger handle.</p>

Configuring NDMP devices

You can back up NDMP data to an NDMP or non-NDMP device in a standalone or library configuration. You can also back up NDMP data to ACSLS controlled silos.

The following sections describe how to configure NDMP devices:

- ◆ [“Configuring a standalone NDMP device” on page 657](#)
- ◆ [“Configuring an NDMP autochanger” on page 658](#)
- ◆ [“Changing the block size of an NDMP device” on page 661](#)

Configuring a standalone NDMP device

Use NMC to configure a standalone NDMP tape device for Direct NDMP backups.

1. In the **Administration** window, click **Devices**.
2. Right-click **Devices** in the navigation tree, and select **New**.
3. In the **Name** attribute, specify the NDMP device in the format:

```
rd=NAS_hostname:NAS_device_handle (NDMP)
```

where:

- *NAS_hostname* is the hostname of the NAS that has the NDMP device attached.
- *NAS_device_handle* is the path of the device.

NOTICE

You must configure the NDMP device as a remote device and you must add **(NDMP)** after the pathname. Otherwise, you will receive a message similar to the following:
 NDMP device name shall be in rd=<snode>:<devname (NDMP)>
 format

4. In the **Media Type** attribute, specify the appropriate device type.
5. Specify a valid NAS administrator account in the **Remote User** attribute.

NOTICE

For an EMC Celerra and VNX filers, specify the trusted account created for backup on each NDMP-Host Data Mover. Some EMC Celerra versions require that you use a trusted account named **ndmp**. *Configuring NDMP on EMC Celerra* on the EMC Online Support web site provides detailed information.

6. Specify the password for the NAS administrator account in the **Password** attribute.
7. Under the **Configuration** tab:
 - a. Select the **NDMP** checkbox. You can only set this attribute when creating the device. You cannot change the NDMP attribute after you create the device. To change the device configuration, you must delete and recreate the device.
 - b. Set the **Target Sessions** attribute to **1**. NDMP devices do not support multiplexing.
 - c. The **Dedicated Storage Node** attribute must remain at the default value: **No**.
8. Under the **Advanced** tab, the **CDI** attribute must remain at the default value: **Not used**.
9. Optionally, change the block size the NDMP device uses. By default, NDMP devices use a block size of 60KB. If required, select a different block size in the **Device block size** field. When you configure the NDMP client, you must set the **NDMP_AUTO_BLOCK_SIZE** environment variable in the **Application Information** attribute. [“Vendor-specific Application Information variables” on page 665](#) provides more information.
10. Click **Ok**.

Configuring an NDMP autochanger

You can use an NDMP autochanger to manage Direct NDMP or Three-party backups with NDMP devices.

Configure an NDMP autochanger by using NMC or the **jbconfig** command:

- ◆ [“Configuring an NDMP autochanger with NMC” on page 659](#)
- ◆ [“Configuring an NDMP autochanger the jbconfig command” on page 660](#)

Configuring an NDMP autochanger with NMC

When you configure an NDMP autochanger in NMC, the NetWorker software first detects the NDMP devices and then configures the library.

To detect NDMP devices and configure an NDMP autochanger:

1. In the **NetWorker Administration** window, click **Devices**.
2. Right-click on the NetWorker Server and then select **Configure All Libraries**.
3. In the **Provide General Configuration Information** window, accept the default library type, **SCSI/NDMP**, and then click **Next**.
4. In the **Select Target Storage Nodes** window, click **Create a new Storage Node**.
5. In the **Storage Node Name** field, specify the hostname of the NAS. If DinoStor TapeServer manages the autochanger, specify the DinoStor hostname.
6. Select **ndmp** in the **Device Scan Type** attribute.
7. In the **NDMP User Name** and **NDMP Password** fields, specify the NAS administrator account. If DinoStor TapeServer manages the autochanger, specify the DinoStor username and password.
8. Click **Start Configuration**.
9. Click **Finish**.
10. Monitor the **Log** window for the status of the device scan.

- a. When you specify an incorrect username and password combination:

- The **Log** status window reports:

```
No configured libraries detected on storage node <storage node name>
```

- The **daemon.raw** file on the NetWorker server reports:

```
NDMP Service Debug: The process id for NDMP service is
0xb6c0b7b0
42597:dvdetect: connect auth: connection has not been
authorized
42610:dvdetect: The NDMP connection is not successfully
authorized on host 'storage_node_name'
```

To resolve this issue, relaunch the **Configure All Libraries** wizard and correct the NDMP username and password combination.

- b. If the **Log** window reports that NetWorker cannot detect the serial numbers for the library, see [“Configuring an NDMP autochanger the jbconfig command” on page 660](#) for detailed instructions.

Configuring an NDMP autochanger the `jbconfig` command

The NMC interface is the preferred method to configure an NDMP autochanger. Use the `jbconfig` command when you cannot configure the autochanger by using the NMC Configure Library wizard.

To configure an NDMP library with the `jbconfig` command:

1. Login to the NetWorker server as root (UNIX) or Administrator (Windows).
2. At the command prompt, type the following command:


```
jbconfig
```
3. Type **3** at the **What kind of jukebox are you configuring** prompt, to configure an autodetected NDMP SCSI jukebox.
4. When prompted for an **NDMP username**, specify the NAS administrator account. If DinoStor Tape Server manages the jukebox, then specify the DinoStor account.
5. When prompted for an **NDMP password**, specify the NAS administrator password. If DinoStor manages the jukebox, then specify the DinoStor password.
6. When prompted for the **NDMP Tape Server Name**, specify the NAS filer hostname. If DinoStor manages the autochanger, then specify the DinoStor hostname.
7. In the **What name do you want to assign to this jukebox device** prompt, provide a name to identify the autochanger.
8. To enable auto-cleaning, accept the default value of **Yes**, otherwise type **No**.
9. In the **Is (any path of) any drive intended for NDMP use? (yes / no) [no]** prompt, type **Yes**.
10. In the **Is any drive going to have more than one path defined? (yes / no) [no]** prompt, type **No** if you will not configure shared devices. Type **yes** to configure shared drives. [“DDS on NDMP nodes in a SAN environment” on page 207](#) provides detailed information about dynamically sharing NDMP devices.
11. The `jbconfig` command prompts for the first pathname for the NDMP devices in the jukebox.

- a. Specify the pathname in the following format:

```
NDMP_tape_server_name:device_path
```

where:

- *NDMP_tape_server_name* is the hostname of the NDMP or DinoStor Tape Server.
- *device_path* is the first device path. [“Determining NDMP device path names” on page 655](#) provides more information.

NOTICE

For a NetApp device, do not type a slash before the device name. Although the `jbconfig` command completes without errors, the NetApp filer will not recognize the tape device or autochanger.

- b. At the **Is this device configured as NDMP** prompt, type **yes**.

- c. Repeat [step a](#) and [step b](#) for all NDMP devices in the autochanger.
- d. Assign a hardware ID when prompted.
- e. To use DDS:
 - Respond to the prompts as required so that the first host will have access to the shared drive.
 - When prompted to share this drive with another host, type: **Yes**
 - When prompted, type the hostname and device path of the second host that will have access to the shared drive.
12. Complete the prompts for the second device.
13. In the **Enter the drive type of drive 1** prompt, specify the number that corresponds to the NDMP device type.
14. If each drive in the autochanger is the same model, then type **Yes**. Otherwise, type **No** then specify the appropriate device types for each additional autochanger device.
15. When prompted to configure another autochanger, type **No**.

The *EMC NetWorker Command Reference Guide* or the UNIX man page provides more information about the **jbconfig** command.

Changing the block size of an NDMP device

By default, the block size used to write data to an NDMP backup is 60KB. With the exception of EMC Celerra, when you specify the `NDMP_AUTO_BLOCK_SIZE=Y` variable for an NDMP client, an NDMP device can use the value defined in its **Device block size** attribute. [“Vendor-specific Application Information variables” on page 665](#) describes how to configure the `NDMP_AUTO_BLOCK_SIZE` variable. Consult the applicable vendor documentation to determine the block sizes supported by the NDMP filer before setting the block size for an NDMP device.

To change the block size defined for the NDMP device:

1. From the **View** menu, select **Diagnostic Mode**.
2. In the **Devices** window, right-click the NDMP device and select **Properties...**
3. Under the **Advanced** tab, select a value in the **Device block size** field.
4. Select a supported value in the **Device block size** attribute. The selected block size must not exceed the block size configured on the NAS filer.
5. Click **Ok**.

Configuring NDMP-DSA devices

When you use DSA, NetWorker sends the NDMP data to an AFTD. The steps to configure an AFTD for NDMP data and non-NDMP data are the same. [“Advanced file type devices” on page 167](#) provides detailed information.

The steps to configure an NDMP-DSA autochanger are the same as configuring a library for non-NDMP data. [“Autodetection of libraries and tape devices” on page 141](#) provides detailed information.

Configuring the Clone Storage Node

When cloning NDMP data, specify the destination storage node, called the clone “write source” (the device that receives the clone data), in the **Clone storage nodes** attribute. Prior to NetWorker 8.0, the Client resource contained the Clone storage node attribute. In NetWorker 8.0 and later, the backup Storage Node resource contains this attribute. [“Storage node selection criteria and settings for writing a clone” on page 355](#) provides details.

Creating resources to support NDMP clients

Table 110 NDMP resource requirements

Resource	Resource requirements for NDMP (page 1 of 2)
Pool	<p>When creating a pool for non-NDMP devices, select only the devices required for the NDMP clients.</p> <p>NetWorker cannot send bootstrap and index backups to an NDMP device. Configure a separate pool to direct the index and bootstrap to a non-NDMP device. “Directing client file indexes and bootstrap to a separate media pool” on page 306 provides more information.</p> <p>Note: When you do not configure a non-NDMP devices or a non-NDMP device is not available to receive the index and bootstrap backups, the NDMP client backup appears to hang.</p> <p>Auto media verification does not supported NDMP.</p> <p>When an NDMP client backup is a member of a clone-enabled group, configure a clone pool with non NDMP devices, local to the NetWorker server to receive the clone bootstrap and index.</p>
Schedule	<p>NetWorker does not support the use of synthetic full backup levels for NDMP data.</p> <p>EMC Celerra, Isilon, VNX, and NetApp filers with NDMP version 4 or later support token-based backups (TBB) to perform NDMP full, incremental, and level 1-9 backups. NetWorker supports the same number of incremental levels that the NAS vendor supports. EMC Celerra, Isilon, and NetApp documentation provide the maximum number of incremental levels that the TBB incremental backup can support.</p> <p>When you configure TBB after you update the NetWorker server from 7.6 SP1 or earlier, the first incremental or level 1 to 9 backups does not occur until after one complete full backup.</p> <p>Filers that do not support TBB, do not support incremental backups. If you select the level incr, the NetWorker server performs a full backup. However, you can schedule level backups to function like incremental backups.</p> <p>For example, define a weekly backup schedule of full on day 1, level 1 on day 2, level 2 on day 3, level 3 on day 4, and so on.</p> <p>Note: Verify the NAS storage vendor supports NDMP incremental backups before you use this feature.</p>

Table 110 NDMP resource requirements

Resource	Resource requirements for NDMP (page 2 of 2)
Group	<p>For Direct-NDMP backups, set the Savegrp parallelism value to the number of available NDMP drives.</p> <p>Note: If you set the Savegrp parallelism attribute to a higher value, there will not be enough drives to support all of the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.</p> <p>When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the Savegrp parallelism attributes.</p> <p>Note: Setting the Savegrp parallelism value for the group overrides the parallelism value defined for the NDMP clients.</p> <p>To enable automatic cloning of NDMP data, enable the Clones checkbox under the Setup tab of the group properties window then select the appropriate clone pool.</p> <p>For NetApp clients that will use checkpoint restart, set the value of the Client retries attribute under the Advanced tab to a number greater than 0.</p> <p>With the exception of Isilon, NetApp, and EMC Celerra filers, when the Force Incremental attribute is set to Yes, and the Interval attribute is set to a value less than 24 hours, a level 1-9 backup is backed up at level Full.</p>
Browse policy	<p>When you define the browse policy an NDMP client, consider the amount of disk space required for the client file index. NDMP clients with several thousand small files will have significantly larger client file indexes on the NetWorker server than a non-NDMP client. A long browse policy for an NDMP client will increase disk space requirements on the filesystem that contains the client file indexes. “Task 3: Set up policies for quick access and long term storage” on page 62 provides more details on configuring browse and retention policies.</p>

Creating and configuring the NDMP client resource

Use the NMC Client Configuration Wizard to create the NDMP client or create the client manually. EMC recommends that you use the NMC Client Configuration Wizard to create NDMP clients.

- ◆ [“Using the Client Configuration wizard” on page 663](#)
- ◆ [“Configuring the NDMP client manually” on page 671](#)

Using the Client Configuration wizard

Review this information before you use the NMC Client Configuration Wizard to create the NDMP client.

- ◆ For an NDMP configuration that includes storage node resources, configure a client resource for each storage node that you define for an NDMP backup and clone operation.
- ◆ For NDMP three-party storage nodes that use NDMP devices, repeat these steps for each NDMP storage node.

- ◆ For NDMP-DSA storage nodes, create the NetWorker client resources in the same manner as non-NDMP clients. [“Task 6: Create a backup Client resource” on page 64](#) provides details on how to create a non-NDMP client resource.

To configure an NDMP client:

1. From the **Administration** window in NMC, click **Configuration**.
2. In the expanded left pane, select **Clients** and then from the **Configuration** menu select **New client wizard**.
3. In the **Client Name** field, specify the hostname of the filer. For EMC Celerra, it is typically the CIFS server name configured on the Data Mover.
4. Select **NDMP client** and click **Next**.
5. In the **NDMP User Name** field, specify a valid NAS administrator account.
 - For an EMC Celerra and VNX filers, specify the trusted account created on each NDMP-Host Data Mover for backups. Some versions of EMC Celerra require you to use an account called **ndmp**. *Configuring NDMP on EMC Celerra* document on the EMC Online Support web site provides more information.
 - For an Isilon, specify the username and password of an NDMP administrator. The *OnFS Users Guide* on EMC Online Support web site describes how to create NDMP administrators.
6. In the **NDMP Password** attribute, specify the password for the NAS administrator account; click **Next**.
[“Unable to connect to NDMP host hostname” on page 674](#) describes how to resolve errors when configuring the NDMP client.
7. In the **NDMP backup type** attribute, select or specify the backup type. [Table 111 on page 664](#) summarizes the supported backup types for each NAS.

Table 111 Supported backup types

NAS	Supported backup types
EMC Celerra and VNX	<ul style="list-style-type: none"> • tar • dump - traverses a file tree in mixed width first and depth-first order. The optimal backup type. • vbb - used to backup the entire volume at the block level rather than at a file level. The vbb backup type reads data blocks in a more efficient method compared to traditional file-based backups. The vbb backup type does not support DDAR, TBB, and Three party backups. • ts - enables a tape silvering backup.
Isilon	<ul style="list-style-type: none"> • tar — use this backup type for Checkpoint Restart. • dump
BlueArc	<ul style="list-style-type: none"> • dump
MiraPoint	<ul style="list-style-type: none"> • image
NetApp	<ul style="list-style-type: none"> • dump - an inode-based backup that traverses a file tree in directory first and file-based order. • smtape - performs a block-level backup of a SnapMirror volume.

8. The **NDMP Array Name** field enables you to configure the same NAS device with multiple NDMP clients that have different host IDs, specify the logical name assigned to the NDMP NAS array.

Note: NDMP clients that use the same NAS device must have the same NDMP array name.

9. Review the **App Info** options and disable options, as required. EMC recommends that the default options remain enabled. Online Wizard help describes each **App Info** option.
10. In the **Advanced App Info** field, specify additional NAS specific environments variables, one per line. [Table 112 on page 665](#) provides a list of the available **Application Information** environment variables for each NAS.

NOTICE

Environment variables are case-sensitive. Use an equal (=) sign to separate the environment variable name from its value.

Table 112 Vendor-specific Application Information variables (page 1 of 4)

NAS	Variables	Definition
EMC Celerra and VNX	DIRECT= n	Optional. When you use DAR or DDAR, you must set this value to y .
	EMC_EDIR <i>nn=string</i>	Optional. This string value identifies a directory to exclude from the backup. You can use asterisk (*) as a wildcard, but only when * is the last character in the string. To include multiple directories, increment the number. For example: <ul style="list-style-type: none"> EMC_EDIR01=/fsX/DIRx EMC_EDIR02=/fsX/DIRy EMC Dart version 5.5 and later supports this variable. Note: This variable is ignored when you perform a vbb backup.
	EMC_EFILE <i>nn=string</i>	Optional. This string value determines which files to exclude from the backup. You can use the asterisk (*) as a wildcard, but only when * is the first or last character in the string, or both. To include multiple files, increment the number. For example: <ul style="list-style-type: none"> EMC_EFILE01=*mp3 EMC_EFILE02=temp* EMC Dart version 5.5 and later supports this variable. Note: This variable is ignored when you perform a vbb backup.

Table 112 Vendor-specific Application Information variables (page 2 of 4)

NAS	Variables	Definition
	OPTIONS=NT	Required. This value ensures backup and recovery all ACLs. In addition to setting this variable, axtrp must exist in the /nas/server/slot_#/netd file.
	SNAPSURE=y	Required. This value ensures that NetWorker can back up open files and generate FH information.
	USE_TBB_IF_AVAILABLE=n	Optional. The NetWorker software automatically enables TBB support for EMC Celerra filers. Specify this variable and value to disable TBB support for incremental backups and when you use the vbb backup type. When you specify this value, the backup reverts to the native level-based backup of the NAS.
	ALLOW_SINGLE_FILE_BACKUP=y	Optional. Specify this variable only when you perform a single file backup.
	NSR_NDMP_RECOVER_NO_DAR=y	Optional. Specify this variable to perform a non-DAR recovery when you set the DIRECT=y variable during the backup.
	NSR_NDMP_DDAR Note: This environment variable must be set in the operating system before invoking the either the recover or winworkr program.	Optional. Specify this variable to perform a DDAR recovery when: <ul style="list-style-type: none"> You set the DIRECT=y variable during the backup The DART version is 5.5 and later Do <i>not</i> specify NSR_NDMP_DDAR when you also use NSR_NDMP_RECOVER_DIR.
	NSR_NDMP_RECOVER_DIR=y Note: This environment variable must be set in the operating system before invoking either the recover or winworkr program.	Optional. Specify this variable to perform a DAR recovery when: <ul style="list-style-type: none"> When you set the DIRECT=y variable during the backup The DART version is 5.5 and later Do <i>not</i> use NSR_NDMP_RECOVER_DIR when you also use NSR_NDMP_DDAR.
	TS=y	Optional. Enables tape silvering.
Isilon	DIRECT=y	Required.
	HIST=F	Required for Checkpoint Restart backups to ensure the use of path based file history.
	FILES= <i>pattern</i>	Optional. Use this variable to back up only files that match the defined <i>pattern</i> . You can use wildcards in the <i>pattern</i> definition.
	PER_DIRECTORY_MATCHING=y	Optional. Use this variable along with the FILES= <i>pattern</i> variable. Matches the pattern defined by FILES across directories.
	USE_TBB_IF_AVAILABLE=n	Optional. The NetWorker software automatically enables TBB for Isilon filers. Specify this variable to disable TBB support for incremental backups and when you use the vbb backup type. When you specify this value, the backup reverts to the native level-based backup of the NAS.

Table 112 Vendor-specific Application Information variables (page 3 of 4)

NAS	Variables	Definition
	NSR_NDMP_RECOVER_NO_DAR=y	Optional. Define this variable to perform an non-DAR recovery when you set the DIRECT=y variable during the backup.
	NDMP_AUTO_BLOCK_SIZE=Y	Optional. Specify this variable to override the default block size of 60 KB when writing NDMP backups to an NDMP device. Uses the block size value defined in the Device block size attribute when you labeled the NDMP volume. “Configuring NDMP devices” on page 657 provides more information.
NetApp	FILESYSYSTEM= <i>path</i>	Optional. Use this variable to define the file system to back up and override the value in the Save set attribute of the client.
	DIRECT= y/n	Optional. When you use DAR or DDAR, you must set this value to y .
	EXCLUDE= <i>string</i>	Optional. This string specifies files to exclude from backup. The following rules apply: <ul style="list-style-type: none"> • The string must be a file name. Use file names, not absolute paths. • You can use the asterisk (*) as a wildcard, only when * is the first or last character in the string, or both. • To list multiple files, separate each name with a comma. A comma cannot appear as part of the file name. You cannot use spaces. • You can specify up to 32 strings.
	EXTRACT_ACL=y	Optional. Specify this variable to recover ACLs when you use DAR.
	SMTAPE_BREAK_MIRROR=Y	Optional when you use SMTape. During the recovery of the mirror, setting SMTAPE_BREAK_MIRROR=Y ensures that the mirror breaks and the volume becomes available for reuse. If you do not set the variable or you specify SMTAPE_BREAK_MIRROR=N, the mirror remains in the same state as at the time of backup.
	SMTAPE_DELETE_SNAPSHOT=Y	Optional when you use SMTape. When backing up the filer volumes, setting SMTAPE_DELETE_SNAPSHOT=Y ensures the removal of the mirror created during the backup, at the end of backup. If you do not set the variable or you specify SMTAPE_BREAK_MIRROR=N, each backup attempt creates a new snap mirror image.
	RECURSIVE=y	Optional for DAR and DDAR recoveries. RECURSIVE=y ensures the correct recovery of ACLs, permissions, and ownerships for all interim directories selected in the recover operation.
RECOVER_FULL_PATHS=y	Optional for DAR and DDAR recoveries. RECOVER_FULL_PATHS=y ensures the NetWorker recovers the ACLs, permissions, and ownerships for each interim directories selected in the recover operation.	

Table 112 Vendor-specific Application Information variables (page 4 of 4)

NAS	Variables	Definition
	USE_TBB_IF_AVAILABLE=n	Optional. The NetWorker software enables TBB automatically. Specify this variable to disable TBB support for incremental backups and when you use the vbb backup type. This value reverts the backup to the native level-based backup of the NAS
	UTF8=n	Optional. Provides support for UTF-8 formatted data. When you do not define this variable, the default value is n. When you set UTF8=Y during an NDMP client backup and the backup contains path names with non-ASCII characters, an index-based recovery of this backup fails with the error: "RESTORE: could not create path <i>pathname</i>" .
	NSR_NDMP_RECOVER_NO_DAR=y	Optional. Specify this variable to perform a non-DAR recovery when you set the DIRECT=y variable during the backup.
	NDMP_AUTO_BLOCK_SIZE=Y	Optional. Specify this variable to override the default block size of 60 KB when writing NDMP backups to an NDMP device. Uses the block size value defined in the Device block size attribute when you labeled the NDMP volume. "Configuring NDMP devices" on page 657 provides more information.
Mirapoint	MIRA_OPTIONS= (fromimagefull=)	Required. The (fromimagefull=) value allows full image and message (file) based backups to use the date of the image when performing the selection.
	NDMP_AUTO_BLOCK_SIZE=Y	Optional. Specify this variable to override the default block size of 60 KB when writing NDMP backups to an NDMP device. Uses the block size value defined in the Device block size attribute when you labeled the NDMP volume. "Configuring NDMP devices" on page 657 provides more information.
BlueArc	NDMP_BLUEARC_FH_NAMETYPE=UNIX	Required. This variable requests that the BlueArc filer provide UNIX-style names when backing up a CIFS share.
	NDMP_AUTO_BLOCK_SIZE=Y	Optional. Specify this variable to override the default block size of 60 KB when writing NDMP backups to an NDMP device. Uses the block size value defined in the Device block size attribute when you labeled the NDMP volume. "Configuring NDMP devices" on page 657 provides more information.

11. In the **Specify the Client Backup Options** window:

- a. To enable DSA backups, select **DSA Backup**.
- b. In the **Target Pool** field, select the pool that will receive the client data.
- c. If you use Data Domain devices, then select **Data Domain Backup**.

- d. For NetApp and Isilon filers only:
- To configure NetWorker to restart a scheduled NDMP backup that fails, select the option **Checkpoint enabled**. “[Checkpoint restart backups](#)” on page 95 provides more information on checkpoint restart backups.
 - To decrease the frequency checkpoints written during an NDMP backup, increase the **Checkpoint Granularity (in bytes)** value to 5.
- e. Click **Next**.
12. Select or specify the objects to backup:
- When the NAS supports **NDMP snapshot management extension**, you can browse and mark individual filesystems to back up. When the client supports browsing, by default, NetWorker selects all objects.
 - When the client does not support browsing, specify the save sets to back up.
 - To back up all of the file systems on the client, type **ALL**.
 - NAS versions earlier than 3 do not support the ALL save set. List the file systems one per line.
 - When you do not use the **ALL** save set, specify the filesystem name, as configured on the NAS.
 - File system names in the save set field are case sensitive.
 - For EMC Celerra backups, do not use the **ALL** save set. List the file systems, on per line excluding the root, or “/” file system.

Note: When you include the root file system, client index updates fail for hidden file systems (directories that start with a “.”) with the error: **Failed to store index entries**.

 - For EMC Celerra block-level backups, specify the entire file system mount point.
 - You cannot specify a share name.
13. To back up large client file systems, optionally schedule each file system to back up separately. For example, create two separate clients with the same name, but with different save sets.
14. Select the **Browse Policy** and **Retention Policy** for the NDMP client.
15. In the **Remote Access** attribute:
- Specify the root (on UNIX) or administrator (on Windows) account of any computer that you use to browse entries for the NDMP computer.
- NOTICE**
- For an Isilon filer in SmartLock compliance mode, specify the compadmin account.
- Specify the NetWorker server administrator account (Windows) or the root account (UNIX).
- Use this format to specify each account:
- account_name@hostname*

16. Click **Next**.

17. In the **Choose the Backup Group** window:

- To add the client to an existing group, select **Add to an existing group** then select the group name.
- To create a new group:
 - Select **Create a new group**.
 - For NetApp clients that will use checkpoint restart, set the value of the **Client retries** attribute to a number greater than 0.
 - In the **Schedule Backup Start Time**, specify the time to start the group backup.
 - Select **Automatically start the backup at the schedule time**.

18. In the **Specify Storage Node options** window:

- To use NDMP devices managed by the NetWorker server, select **Backup to NetWorker server only**.
- To backup to a remote storage node, select **Backup to the Following storage nodes**.

19. In the **Backup Configuration Summary** window, review the attributes and click **Create**.

20. Review the configuration summary then click **Create**.

21. Click **Finish** to exit the wizard.

Performing post Client Configuration Wizard steps

After the Client Configuration Wizard creates the NDMP client, modify the properties of the new NDMP client.

Modify the client parallelism

On the **Globals (1 of 2)** tab, modify the client parallelism value to the recommended value for the NDMP configuration.

- ◆ For Direct-NDMP, set the **Parallelism** attribute to **1**.
- ◆ For NDMP-DSA, the parallelism value depends on the NAS capabilities and you must set parallelism to a value that is appropriate for the NAS. Parallelism values of 4 to 8 are common. In general, the optimal parallelism setting depends on filer configuration and the amount of installed RAM.
- ◆ For an EMC Celerra and VNX filers, the parallelism value differs for the DartOS version:
 - For an EMC Celerra using DartOS v.5 and earlier, the **Parallelism** attribute cannot exceed **4**. *Configuring NDMP Backups on EMC Celerra* provides more information.
 - For VNX using DartOS v.6 and later, the maximum parallelism value is **8**. The optimal parallelism value depends on:
 - The amount of physical memory on the Data Mover.
 - The amount of physical memory allocated to the NDMP PAX configuration.
 - The value defined for the **concurrentDataStreams** parameter on the filer. *Configuring NDMP Backups on VNX* provides more information.
- ◆ For a NetApp filer, the recommended parallelism value is **8**. If required, you can use a higher **Parallelism** value however, for best performance, do not exceed **12**.

Modify the Storage Node

On the **Globals (2 of 2)** tab, specify the appropriate storage node in the **Storage Nodes** attribute. The value depends on the type of backup:

- ◆ When you perform Direct-NDMP backups with NDMP devices, specify the hostname of the NAS that manages the tape device or autochanger.
- ◆ For three-party backups, list the destination host first.
- ◆ For NDMP-DSA backups, specify the hostname of the storage node that manages the tape device or autochanger. If the NetWorker server is the storage node, specify **nsrserverhost**.
- ◆ For a DinoStor-managed NAS, specify the hostname of the DinoStor server first.

NOTICE

In NetWorker 8.0 SP1 and later, for NDMP-DSA backups, the NetWorker software uses the **Storage Node** attribute field of the NDMP client to determine which host receives the backup data. The **nsrmdmp_save** command does not require the **-M** and **-P** options. If you specify the **-M** and **-P** options, they will override the **Storage Node** attribute value.

Configuring the NDMP client manually

EMC recommends that you create a new NDMP client by using the Client Configuration Wizard. If you create the NDMP client manually, then the configuration details for each attribute in the Client Configuration Wizard apply when you create the client manually.

Review this information before you configure an NDMP client manually:

- ◆ NDMP does not support the use of directives including AES encryption. The NetWorker software ignores any value you define in the Directives attribute for an NDMP client.
- ◆ When you select **Checkpoint enabled** on the **General** tab, do not modify the **Checkpoint granularity** attribute. NDMP backups do not support checkpoint granularity and the NetWorker software ignores any value that you define for this attribute.
- ◆ In NetWorker 8.0 and later, if the NAS supports NDMP snapshot management extension, then you can browse and mark individual filesystems for backup instead of specifying the save sets in the **Save Set** attribute. You cannot use the Save set browse icon to browse the NDMP file system until you:
 - Select the **NDMP** checkbox, on the **Apps & Modules** tab.
 - Specify the NDMP username and password in the Remote user and password fields on the **Apps and Modules** tab.

Performing NDMP backups

After you configure the NetWorker server for NDMP backup data operations, you can perform scheduled or manual NDMP backups.

The steps to configure a scheduled NDMP backup are the same as configuring non-NDMP scheduled backups.

On Windows, you can manually back up NDMP data by using the **NetWorker User** program. The method to backup NDMP data is the same as a non-NDMP local backup. [“Performing a manual backup on Windows” on page 71](#) provides more information.

NOTICE

You cannot perform a three-party backup with the **NetWorker User** program

On Windows and UNIX you can perform a manual backup from a command prompt by using the **nsrndmp_save** command. [“Performing an NDMP backup from the command line” on page 672](#) provides more information.

Before performing a manual backup by using the **nsrndmp_save** command or the **NetWorker User** program, review these requirements:

- ◆ You can only perform manual Direct-NDMP backups from a NetWorker server.
- ◆ You can start a manual NDMP-DSA backups from a NetWorker server, storage node, or client. When you do not start the NDMP-DSA backup from the NetWorker server, the **servers** file on the NetWorker server and storage node, must contain the hostname of the host that initiates the backup.
- ◆ Before you perform a manual backup, you must configure the NDMP client on the NetWorker server. Manual backups use client configuration information for example, the variables defined in the **Application Information** attribute of an NDMP client.
- ◆ Direct-NDMP and three-party NDMP backups support manual DAR backups when the NDMP client contains the **DIRECT=Y** and **HIST=Y** environment variables in the **Application Information** attribute for the NDMP client.

NOTICE

To use DAR, the NAS filer must use NDMP version 4. The *EMC NetWorker Software Compatibility Guides* describes how to determine if a particular NDMP vendor supports DAR.

Performing an NDMP backup from the command line

Use the **nsrndmp_save** command to perform a manual command line NDMP backup.

The **nsrndmp_save** command does not back up the bootstrap. Without the bootstrap, you cannot perform a disaster recovery of the NetWorker server. To back up the bootstrap, run the **savegrp -G group_name** command from the NetWorker server. The **savegrp** command uses the attribute values specified for the group. For example, the pool and schedule values.

To perform an NDMP backup from the command prompt, use the following syntax:

```
nsrndmp_save -T backup_type -s NetWorker_servername -c clientname -l backup_level -t date_time -g savegroup path
```

where:

- ◆ *backup_type* is a supported backup type for the NAS filer:
 - NetApp supports the **dump** and **smtape** backup types.

- The optimal backup type for the EMC Celerra NAS is **tar** or **dump**. Use the **vbb** backup type to back up the entire volume at the block level.
 - Isilon supports the **dump** and **tar** backup types
 - BlueArc only supports the **dump** backup type.
 - MiraPoint only supports the **image** backup type.
- ◆ *backup_level* is the **full** for a full backup, **incr** for an incremental backup, or the appropriate backup level in the range 1-9. Each NAS supports full and level 1-9 backups. EMC Celerra, Isilon, and NetApp filers only support incremental level backups.
 - ◆ *date_time* is the date and time of the last backup, enclosed in double quotes. You must specify this value for **incr** level backups, but not for level 1-9 backups. When you do not specify the date and time, the backup is a native NDMP level-based backup.

NOTICE

During a NetWorker scheduled group backup, the NetWorker software supplies the date and the time information, and incremental and level backups work as expected.

Use one of the these methods to determine the date and time of the last NDMP backup:

- Review the **daemon.raw** file on the NetWorker server or the savegroup completion report for a line similar to the following:

```
42920:nsrndmp_save: browsable savetime=1296694621
```

Use the value after *savetime=* with the **-t** option.

- Specify the date and time of the last backup reported by the **mminfo** command for the NDMP save set.

Example 62

To perform an incremental backup of a NetApp client named mynetapp:

1. Determine the time of the last full backup:

```
mminfo -v -c mynetapp
```

client	date	time	size	ssid	fl	lvl	name
mynetapp	02/16/11	15:23:58	1853MB	3864812701	cbNs	full	/.../set1
mynetapp	02/17/11	15:39:58	815MB	3848036430	cbNs	incr	/.../set2

2. Specify the last backup time in **nsrndmp_save** command:

```
nsrndmp_save -T dump -s my_nwserver -c mynetapp -l incr -t "02/16/11 15:23:58" -g mygroup path
```

In NetWorker 8.0 SP1 and later, for NDMP-DSA backups, the NetWorker software uses the **Storage Node** attribute field of the NDMP client to determine which host receives the backup data. The **nsrndmp_save** command does not require the **-M** and **-P** options. If you

specify the **-M** and **-P** options, they will override the **Storage Node** attribute value. The *EMC NetWorker Command Reference Guide* and the `nsrndmp_save` man page on UNIX provide more information.

Troubleshooting NDMP configuration and backup failures

This section provides a list of the possible causes and the resolutions for NDMP backup failures:

- ◆ “Unable to connect to NDMP host hostname” on page 674
- ◆ “NetWorker features not supported on NetApp NDMP v3 and earlier” on page 674
- ◆ “No PAX threads available” on page 675
- ◆ “Failed to store index entries” on page 675
- ◆ “IO_WritePage write failed - No space left on device (28): No space left on device” on page 676
- ◆ “Error reading the FH entries from save through stdin” on page 676
- ◆ “Cannot find file history info for filename...You may still be able to recover this file with a saveset recovery” on page 676
- ◆ “nsrndmp_save: data connect: failed to establish connection” on page 677
- ◆ “nsrndmp_save: get extension list: communication failure” on page 678

Unable to connect to NDMP host *hostname*

This message appears when the NetWorker server cannot create or modify an NDMP client.

To resolve this issue ensure that the:

- ◆ Username and password specified for the client is correct and has sufficient permissions to perform NDMP operations.
- ◆ NDMP service is running on the filer.

NetWorker features not supported on NetApp NDMP v3 and earlier

Features such as Checkpoint restart require NDMP v4.

To verify the NDMP version:

1. Log in to the NetApp host as root or as a Windows Administrator.
2. Display the NDMP version:

```
ndmpd version
```

To change the NDMP version:

1. Log in to the NetApp host as root or as Windows Administrator.
2. Stop the **NDMP** process:

```
ndmpd off
```

3. Change the **NDMP** version:

```
ndmpd version 4
```

4. Restart the **NDMP** process:

```
ndmpd on
```

Cannot perform NDMP backup after the NetWorker server licenses expire

If a NetWorker sever running in evaluation mode expires before you authorize the server, NDMP devices remain disabled after the addition of the required licenses and authorization of the NetWorker server.

To re-enable NDMP devices:

1. Use NMC to connect to the NetWorker server and click the **Devices** button.
2. In the **Devices** windows, right-click the NDMP device and select **Properties**.
3. Click the **Configuration** tab and set the **Target Sessions** attribute to **1**.
4. Click the **General** tab and in the **Enabled** section, select **Yes**.
5. Click **Ok**.

No PAX threads available

This error message appears in the server_log on the NDMP Data Server when the client parallelism value for an EMC Celerra client exceeds what the EMC Celerra can support.

To resolve this issue adjust the client parallelism attribute to a value that the Celerra supports:

- ◆ For an EMC Celerra client that runs DartOS v5.0 or earlier, the client parallelism value cannot exceed 4.
- ◆ For an EMC Celerra client that runs DartOS v6.0, the maximum parallelism value supported is 8, or the value defined in the **concurrentDataStreams** variable on the EMC Celerra. By default, the **concurrentDataStreams** variable is 4.
- ◆ The maximum parallelism value also depends on the available physical memory and the amount of memory allocated to the PAX configuration. *Configuring NDMP Backups on EMC Celerra* on the EMC Online Support web site provides more information.

Failed to store index entries

This error message occurs in the **daemon.raw** file when an index backups fails due to an insufficient amount of swap space.

To resolve this issue, increase the amount of swap space available to the NetWorker server.

NOTICE

You cannot use the **NetWorker User** program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

IO_WritePage write failed - No space left on device (28): No space left on device

This error message appears in the **daemon.raw** file when the index backup fails. There is insufficient temporary space to store the index entries before the NetWorker software commits the information into the client file index.

To resolve this issue, specify a new the temp directory with sufficient disk space in one of the following ways:

- ◆ Define the **NSR_NDMP_TMP_DIR** environment variable in the **Application Information** attribute of the client.
- ◆ Define the **NSR_NDMP_TMP_DIR** as an operating system environment variable on the NetWorker server.

[“Memory and space requirements for NDMP FH updates” on page 653](#) describes how to determine the amount of disk space the NetWorker software requires to temporarily store client files index entries.

NOTICE

You cannot use the **NetWorker User** program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

Error reading the FH entries from save through stdin

This error message appears in the **daemon.raw** file of the NetWorker server when there is a communication error between **nsrndmp_save** and **nsrndmp_2fh** processes.

Resolve any communication or connection issues, then retry the backup.

NOTICE

You cannot use the **NetWorker User** program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed. [“Performing an NDMP save set recovery from the command prompt” on page 689](#) describes how to use a save set recover to restore the data.

Cannot find file history info for *filename*..You may still be able to recover this file with a saveset recovery

This error message appears in the **daemon.raw** file of the NetWorker server when FH information is missing or corrupted for the file specified in the error message. For example, NetWorker cannot update the client file index(CFI) with FH information when a backup process interruption occurs during the failover of a clustered NetWorker environment.

You cannot perform an NMDP file-by-file recover or a save set recover when the CFI does not contain the associated FH information.

To recover this file, perform a save set recover from the command prompt. [“Performing an NDMP save set recovery from the command prompt” on page 689](#) provides for further information.

NOTICE

The NetWorker server does not delete the FH files stored in the tmp directory when the CFI updates fail.

nsrndmp_save: data connect: failed to establish connection

This error message appears in the **daemon.raw** file of the NetWorker server for several reasons:

- ◆ Network connectivity or name resolution issues exist between the NetWorker server and the NDMP client.
- ◆ You specified an incorrect NDMP username or password specified for the NDMP client.
- ◆ The NDMP service is not started on the NAS filer.
- ◆ The NetWorker server cannot communicate with the NAS filer over port 10000.
- ◆ A free port in the NetWorker server's default port range (7937-9936) is not available during an NDMP-DSA backup.

[“Firewall Support” on page 853](#) provides more information about NDMP port requirements and configuration.

- ◆ A misconfigured loop router. For an EMC Celerra filer, the server **route** command utility configures the loop router. For NetApp, the **route** utility configures loop back router. The value of this setup is network-specific and depends on the number of switches and hubs between the NAS filer, NetWorker server, and NetWorker storage node.
- ◆ On the host where DSA is running, if the hostname is present in the **hosts** file, the **nsrdsa_save** process uses this name during backup. The DSA host passes the loopback entry to the NDMP data server and the connection fails. To resolve this issue, remove the hostname from the localhost list.

Knowledge base article **esg11713** on the EMC Online Support Site provides detailed troubleshooting information for this error message and other failed to establish connection failures encountered during an NDMP backup.

nsrndmp_save: get extension list: communication failure

This message appears during a NDMP local backup when NetWorker cannot determine the filer name.

To resolve this issue:

1. From a command prompt on the NetWorker server, type:

```
nsrndmpsup -c NDMP_hostname -o output_filename
```

For example:

```
nsrndmpsup -c myfiler.mnd.com -o nsrndmpsup.txt
```

2. Edit the output file that the **nsrndmpsup** command generates and search for the string **Vendor Name**. Make note of the reported Vendor Name.

For example:

```
Vendor Name = BlueArc Corp
```

3. Change to the **/nsr/debug** directory on UNIX or the *NetWorker_installation_dir\nsr\debug* directory on Windows.
4. Create new empty file and name it with the following format:

```
ndmpgettextlist_disable_VENDOR_NAME
```

where you replace *VENDOR_NAME* with the vendor name of the filer reported in the **nsrndmpsup** output file.

For example, to create this file for a BlueArc filer on UNIX, type:

```
touch "ndmpgettextlist_disable_BlueArc Corp"
```

Cloning NDMP save sets

You can clone Direct-NDMP and NDMP-DSA save sets by using the same methods used to clone non-NDMP save sets.

- ◆ [“Creating resources to support NDMP clients” on page 662](#). describes how to configure automatic cloning of NDMP data immediately after a group backup.
- ◆ [Chapter 12, “Cloning”](#) provides information on other cloning methods.

Before you clone NDMP save sets, review these requirements:

- ◆ To clone Direct-NDMP or Three-party backup data:
 - The source NAS must run NDMP version 3 or later.
 - The destination NAS can run any version of NDMP, but you cannot clone a volume cloned with NDMP earlier than version 3 to another volume.
 - You cannot clone NDMP save sets to a non-NDMP device.
 - You can clone NDMP tapes from one NDMP host to another NDMP host of the same type. For example, you can clone tapes from a NetApp filer with an attached library to another NetApp filer or to the same filer.

- ◆ You require two NDMP devices to clone the NDMP save sets, one device to perform the read operation and one device to perform the write operation.
- ◆ A scheduled or automatic clone operation clones the index and bootstrap save sets. You must have a non-NDMP device available to receive the cloned copy of the index and bootstrap backups. This non-NDMP device is in addition to the non-NDMP device that contains the source bootstrap and index backup. When you manually clone NDMP data, the clone operation does not clone the bootstrap and index data.
- ◆ You must clone NDMP-DSA backups to non-NDMP devices. You can however, clone NDMP-DSA save from one type of tape device to another. For example you can clone save sets on a DLT device to an AIT device.
- ◆ Use the **nsrclone** program to clone NDMP save sets from a command prompt. The *NetWorker 8.0 Command Reference Guide* or the *UNIX man pages* provide more information on **nsrclone** usage.

Reporting NDMP Data

The NetWorker software reports Information about NDMP clients, data, and volumes in two ways:

- ◆ The NMC reporting feature—Reports NDMP data in the same manner as non-NDMP data. Refer to [Chapter 15, “Enterprise reporting and events monitoring,”](#) provides more information.
- ◆ The **mminfo** command. Use the **mminfo** program to query the media database for NDMP volume and save set information:
 - [“Querying the NDMP volumes by backup type with the mminfo command” on page 679](#)
 - [“Querying the NDMP save sets with the mminfo command” on page 680](#)

Querying the NDMP volumes by backup type with the mminfo command

You can query save sets by backup format (NDMP or DSA) to display volume information.

For example:

- ◆ To query NDMP volumes, type:

```
C:\Users\Administrator>mminfo -q ndmp
volume      client      date      size      level  name
005D0000    sim1cifs1  6/22/2011 1036 MB   full   /fs1
005D0001    sim1cifs1  6/22/2011 173 MB    full   /fs1
005D0001    sim1cifs1  6/22/2011 862 MB    full   /fs1
005D0002    sim1cifs1  6/22/2011 348 MB    full   /fs1
```

- ◆ To query NDMP -DSA volumes, type:

```
C:\Users\Administrator>mminfo -q dsa
volume      client      date      size      level  name
NDMP.001    10.8.67.219 12/13/2011 644 MB    full   /vol/vol0
NDMP.001    10.8.67.219 12/13/2011 402 MB    full   /vol/vol1
NDMP.001    10.8.67.219 12/13/2011 402 MB    full   /vol/vol1
NDMP.001    10.8.67.219 12/13/2011 36 MB     full   /vol/vol2
```

Querying the NDMP save sets with the mminfo command

Query the media database to determine which save sets are NDMP save sets and the status of an NDMP save set in the media database. NDMP save set status information is important when performing NDMP recoveries.

- ◆ To perform a browseable NDMP recover, the `ssflags (fl)` field for an NDMP save set must contain a `b`. The `b` value denotes a browsable save set.
- ◆ To perform a save set recover from the NetWorker User program, the `ssflags (fl)` field for an NDMP save set must contain a `b`.
- ◆ An NDMP save set contains an `N` attribute in the `ssflags (fl)` field.
- ◆ An NDMP-DSA save set contains an `s` attribute in the `ssflags (fl)` field.

In the following example, the NDMP save set status is recoverable (`r`). To recover the data, you can only perform a save set recovery from a command line.

```
mminfo -av
```

volume	type	client	date	time	size	ssid	fl	lvl	name
vol1	dlt	clnt	6/22/2011	3:15:12	1036MB	3842140553	hr N	full	/fs1

In the following example, the NDMP-DSA save set status is browsable (`b`). Recover the data by using the **NetWorker User** program, or from the command line. A browseable NDMP-DSA save set supports browsable and save set recoveries.

```
mminfo -av
```

volume	type	client	date	time	size	ssid	fl	lvl	name
vol1	dlt	clnt	6/22/2011	3:15:12	36MB	4259813785	cb Ns	full	/fs1

Performing NDMP recoveries

NetWorker uses the `nsrndmp_recover` program to coordinate recover operations between the NetWorker software and the NDMP client. The `nsrndmp_recover` program does not move data to the NDMP client. When the `nsrndmp_recover` program identifies an NDMP-DSA save set, `nsrndmp_recover` automatically invokes the `nsrdsa_recover` program on the same host that runs the `nsrndmp_recover` command.

To recover NDMP data, you can run the `nsrndmp_recover` program from a command prompt, or use one of following programs, which automatically start the `nsrndmp_recover`:

- ◆ **recover**—The command line program on Windows and UNIX.
- ◆ **winworkr**—The **NetWorker User** GUI on Windows.
- ◆ The NMC Recovery Wizard. [“Using the Recovery Wizard” on page 374](#) describes how to use the Wizard to recover data.

During the recovery process, the `nsrndmp_recover` program passes `nlist` information to the NDMP client. There are three methods to recover NDMP backups:

- ◆ Index-based file-by-file recover — The nlist includes file offset and ACL information. When you recover many files, the recover process uses a significant amount of system resources on both the NetWorker server and the NDMP client to build and process the nlist information. [“Performing an NDMP index-based file-by-file data recovery” on page 683](#) provides more information
- ◆ Full save set recovery—The nlist only includes the path to the recovery directory, down to and including the mount point. When you recover many files, the recover process uses less system resource intensive than an index-based NDMP recover to build and process the nlist information. [“Performing a full or Directory Restore of NDMP data by using a save set recovery” on page 687](#) provides more information
- ◆ NDMP directory restore — A partial save set recovery of a single file or single directory. [“Performing destructive save set recoveries for vbb backups” on page 690](#) provides more information.

For example, when the NetWorker software writes NDMP data a remote storage node, start the **recover** program on the NetWorker storage node to prevent the data from traversing the network.

Note: When you start the **recover** program on the NetWorker server, the data flows from the storage node to the NetWorker server and from the NetWorker server to the NDMP client, over the network.

NDMP recovery requirements

[Table 113 on page 681](#) summarizes the requirements for each recovery feature.

Table 113 Requirements of each NDMP recovery feature

Feature	Requirement (page 1 of 2)
scanner	<ul style="list-style-type: none"> • You cannot use the scanner command with the -i, -f and -r options on an NDMP volume. • You cannot use the scanner command on a volume that contains NDMP and non-NDMP save sets when you load the volume in an NDMP device. The <i>Scanner command usage</i> technical note provides more information about using the scanner command with NDMP data.
Cross platform recoveries	You can recover NDMP data to different NDMP client however, you cannot perform a cross platform recover. Recover NDMP data to an NDMP client that is the same brand, a compatible model, and the same operating system as the original NDMP client.
Devices	<p>Recover Direct-NDMP and Three-party backups performed to an NDMP device from an NDMP device. To improve recover performance from an NDMP tape device, configure the tape device to support variable length records.</p> <p>Recover NDMP-DSA backups from a non-NDMP device.</p>
Localized environments	When recovering data in a localized NDMP environment, the Index Recover status window shows the process in English and not the localized language.
NDMP-DSA	For better recovery performance, start the recover process on the NetWorker host where the backup volume resides.
Immediate recoveries	Run the nsrmdmp_recover program on the storage node with the locally attached backup device to perform an immediate recovery of NDMP-DSA data.

Table 113 Requirements of each NDMP recovery feature

Feature	Requirement (page 2 of 2)
EMC Celerra and VNX	<p>During recover operation the filer skips char and block special files. The following error message appears: Warning: /fs1/SPE_REL/my.char_file has an unknown file type, skipping</p> <p>When you recover named pipes:</p> <ul style="list-style-type: none"> • If the recover directory contains 10,000 or more named pipes, then the recover will fail. Ensure that the recover directory contains less than 10,000 named pipes. • The recover process changes the file permissions. The NetWorker software recovers named pipes as normal files. <p>The <i>Configuring NDMP for VNX</i> on EMC Online Support describes how to recover a tape silvering backup to a different data mover.</p>
Blue Arc	<p>The recover process creates a \$__NDMP__ directory at the root level of the recovery file system when you recover more than 1,024 files. The directory contains the file list that the NetWorker server uses for an index recovery. Do not change the directory and its contents during an active recovery operation. When a recovery is not in progress, you can delete the directory.</p> <p>While performing NDMP backup and recover operations, a message similar to the following may appear: NDMP session-Unknown environment variable name ignored. You can ignore this message.</p>
Mirapoint	<p>After a full backup recovery for a Mirapoint system, reboot the Mirapoint system. An incremental recovery does not require a reboot.</p>
vbb	<p>When you set backup type for an EMC Celerra or VNX filer to vbb, the NetWorker software performs a block-based backup.</p> <p>The NetWorker software recovers the data to a raw device. Use the Relocate data option to specify the raw device.</p> <p>When you use deduplication on the source or target file system, you cannot perform an index based file-by-file recover.</p> <p>When you perform a destructive save set recover, the recover process:</p> <ul style="list-style-type: none"> - Recovers the data to the original location or an alternate location. - Overwrites existing data. - Overlays the data at the file system level and reimposes the saved image on the file system. <p><i>Configuring NDMP backups on Celerra</i> and <i>Using Celerra Data Deduplication</i> on EMC Support Online provides detailed information about how to prepare the filer before you perform FDR.</p>

DAR and DDAR

By default, the NDMP recover process reads an entire tape from start to finish. The recover process extracts the data as it encounters the data on the tape. For large backup images, recovery is slow.

The DAR/DDAR recovery process:

- ◆ Provides the ability to recover a file or directory from the exact location on a tape.
- ◆ Only passes the directory path to the NAS filer.
- ◆ Reduces the size of the nlist information that the recover process stores in memory. During the recover process, the NAS filer assumes that the directory path includes all cataloged files and directories.

- ◆ Does not sequentially read the file or record numbers on the tape to locate the data. This reduces the amount of time you require to recover specific files from a backup.

EMC Celerra with DART version 5.5 and later and NetApp filer with NDMP v4 and OnTap version 6.4 and later support DDAR.

Note: [“Creating and configuring the NDMP client resource” on page 663](#) describes how to configure the DAR and DDAR Application Information attributes for NDMP clients.

Use the **recover** command or the **NetWorker User** program to perform DAR and DDAR recoveries. You cannot use the **nsrmdmp_recover** program to perform DAR/DDAR recoveries.

When not to use DAR or DDAR

The DAR and DDAR recoveries send multiple path names across the network to the NDMP Data Server and, in three-party configurations, to the NetWorker server. The recover process stores the path names in memory on the NDMP Data Server. Recoveries of a large amount of data from a large save set can negatively impact the network and the NDMP Data Server resources.

Do not use DAR and DDAR to recover:

- ◆ Several thousands of files in a single index-based recover operation.
- ◆ A specific directory structure containing several thousand or millions of files.

To perform a non-DAR-based recovery of a save set when you set the **DIRECT=y** at the time of backup, first define the **NSR_NDMP_RECOVER_NO_DAR=y** variable in the **Application Information** attribute of the NDMP client.

Performing an NDMP index-based file-by-file data recovery

Perform an NDMP index based file-by-file recover in the same manner as a non-NDMP data recover. You can restore the data to the original NDMP client or directed to a different NDMP client.

Before you perform an index-based file-by-file recover, review the following information:

- ◆ Set the **HIST=y** in the application information attribute of the NDMP client at the time of the backup. [Table 112 on page 665](#) provides more information about the NDMP **Application Information** attributes.
- ◆ The NDMP save set must be browsable. You cannot perform a browsable recover of a **recoverable** or **recyclable** save set. [“Reporting NDMP Data” on page 679](#) describes how to determine the status of an NDMP save set.
- ◆ Do not use an index-based recovery to recover a large numbers of files or directories. For better recovery performance, use a save set recover. [“Performing a full or Directory Restore of NDMP data by using a save set recovery” on page 687](#) provides more information.

- ◆ To perform an index-based file-by-file recover:
 - Use the **NetWorker User** program on a Windows host. [“Performing an NDMP index-based file-by-file recover using the NetWorker User program” on page 684](#) provides detailed information.
 - Use the **recover** program. [“Performing an NDMP index-based file-by-file recover from a command prompt” on page 686](#) provides detailed information.

Performing an NDMP index-based file-by-file recover using the NetWorker User program

On Windows, to recover data to the original NDMP client or to a different NDMP client:

1. Open the **NetWorker User** program and connect to the NetWorker server.

NOTICE

If you receive the error **“No file indexes were found for client *client_name* on server *server_name*. Try connecting to a different NetWorker server”** and you selected the correct NetWorker server, then ensure that you selected a browseable save set. Alternatively, perform a save set recover. [“Performing an NDMP save set recovery from the command prompt” on page 689](#) provides the instructions to perform an NDMP save set recover from a command prompt.

2. Select **Recover** to open the **Source Client** window.
3. Select NDMP client with the data to recover and click **OK**. The local client is the default selection.
4. Select the destination client for the recovered data and click **OK**. If the destination client is not the source client, ensure the NAS filer is the same brand, a compatible model and the same operating system as the source NDMP client.
5. Optionally, recover the data from an earlier backup time. The **Recover** window appears with the latest version of the backup files. To recover data from an earlier backup, change the date and time of backup using one of the following methods:
 - a. Change the browse time for all files in the recover window:
 - From the **View** menu, select **Change Browse Time**.
 - In 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, a minute, and the letter a for A.M. or p for P.M. Use the 12-hour format.
 - Click **OK**.
 - b. View all versions of the selected file system object:
 - Highlight the file or directory for review.
 - From the **View** menu select **Versions**.
 - Once you locate the version to recover, change the browse time. To change the browse time, highlight the volume, directory, or file and click **Change Browse Time**. The **Version** window closes and the Recover window reflects the new browse time.

6. Optionally, search for the files. To search for and recover the most recently backed-up version of a file or directory:
 - a. From the **File** menu, select **Find**.
 - b. Type the name of the file or directory. Use wildcards to expand the search; without wildcards, partial filenames do not provide any results.
7. Mark the data to recover. To select file system objects to recover:
 - a. In the left pane of the **Recover** window, click the appropriate directory folder.
 - b. Mark each directory or file to recover by selecting the checkbox next to each directory or file.
8. Optionally, relocate the data to a different location. By default, the recover process recovers the selected files to the original location.

NOTICE

NDMP recoveries will *always* overwrite existing files. EMC recommends that you recover the NDMP data to a different location, to avoid data loss.

To relocate the files to a different location:

- a. Select **Recover Options** from the **Options** menu.

NDMP recovery do not support the following options:

- Rename recovered file
- Discard recovered file
- Prompt for every file conflict

NOTICE

NDMP recoveries will *always* overwrite existing files. EMC recommends that you relocate the NDMP data to a different location, to avoid loss.

- b. In the **Relocate Recovered Data To** field, type the full pathname of the target directory, click **OK**.

NOTICE

The target directory is a literal string and *must* match the path as seen by the NAS filer in its native OS, exactly. Otherwise, the recover process uses the original location and overwrites existing files with the same name.

9. Optionally, to view the volumes required to recover the marked file system objects, from the **View** menu, select **Required Volumes**.
10. Click **Start** to begin the recovery. If any required volume is not available to the NetWorker server, a volume status warning appears.

When this warning appears:

- a. Click **No**.
- b. From the **View** menu, select **Required Volumes**.

- c. Ensure that the NetWorker software can mount each listed volumes in an available device.
- d. Reattempt the recover operation.

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 completes successfully, 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 an NDMP index-based file-by-file recover from a command prompt

This section applies to command line recoveries from a Windows and UNIX client.

To avoid using the Windows version of **recover.exe** on Windows operating systems, perform one of the following actions:

- ◆ Specify the full path to the **recover** program. For example: C:\Program Files\EMC NetWorker\nsr\bin\recover.exe.
- ◆ Ensure that the **\$PATH** environment variable contains the *NetWorker_install_path*\bin directory before %SystemRoot%\System32.

To recover NDMP data from a command prompt on a UNIX or Windows NetWorker host:

1. From the command prompt, type:

```
recover -s NetWorker_servername -c client_name
```

where:

- the **-s** *NetWorker_servername* option specifies a particular NetWorker server on the network to use when recovering data.

When you do not use the **-s** option, the **recover** program tries to connect to the first computer listed in the servers file. When the servers file does not contain any servers, or lists more than one server, the **Change Server** window appears, and you can select the server.

- the **-c** *client_name* option specifies the source NDMP client.

2. When prompted, type the directory to browse, for example:

```
cd /mydirectory
```

3. Use the **add** command to add the required files or folders to the recover list. The *EMC NetWorker 8.0 Command Reference Guide* provides a complete list of options for the **recover** command.

4. When restoring NDMP data, EMC recommends that you relocate the NDMP data to a different location. The NDMP protocol does not allow you to handle naming conflicts. The recover operation overwrites existing files with the same name.

- To relocate the data to a different directory, type:

```
relocate destination_directory_name
```

The target pathname for *destination_directory_name* is a literal string and must match the path as seen by the NAS filer in its native OS, exactly. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

- To recover the data to a different host, type:

```
relocate target_hostname::/mount_point
```

NOTICE

Data ONTAP may require you to add a backslash (\) after the mount point. For example, *target_hostname::\mount_point*.

5. After you add all the required files, type:

```
recover
```

Performing a full or Directory Restore of NDMP data by using a save set recovery

You perform an NDMP save set recover in the same manner as a non-NDMP save set recovery. You can recover data to the original NDMP client or directed the data to a different NDMP client of the same platform.

Before you perform a full save set recover, review the following information:

- ◆ Use a full save set recovery to recover all files and folders in an NDMP data save set or to recover an entire directory within an NDMP save set. You cannot use the **NetWorker User program** to perform an **NDMP Directory Restore**. [“Performing an NDMP save set recovery from the command prompt” on page 689](#) provides more information.
- ◆ To use the **NetWorker User** program on Windows, a client file index entry for the save set must exist. When the index entry for the save set does not exist, the recover fails with an **“index not found”** error. When the client file index entries do not exist for the save set, use the `nsrndmp_recover` program with the **‘-v off’** option. [“Performing an NDMP save set recovery from the command prompt” on page 689](#) provides more information.
- ◆ You cannot perform a save set recover from the **NetWorker User** program when the save set status is eligible for recycling (E). The recover process requires a recoverable (r) or browsable (b) save set status.
 - [“Reporting NDMP Data” on page 679](#) describes how to determine the status of an NDMP save set.
 - [“Recovering a save set entry into the client file index and media database” on page 403](#) describes how to change the status of a save set.

- ◆ A save set recover reads the entire tape set, from beginning to end, to find and recover the requested files. The recovery process completes when the recover operations reads all required tapes in their entirety.
- ◆ As each file recovers, the file name appears on the target share but the file size is 0 KB. The actual file size update occurs after the recovery completes.
- ◆ The following sections describe how to perform a full save set recover:
 - “[Performing an NDMP save set recover by using the NetWorker User in Windows](#)” on [page 688](#) describes how to recover data from a Windows host by using the **NetWorker User** program.
 - “[Performing an NDMP save set recovery from the command prompt](#)” on [page 689](#) describes how to recover NDMP data from a command prompt on Windows and UNIX hosts.
 - “[Performing destructive save set recoveries for vbb backups](#)” on [page 690](#) describes how to recover from a vbb backup from a command prompt by using the **nsrndmp_recover** program on Windows and UNIX hosts.

Performing an NDMP save set recover by using the NetWorker User in Windows

To perform a save set recovery of a Windows NDMP client:

1. Start the **NetWorker User** program.
2. In the **Change Server** window, select the NetWorker server and click **Ok**.
3. Select **Options > Recover Save Sets**.
4. In the **Source Client** window, select the appropriate NDMP client and click **Ok**.
5. In the **Save Sets** window, select the name of the save set.
6. Select the version of the save set, if there are multiple versions. You can also select the cloned version of a save set, if applicable.
7. To recover specific files and directories instead of the entire save set:
 - a. Click **Files**.
 - b. Specify the files and directories, one per line.
 - c. Click **Ok**.

NOTICE

Do not use this method to mark tens of thousands of files. Instead, use perform an NDMP Directory Restore. Marking many files and directories generates a large **nlist** and requires intensive resources on both the NetWorker server and the NAS filer.

8. Click **Recover Options**.

An NDMP data recovery does not support the following options:

- Rename recovered file
- Discard recovered file
- Prompt for every file conflict

NOTICE

EMC recommends that you relocate the NDMP data to a different location. NDMP recoveries *always* overwrite existing files.

9. To recover the data to a pathname that is different from the original backup location, in the **Relocate Recovered Data To** field, type the full pathname of the destination directory, then click **Ok**.

NOTICE

For NDMP data recoveries, the target pathname is a literal string and *must* exactly match the path as seen by the native OS on the NAS filer. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

- Use the **Relocate recovered data to this raw device** option when performing a SnapImage destructive restore. The *NetWorker SnapImage Module 2.5 Solaris Version Installation and Administration Guide* and the *NetWorker SnapImage Module 2.5 Windows Version Installation and Administration Guide* provides more information.
10. To recover the data to a different NDMP client, specify the name of the client to receive the NDMP data in the **Destination Client** field.
 11. To view the volumes required to perform the recover, select **View > Required Volumes**
 12. Click **OK** to begin the recovery. The recovery status appears in the **Recover Status** window.

NOTICE

When the recover operations fails with the error: **“Failed to propagate handle <number> to child process: Access is denied”** the save set is not in the client file index of the NDMP client. Perform a save set recover from a command prompt. [“Performing an NDMP save set recovery from the command prompt” on page 689.](#) provides more information.

Performing an NDMP save set recovery from the command prompt

To perform a save set recovery to the original NDMP client or to a different NDMP client, use the **nsrndmp_recover** command.

For example:

```
nsrndmp_recover -s NetWorker_server -c source_ndmp_client -S ssid/cloneid -v
off -m target_ndmp_client::/target_path /source_path
```

where:

- ◆ *source_ndmp_client* is the hostname of the source NDMP client.
- ◆ *target_ndmp_client* is the hostname of the destination NDMP client.
- ◆ */source_path* is the original location of the data.
- ◆ */target_path* is the location to recover the data.

NOTICE

EMC recommends that you relocate the NDMP data to a different location. NDMP recoveries *always* overwrite existing files. The `/target_path` is a literal string and must exactly match the path as seen native OS on the NAS filer. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

- ◆ **-v off** allows you to restore data when client file index of the NDMP client does not contain information about the NDMP save set.

In the following examples, the NetWorker server is **mars** and the backup client is **venus**.

- To recover a mount point `/mnt` from a backup of NDMP host **venus** to a directory `/newmnt` on NDMP host **jupiter**, type:

```
nsrndmp_recover -s mars -c venus -S 123456789 -v off -m
jupiter::/newmnt
```

- To recover a mount point `/mnt` from a backup of NDMP host **venus** to NDMP host **pluto**, type:

```
nsrndmp_recover -s mars -c venus -R pluto -S 123456789 -v off -m
/mnt
```

NOTICE

Data ONTAP may require you to add a slash (`/`) after the mount point. For example, `target_hostname::mount_point/`.

Performing destructive save set recoveries for vbb backups

Use the `nsrndmp_recover` command with the `-r raw_device` and `-m mount_point` options to perform a destructive save set recovery.

NOTICE

Do not perform a NDMP Directory Restore from a vbb backup of a Celerra de-duplicated filesystem.

For example:

On a Microsoft Windows system to perform a destructive save set recovery to the `/data` drive:

```
nsrndmp_recover -s mars -c venus -m /data -r raw_device_name -S
2674606849
```

On a UNIX system, the following command performs a destructive save set recovery to the `/dev/c1t1d0s0` device, mounted at the `/` file system:

```
nsrndmp_recover -s mars -c venus -r /dev/c1t1d0s0 -S 2674606849 -m /
```

The *EMC NetWorker 8.0 Command Reference Guide* or the UNIX man page provides more information about the `nsrndmp_recover` command.

If you do not specify the `-r` option when you use the `-m`, the recover operation:

- ◆ Is nondestructive.
- ◆ Operates at the file or directory level, rather than the file system level.

This nondestructive restore overwrites existing files on the destination that have the same names as those in the recovery list. Other data remains untouched on the file system.

Use this nondestructive method to:

- ◆ Perform a directory level recovery on a high density file system.
- ◆ Recover many files in one directory.

Troubleshooting NDMP recover

This section provides a list of the possible causes and the possible resolutions for NDMP recovery issues.

- ◆ [“RESTORE: could not create path pathname” on page 691](#)
- ◆ [“These files were not restored \(Restore failed with error, or file/directory specified but not found in backup\)” on page 691](#)

RESTORE: could not create path *pathname*

This error message appears when restoring NetApp data. This error, when encountered, appears in the **daemon.raw** file of the NetWorker server and the recovery output.

To resolve this issue:

- ◆ Ensure that you specify a source and a target path during the recover that exists on the target filer.
- ◆ If you set the **UTF8=Y** application information variable during an NDMP client backup and the backup contains path names with non-ASCII characters, then perform a save set recover. Index-based recoveries will fail with this error message.

These files were not restored (Restore failed with error, or file/directory specified but not found in backup)

This error message appears in the **daemon.raw** file of the NetWorker server and the in the recovery output.

To resolve this issue:

- ◆ Ensure that the file or directory specified during the recover, exists in the save set.
- ◆ Ensure the pathname specified to relocate the data exists on the destination filer. For NDMP data recoveries, the target pathname is a literal string and *must* exactly match the path as seen by the native OS on the NAS filer.

CHAPTER 22

SNMP Module

This chapter covers these topics:

- ◆ [SNMP traps](#) 694
- ◆ [Configuring NetWorker SNMP notifications](#) 694
- ◆ [Configuring SNMP management software](#) 696

SNMP traps

The NetWorker Simple Network Management Protocol (SNMP) Module allows NetWorker servers to send notification messages to SNMP management agents.

SNMP-enabled network management software must be configured 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 (such as the NetWorker server) to the SNMP event manager.

The types of traps that the NetWorker server sends are determined when the NetWorker SNMP notification is configured within the NetWorker server. Typical traps include warnings, critical errors, and other messages from the NetWorker server. [“Configuring NetWorker SNMP notifications” on page 694](#) provides instructions on configuring the NetWorker SNMP notification.

Configuring NetWorker SNMP notifications

NetWorker software provides notifications to a variety of resources about NetWorker server events. The NetWorker SNMP Module is one of those resources. The module then forwards the notification to the SNMP management software by using the **nsrtrap** program. When you configure the SNMP notification, you include the IP address or hostname of the SNMP management server, along with other **nsrtrap** command line options, such as the SNMP community and the trap type.

In order to configure the NetWorker SNMP notification, you must first enable the NetWorker SNMP module. The *NetWorker Installation Guide* provides information on enabling and licensing the NetWorker software.

Command line options for nsrtrap

The NetWorker SNMP Module uses the **nsrtrap** program to communicate SNMP traps from the NetWorker server to the SNMP management software. [Table 114 on page 695](#) lists the command line options that can be included in the Action attribute when the SNMP notification is configured.

Table 114 Command-line options for nsrtrap

Option	Description
-c <i>community</i>	Specifies the SNMP community that is authorized to receive traps from the NetWorker server. SNMP communities are configured 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 SNMP servers to limit the communities from which the server accepts traps. If a community other than Public is configured on the SNMP server, include the appropriate community name by using this option when you configure the SNMP notification.
-t <i>trap_type</i>	Sets the type of trap the NetWorker SNMP Module sends to the SNMP server. The default setting is 6, which means that this is an “enterprise-specific” trap. Because traps that the NetWorker server sends are notifications (for example, error messages), the default setting is normally correct and should not be changed. This option should be used only if you intend to send a specific trap other than a normal NetWorker notification.
-s <i>specific_type</i>	A generic setting that can be used to identify the type of trap the NetWorker server is sending. This option can be set to any integer value and may be used in conjunction with different SNMP notifications to distinguish different traps from the NetWorker server. For example, you can create multiple SNMP notifications: one for critical messages, another for warnings, and another for other events or priorities. You can then use the -s option to differentiate the various notifications so that the SNMP management software can determine which type of trap is being sent. You could create one notification called Critical SNMP Notification, and include the -s option in the Action attribute: nsrtrap -s 1 host With this setting, the SNMP management software can be configured to recognize that NetWorker traps with the specific trap type of 1 are critical messages. Additional SNMP notifications can have other settings for the -s option to further differentiate various traps from the NetWorker server.
-v	Sets the output mode to verbose. In verbose mode, nsrtrap echoes the community, trap type, specific trap type, and the hostname or IP address at the command-prompt.

Modifying preconfigured NetWorker SNMP notification

The NetWorker server has a preconfigured SNMP notification that can be modified if necessary. The only modification that can be made to this notification is to add or remove command line options to the Action attribute.

To modify the preconfigured notification request:

1. In the **Administration** window, click **Configuration**.
2. Select **Notifications**.
3. Right-click the **SNMP notification request** and select **Properties**.

4. In the **Action** attribute, enter any necessary options for the **nsrtrap** command, such as the SNMP community. “[Command line options for nsrtrap](#)” on page 695 provides information about command-line options.
5. Click **OK**.

The events and priorities associated with the preconfigured SNMP notification cannot be modified. “[Creating NetWorker SNMP notifications](#)” on page 696 provides instructions on how to set different events and priorities for the SNMP notification.

Creating NetWorker SNMP notifications

To create additional NetWorker SNMP notifications:

1. In the **Administration** window, click **Configuration**.
2. Select **Notifications**.
3. Right-click the **SNMP notification request** and select **New**.
4. Enter a name for the **SNMP notification**.
5. (Optional) For the **Comment** attribute, enter a description of the notification.
6. Select the events and priorities that the notification should communicate to your SNMP server.

Note: The events and priorities cannot be modified after the notification is created.

7. For the **Action** attribute, type:

- Windows servers:

```
%NetWorker_install_path%\bin\nsrtrap network_management_station
```

- UNIX/Linux servers:

```
/usr/sbin/nsrtrap network_management_station
```

where *network_management_station* is the DNS name or IP address of the host on which the SNMP management software is running.

Include options for **nsrtrap**, such as **-c community**, in this attribute if necessary. “[Command line options for nsrtrap](#)” on page 695 provides more information about command-line options.

8. Click **OK**.

Configuring SNMP management software

In order for the SNMP management software to accept traps sent by NetWorker servers, it must be configured to recognize the traps. Configuration procedures vary by the type of management software you are using.

For specific instructions on configuring the types of acceptable traps, refer to the SNMP management software documentation.

NetWorker SMI Network Management Private Enterprise Code

When configuring management software to accept traps, you must also indicate the specific type of trap to accept. 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

Typically, once the network management software is configured to accept traps from NetWorker servers, an icon of each NetWorker server appears on the network management console. These examples show how the software can then be configured:

- ◆ To indicate that a trap was received (for example, the NetWorker server icon may blink or change color).
- ◆ Track pending, alert, and other configured messages.
- ◆ 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.

You may also want to create additional SNMP notification schemes in the NetWorker Administrator program that have different priorities and events. You can use the **-s *specific-type*** command line option for **nsrtrap** so that the SNMP management software can differentiate the traps sent by the various notification schemes. [“Command line options for nsrtrap” on page 695](#) provides more information about setting the **-s *specific-type*** option.

CHAPTER 23

DiskXtender Data Manager File System Support

This chapter covers these topics:

- ◆ Supported configurations 700
- ◆ DiskXtender Data Manager file system overview 700
- ◆ Backup of DXDM file systems 702
- ◆ Recovery of DXDM file systems 705

Supported configurations

These configurations are supported for the backup and recovery of DiskXtender Data Manager (DXDM) file systems:

- ◆ DXDM software and NetWorker server, client, or storage node software installed on the same computer.
- ◆ DXDM software installed on a computer that is a client of a computer that is running NetWorker server, client, or storage node software.

NOTICE

The Archive feature does not work for DXDM file systems.

The *EMC NetWorker Software Compatibility Guide* provides information about supported operating systems and file systems.

Path information

NetWorker backup and recovery requires the `/etc/dxuldm.path` file when supporting a DXDM file system. The file is automatically created during DXDM installation.

Permissions

The NetWorker daemon `nsrexecd` controls automated backup and recovery of DXDM file systems. The daemon is configured to run with set user ID (suid) root permissions. Root permissions are also required to perform manual NetWorker backup and recovery operations with DXDM file systems.

DiskXtender Data Manager file system overview

DXDM file systems use the following enhancements that are not found in standard file systems:

- ◆ The creation and retention of DMAPI information for each file.
- ◆ The ability to migrate files to a storage target.
- ◆ The ability to purge file data from the file system after migration.
- ◆ The retention of a data stub for each purged file.

Through these enhancements DXDM provides file system access to large numbers of files while storing the bulk of the file data on one or more storage target systems.

File data in a DXDM file system

When a file is placed into a DXDM file system it is available to list, view, access, and change the same as in any standard UNIX file system. The difference is that as soon as a file is placed into a DXDM file system, DMAPI metadata is created for the file. This extended metadata permits the transparent archiving of the file's data outside of the file system.

Some of the DMAPI information that is tracked on each file includes:

- ◆ The file's migration status.
- ◆ The file's purge status.
- ◆ The file's data stub size.

After a period of time the file's data is migrated to a storage target. As shown in [Figure 49 on page 702](#), the storage target can be a DXSM system or an EMC Centera Storage System (EMC Centera). After a file's data is migrated it exists on both the DXDM file system and the storage target.

After a period of time, unchanged files are purged from the file system. The file system retains the file's DMAPI metadata and data stub. The data stub consists of a user-configured number of bytes from the beginning of the file.

User access to file data

[Figure 49 on page 702](#) depicts eight DXDM file systems on one host system. Four of the file systems use a DXSM storage target and four use a EMC Centera target. The diagram shows that the data from each DXDM file system exists in its own namespace on the storage target.

- ◆ On a DXSM storage target, relative pathnames and capabilities are used to identify files.
- ◆ On a EMC Centera storage target, the EMC C-Clip™ data is created to identify each file.

This storage target namespace information is not required by DXDM file system users for access to the files. Users need only know the full pathname of a file on the DXDM host system or, when access is provided over NFS, its relative pathname from the file system's NFS mount point.

When a file on a DXDM file system is changed it is marked as not migrated. Any previously migrated data for the changed file is retained on the storage target but renamed with a version label. When a file is deleted from a DXDM file system it is fully removed. Any previously migrated data is retained on the storage target and renamed with a version label and a deleted label.




Figure 49 Prototypical DXDM installation

Backup of DXDM file systems

A DXDM file system consists of files and metadata. The files can be in a variety of migration and purge states. The metadata is contained in the file systems’s DMAPI attributes file. [Table 115 on page 702](#) shows the types of files and metadata that can exist in a DXDM file system and notes the types that are included in a NetWorker backup.

Table 115 DXDM file system data types included in a NetWorker backup

Data type	Included
New files which have not been migrated.	Yes
Changed files which have not been migrated since being changed.	Yes
Files which have been migrated but not purged.	Yes
Data stubs for files which have been migrated and purged.	Yes
File data which exists only on a storage target system.	No
DMAPI metadata.	Yes ¹

1. The DMAPI metadata is included by default. It can safely be excluded. [“Excluding the DMAPI attributes file” on page 703](#) provides more information.

File data which has not changed since being migrated is protected by the storage target. NetWorker backup automatically excludes this data from its client/save sets when the data has been purged. This results in these benefits:

- ◆ Smaller client/save sets.

Purged data which is protected on the storage target is excluded from the client/save set. This is normally the bulk of the data in a DXDM file system.
- ◆ Much faster backups.

In addition to the time savings derived from the smaller client/save set, NetWorker backup prevents the lengthy process of retrieving purged file data from the storage target.

DXDM file system client/save sets should exclude the DMAPI metadata. [“Excluding the DMAPI attributes file” on page 703](#) provides more information.

Backups of a DXDM file system are performed in the same manner as backups of a standard file system. Backups can be performed on a scheduled basis or manually. Both methods require root permissions. [Chapter 7, “Backup Groups and Schedules”](#) provides more information on scheduling backups. [Chapter 2, “Backing Up Data”](#) provides more information on information on manual backups.

Excluding the DMAPI attributes file

Every DXDM file system has a DMAPI attributes file which stores volatile DMAPI metadata. This file does not need to be backed up because it is re-created during the recovery process. Since these files can become quite large, create a global Directive resource to simplify the exclusion process.

To exclude a DMAPI attributes file:

1. Create a global Directive resource which excludes .DMATTR files. DXDM stores DMAPI attribute information in files located at the top level of each file system. The filename for these files is .DMATTR.
2. When creating the global Directive resource, use the Application Specific Module named Skip to exclude .DMATTR files. [Chapter 9, “Directives”](#) provides more information on directives.
3. Apply the **Directive** resource when creating client/save sets for the file system.

Aborted backups

A DMAPI process is initiated whenever when a backup of a DXDM file system occurs. DXDM assigns a DMAPI session label of DXULDMLIB to this type of process. At the conclusion of a successful backup the process is removed.

If a backup is aborted, the DXULDMLIB process exists as a defunct process. Defunct DXULDMLIB processes exist as process table entries and use no system resources. These processes normally do not cause problems and are removed when the system is rebooted. However, they can be removed manually. [“How to remove a defunct DXULDMLIB process” on page 704](#) provides details. [“Viewing existing DMAPI processes” on page 704](#) provides information on how to determine whether a defunct DXULDMLIB process exists.

Viewing existing DMAPI processes

To view all DMAPI processes on a DXDM host system:

1. Log in as root on the DXDM host system.
2. Create the DXDM environment.

Use the correct command and file for the current shell:

- C shell (csh):

```
source /opt/dxuldm/etc/dxuldm.login
```

where */opt/dxuldm* is the full path to the DXDM installation directory.

- Korn shell (ksh) or Bourne shell (sh):

```
./opt/dxuldm/etc/dxuldm.profile
```

where */opt/dxuldm* is the full path to the DXDM installation directory.

3. Type the **prtdmession** command:

```
prtdmession
```

The **prtdmession** command lists each existing DMAPI session.

How to remove a defunct DXULDMLIB process

NOTICE

Do not remove active DMAPI sessions. If active sessions are removed, DXDM processes terminate and file system activity is blocked until those processes are restarted.

To remove a defunct DXULDMLIB process:

1. Log in as root on the DXDM host system.
2. Create the DXDM environment. [“Viewing existing DMAPI processes” on page 704](#) provides details.
3. Type the **prtdmession** command.
4. Find the listing which has **DXULDMLIB** in the **Session Info** field and note the integer in the listing’s Session ID field.
5. Type the **deldmession** command:

```
deldmession -s DXULDMLIB -i session-id -n
```

where *session-id* is the integer from the defunct DXULDMLIB process’s Session ID field.

Recovery of DXDM file systems

Recovery of a DXDM file system consists of restoring data from a client save set and synchronizing the file system's DMAPI information with the data on the storage target.

Initiating a recovery

Recovery of backed-up data from a DXDM file system is performed by using the same procedures that are used for a standard file system. [Chapter 14, "Recovering Filesystem Data"](#) provides details. DXDM file system recovery can consist of individual files, directories, or a file system.

With DXDM file systems, file data must be recovered into the same file system from which it is backed up. Attempting to recover backed up data into a new file system causes:

- ◆ The unavailability of all data on the storage target.
- ◆ The DXDM processes to stop responding.

This requirement is based on the nature of the DMAPI metadata that DXDM file systems use to locate and access data on the storage target. The DMAPI metadata uses full pathnames. If the file-system mount point changes because of a recovery into a new file system, the pathname information in the metadata becomes inaccurate.

Recovered files list

As part of the recovery process each restored file's pathname is listed in a file on the DXDM host system. DXDM uses this information to synchronize its metadata. ["File system synchronization" on page 706](#) provides details.

When a recovery is complete, all data in the file system is available without waiting for the synchronization process. Data on the storage target becomes available after synchronization.

This data require synchronization:

- ◆ Migrated and purged file data.
- ◆ Purged and deleted file data. ["Restoring deleted files and previous file versions" on page 705](#) provides details.

Restoring deleted files and previous file versions

DXDM file systems retain file versions. This allows administrators to restore specific versions of files. In a DXDM file system, a file is migrated when it is first placed into the file system and again after each change. Each of these migrations cause a new version of the file to be created on the storage target.

To restore a previous version of a file, whether or not the file still exists in the file system:

1. Recover the file from the client/save set that backed up the desired version. ["Recovering the data" on page 373](#) provides information on this procedure.
2. Complete the synchronization process. ["File system synchronization" on page 706](#) provides details.
3. To expedite the synchronization of a particular file, complete the procedure described in ["Manually synchronizing a file" on page 707](#).

The synchronization process restores the data for the version of the file that is recovered. It will also restore data for a file which has been inadvertently deleted. At the conclusion of the synchronization process the recovered file version or deleted file is available.

File system synchronization

As part of a recovery of one or more files into a DXDM file system, the **recover** program creates a file that contains a list of each recovered file.

This list file has a pathname with this format:

```
/opt/dxuldm/adm/recdir/rec.date.pid
```

where:

- ◆ /opt/dxuldm is the installation directory for DXDM.
- ◆ *date* is the date of the recovery.
- ◆ *pid* is the process ID of the NetWorker client process.

DXDM uses the list file to rebuild the file system's DMAPI attributes file. This process synchronizes the file system's metadata with the data on the storage target.

NOTICE

Do not remove a list file created by the **recover** program. Automatic file synchronization will not occur if a list file is removed.

Manual synchronization can be conducted without a list file. [“Manually synchronizing a file” on page 707](#) provides more information. [“Automatic synchronization” on page 706](#) describes that synchronization is an automatic process which does not require administrative intervention.

[“Manually synchronizing a file” on page 707](#) describes how make a file's data available more quickly.

If NetWorker **recover** cannot write to the recdir directory, each recovered file is synchronized as it is recovered. This file-by-file synchronization significantly slows down the recovery process.

The recdir directory is created during the installation of DXDM. NetWorker **recover** will be unable to write to it if the directory is removed or if the partition on which it is mounted becomes full.

Automatic synchronization

DXDM uses the list file created by NetWorker **recover** to determine which files require synchronization. Each file on the list is synchronized. If synchronization does not complete, DXDM retries until all files on the list have been synchronized. After all files are synchronized, the list file is removed.

Automatic synchronization is performed by the script **dxuldmcronscript** which is invoked by root's **crontab** every 10 minutes. The **cron** job to invoke **dxuldmcronscript** is placed in root's **crontab** when DXDM is installed.

Manually synchronizing a file

To make a file available before automatic synchronization is complete, use manual synchronization. This uses the DXDM command-line utility **dxuldmrecover**.

To manually synchronize a file:

1. Log in as root on the DXDM host system.
2. Create the DXDM environment. Step 2 in [“Viewing existing DMAP processes” on page 704](#) provides details.
3. Run the **dxuldmrecover** utility:

```
dxuldmrecover -p filepath
```

where *filepath* is the full path to the file that is being synchronized.

The **dxuldmrecover** utility takes the full path of a file as its argument. It can be invoked from a script to manually synchronize multiple files.

When synchronization is complete this message appears:

```
Successfully recovered file filepath.
```

Note: A DXDM file system can be manually synchronized by using the **dxuldmrecoverfs** utility. The *EMC DiskXtender Data Manager, Release 2.6, Installation and Administrator's Guide* provides more information.

CHAPTER 24

Recovery Support for Windows XP and 2003 Automated System Recovery

This chapter covers these topics:

- ◆ Microsoft Automated System Recovery 710
- ◆ NetWorker support for ASR disaster recovery of Windows XP and 2003 clients 710
- ◆ ASR Limitations and special considerations 711
- ◆ Data and configuration changes since the last backup 712
- ◆ Creating an ASR disk 712
- ◆ Using the ASR disk to recover a NetWorker client 715

Microsoft Automated System Recovery

Microsoft Automated System Recovery (ASR) for Windows XP and Windows Server 2003 enables backup and recovery applications to implement an automated disaster recovery solution.

ASR is similar to the Windows NT Emergency Repair Disk (ERD), but ASR has additional features. Both ERD and ASR require that you prepare a recovery disk in advance. While the ERD requires user interaction to repair selected components of the Windows operating system, ASR provides an automated solution for complete disaster recovery of a failed computer.

ASR is installed as a standard component of the Windows XP and Windows Server 2003 operating systems. No additional Microsoft software is required.

NOTICE

The automated disaster recovery feature for Microsoft Windows Server 2008 and Windows 7 is known as Windows Bare Metal Recovery and is not covered in this chapter. NetWorker support for Windows Bare Metal Recovery is covered in [Chapter 25, “Windows Bare Metal Recovery.”](#)

Microsoft ASR documentation

Microsoft recommends use of ASR as a last resort, after all other system recovery options (such as Safe Mode Boot and Last Known Good) have been exhausted. However, ASR recovery is appropriate in a disaster recovery situation, such as a failure of the system drive.

NetWorker support for ASR disaster recovery of Windows XP and 2003 clients

Beginning with NetWorker 8.0 clients, ASR backup for Windows XP Professional and Windows Server 2003, is no longer supported. However, ASR recovery using pre-NetWorker 8.0 client ASR save sets is still supported, as described in this chapter.

NetWorker ASR save set

The NetWorker ASR save set contains all the information necessary to return the failed computer to its condition at the time of the last ASR backup, including:

- ◆ An automated reinstallation of Windows
- ◆ Restoration of the system configuration
- ◆ Recovery of one or more disk volumes

Network connection names

Microsoft assigns a default name to each client network connection. If possible, do not rename these network connections. If a default network connection name is renamed, one must edit the net.cfg file on the ASR disaster recovery diskette so that the new names are replaced with the default names that were originally assigned by Microsoft. [“Posting ASR disk creation task” on page 714](#) provides details.

ASR Limitations and special considerations

This section describes limitations and special considerations that apply to ASR backup and recovery.

NOTICE

Refer to the Microsoft Knowledge Base article, 818903, for information about supported configurations of Windows XP at the time of backup and recovery.

FAT16 partitions are not supported

Microsoft ASR does not support recovery of disk partitions in FAT16 (also called FAT) format.

To perform an ASR recovery on a computer that has a FAT16 partition:

- ◆ Select the **pause during recovery** option while creating the ASR disk. [“Creating an ASR disk” on page 712](#) provides more information about this option.
- ◆ When the ASR recovery operation pauses, clear the FAT16 partition to exclude it from the recovery. [“Using the ASR disk to recover a NetWorker client” on page 715](#) provides more information about performing ASR recoveries.

After you have completed the ASR recovery and rebooted, recover the FAT16 partition in a separate, non-ASR NetWorker recovery operation. [Chapter 14, “Recovering Filesystem Data”](#) provides information about recovering data.

OEM recovery CDs are not supported

Many computer manufactures, such as Dell, Hewlett Packard, and IBM provide a recovery CD or DVD with each system. These recovery disks typically contain the Windows installation files, plus any additional software included with the system. Although these recovery disks contain a complete set of Windows installation files, they cannot be used to perform an ASR recovery.

NOTICE

To perform an ASR recovery, you must have an official Microsoft Windows installation CD for the version of Windows you are recovering. The *NetWorker Procedure Generator* and the OEM documentation provide information.

Vendor-specific drivers must be installed after Windows installation

ASR recovery can only be done on hardware components supported by the Windows installation media. Drivers for vendor-specific hardware must be installed *after* the Windows installation is complete.

For example, the IBM Thinkpad network interface card (NIC) is not supported by the Windows installation media, and will cause ASR recovery to fail. For more information on hardware supported by the Windows installation media, refer to the Microsoft documentation.

Data and configuration changes since the last backup

When you use the ASR disk to recover a NetWorker client computer, any data and configuration changes made since the last ASR backup will be lost.

Creating an ASR disk

The information in this section refers to using the NetWorker User program on a pre-NetWorker 8.0 client to create an ASR disk. If NetWorker Module for Microsoft Applications is installed on the client computer, ASR is not supported. The *EMC NetWorker Module for Microsoft Applications Administration Guide* provide information about the NetWorker Module for Microsoft Applications program.

Before an ASR recovery, create an ASR disk for the NetWorker client computer that will be recovered. An ASR disk is created by using the NetWorker User program either locally, or as a directed recovery of the ASR save set.

Prerequisites

- ◆ The computer used to create the ASR disk must be running NetWorker release 7.x software.
- ◆ For critical systems, back up the ASR DISK save set frequently and create an ASR disk whenever the client computer drive configuration changes. [“Data and configuration changes since the last backup” on page 712](#) provides more information about backing up an ASR DISK save set.
- ◆ Multiple floppy diskettes will be required to create an ASR disk.

Create an ASR disk locally

If the NetWorker client computer for the new ASR disk is not functional, perform a directed recovery of the ASR save set. [“Creating an ASR disk by using directed recovery” on page 713](#) provides details.

To create an ASR disk locally:

1. Log in with administrator privileges to the NetWorker client computer for which you want to create the ASR disk.
2. In the NetWorker **User** program, click **Recover**.
3. In the **Source Client** dialog box, click **OK** to select the local client.

4. In the **Destination Client** dialog box, click **OK** to select the local client.
5. In the **Recover** window, mark the **ASR: save set for recovery**.
6. Click **Start**. The files for the ASR disk are saved locally in the %temp% directory.
7. When prompted to create an ASR disk for this client, click **Yes**.
8. When prompted, insert the first blank, formatted disk into the **A:** drive and click **OK**.
9. When prompted about pausing the ASR recovery to select which save sets to restore, choose one of these options:
 - **Yes** — During recovery of the NetWorker client host computer, ASR pauses and prompts for the save sets to restore. If the ASR save set was backed up as a component of save set All during a scheduled backup, or as a component of a manual backup, *all* of the client's save sets are available to select. If the ASR save set was backed up individually during a scheduled backup, only those save sets included in the scheduled backup are available for selection.
 - **No** — ASR performs a fully automated recovery of the NetWorker client host computer, without a pause. If the ASR save set was backed up as a component of save set All during a scheduled backup, or as a component of a manual backup, *all* of the client's save sets are restored. If the ASR save set was backed up individually during a scheduled backup, only those save sets included in the scheduled backup are restored.
10. Label each ASR disk after it is created.
11. If the network connections for this client were renamed, complete the procedure described in [“Posting ASR disk creation task” on page 714](#).

Creating an ASR disk by using directed recovery

If the NetWorker client computer is not functional, you can use the directed recovery method to create an ASR disk. The directed recovery method is also useful for performing centralized NetWorker administration, and in cases where you cannot access the %SystemRoot%\Repair\nsr directory on the NetWorker client because the system is damaged.

To create an ASR disk by using directed recovery:

1. Log in with administrator privileges to a NetWorker client computer.
2. In the NetWorker **User** program, click **Recover**.
3. In the **Source Client** dialog box, select the NetWorker client for which you are creating the ASR disk and click **OK**.
4. In the **Destination Client** dialog box, click **OK** to select the local client.
5. In the **Recover** window, mark the **ASR: save set for recovery**.
6. Click **Start**. The files for the ASR disk are saved in the %temp% directory on the destination client (the computer you are using to perform the directed recovery).
7. When prompted to create an ASR disk for this client, click **Yes**.
8. When prompted, insert the first blank, formatted disk into drive A:\ and click **OK**.

9. When prompted about pausing the ASR recovery to select which save sets to restore, choose one of these options:
 - **Yes** — During recovery of the NetWorker client host computer, ASR pauses and prompts for the save sets to restore. If the ASR save set was backed up as a component of save set All during a scheduled backup, or as a component of a manual backup, *all* of the client's save sets are available to select. If the ASR save set was backed up individually during a scheduled backup, only those save sets included in the scheduled backup are available for selection.
 - **No** — ASR performs a fully automated recovery of the NetWorker client host computer, without a pause. If the ASR save set was backed up as a component of save set All during a scheduled backup, or as a component of a manual backup, *all* of the client's save sets are restored. If the ASR save set was backed up individually during a scheduled backup, only those save sets included in the scheduled backup are restored.
10. Label each ASR disk after it is created.

If the network connections for this client were renamed, complete the procedure described in [“Posting ASR disk creation task” on page 714](#).

Posting ASR disk creation task

When an ASR save set is created, client network connections are saved with the current client network connection name. During a disaster recovery, Microsoft ASR requires that the network connections be reconfigured by using the default name assigned by Microsoft. If the client network connections were renamed before the ASR save set was created, update the ASR disk to use the original names assigned by Microsoft.

To update the ASR disk to use the default Microsoft ASR client network connection names:

1. Using a text editor, open the net.cfg file on the ASR disk.
2. Replace all instances of the renamed client network connection with the default Microsoft network connection name.

For example, suppose the Microsoft default network connection name *Local Area Connection* was renamed *Backup Subnet*. In this case, one would replace all instances of the name *Backup Subnet* with *Local Area Connection*.

Microsoft uses these conventions to name client network connections:

- Local Area Connection
- Local Area Connection *x*

where *x* is the second, third, or fourth connection and so on.

3. Close and save the net.cfg file.

Using the ASR disk to recover a NetWorker client

ASR recovery should typically be used only if the failed computer requires a full disaster recovery (for example, if the system drive has failed). For information about other disaster recovery options, refer to the *NetWorker Procedure Generator*.

During an ASR recovery, this occurs:

- ◆ All disk signatures, volumes, and partitions are restored.
- ◆ The Windows operating system is reinstalled.
- ◆ If you selected “pause during recovery” while creating the ASR disk, you are prompted to select which NetWorker backup save sets to recover.
- ◆ The NetWorker software recovers the selected backup save sets.

If you perform an ASR recovery of a Windows XP Professional client that has any type of Microsoft Windows licensing other than Enterprise licensing, you are then prompted to reactivate the Windows license.

Requirements for an ASR recovery

An ASR recovery requires:

- ◆ A current ASR disk for the computer being recovered. [“Creating an ASR disk” on page 712](#) provides more information.

If a current ASR disk or ASR save set is not available for the failed computer, use the legacy NetWorker disaster recovery method. The *NetWorker Procedure Generator* provides details.

- ◆ The Microsoft Windows XP Professional or Windows Server 2003 installation CD for the computer you are recovering.

NOTICE

Recovery CDs provided by computer manufacturer are not supported for ASR recovery. [“OEM recovery CDs are not supported” on page 711](#) provides more information about this restriction.

- ◆ The latest NetWorker backup for the computer you are recovering.

Note: If you routinely move NetWorker backup media to an offsite location for safekeeping, ensure that all necessary volumes are available before you start the recovery. To list the media associated with the files you want to recover, run `mminfo -mv` from the command prompt. The EMC NetWorker Command Reference Guide provides more information about the `mminfo` command.

Performing an ASR recovery

If the steps in this section do not work when recovering a Windows 2003 x64 host, refer to the steps listed in [“Performing a manual recovery on Windows 2003 x64 hosts” on page 717](#).

To perform an ASR recovery on a NetWorker client computer:

1. Start the target computer from the Windows XP Professional or Windows Server 2003 installation CD.

Note: You may need to run the BIOS setup program to configure the computer to boot from the CD-ROM drive. For instructions, refer to the computer manufacturer’s documentation.

2. Watch closely at the beginning of the boot process. If prompted, press a key to boot from the CD-ROM drive.
3. During the text-mode phase of Windows setup, watch the lower portion of the screen. When prompted, press [F2] to display the ASR Recovery menu. Follow the instructions on the screen.
4. When prompted, insert the ASR disk into the **A:** drive and press a key to continue. ASR formats the system partition, copies files, and begins the Windows installation.

NOTICE

Due to a Microsoft Windows 2003 ASR mode problem, when prompted to insert the ASR disk and press a key to continue, you may need to press a key several times before the system recognizes the disk and proceeds with the recovery. This problem does not occur on Windows Server 2003 systems.

If multiple diskettes are required for recovery, the recovery process may not prompt explicitly for the next diskette. Instead, a message may appear that is identical to the message that prompted you to insert the ASR diskette at the beginning of the recovery procedure. In this case, insert the next diskette and press a key to continue.

5. If you did not select the "pause during recovery" option while creating the ASR disk, a fully automated recovery is performed. There is no pause and you are not prompted to select the save sets to recover.

If you selected the “pause during recovery” option while creating the ASR disk, the NetWorker ASR Client dialog box appears during the graphical phase of the Windows installation. Expand **My Computer** to view the save sets to recover, then click **Continue**.

The save sets are marked by default and include these legacy save sets:

- SYSTEM STATE:\
- SYSTEM DB:\
- SYSTEM FILES:\

If VSS is licensed and enabled, these VSS save sets are included:

- VSS SYSTEM BOOT:\
- VSS SYSTEM FILESET:\

VSS USER DATA, VSS OTHER, and VSS SYSTEM SERVICES do not appear because they are not required to boot from ASR mode. [“Creating an ASR disk” on page 712](#) provides more information about the “pause during recovery” option.

By default, the displayed save sets represent the most recent backup. You can view and select previous backups by typing a new browse time in the Browse Time field. The browse time must be entered in time and date the `nsr_getdate` format. For example, a date can be specified by using the format mm/dd/yy or month dd, yy. The *EMC NetWorker Command Reference Guide* provides more information about `nsr_getdate`.

The VSS components and certain components cannot be correctly restored during ASR recovery. [“Components that require special handling after an ASR recovery” on page 717](#) provides information on how to ensure that all necessary components are properly recovered.

Performing a manual recovery on Windows 2003 x64 hosts

For Windows 2003 x64 hosts, if the ASR recovery cannot be performed automatically from the disk as outlined in the section [“Performing an ASR recovery” on page 716](#) use the following steps to manually perform the ASR recovery:

1. Create the ASR disk using the normal method outlined in [“Create an ASR disk locally” on page 590](#).
2. Install the operating system.
3. Create the drives as they were previously.
4. Copy all the ASR disk contents to the %TEMP% folder.
5. Open a command prompt and `cd` to the %TEMP% folder.
6. Type the following in the command prompt:

```
extractlib.cmd
extract.exe asr.sif
```

The **Recover** window now displays. From this window, you can mark the save set for recovery and start the recovery process. This will also configure the network of this system.

Components that require special handling after an ASR recovery

Due to limitations in Microsoft ASR functionality, these system state components cannot be correctly restored during ASR recovery:

- ◆ COM+ Registration Database
- ◆ Disk Quota Database
- ◆ Windows Management Instrumentation Database
- ◆ VSS writers

If the NetWorker client being recovered uses any of these components, perform this procedure after an ASR recovery:

1. Log in with administrator privileges to the target computer.
2. Start the NetWorker **User** program.
3. Click **Recover**.
4. In the **Source Client** dialog box, click **OK** to select the local client.
5. In the **Destination Client** dialog box, click **OK** to select the local client.
6. In the **Recover** window, if VSS is licensed and enabled, mark all VSS save sets for recovery, *except* VSS ASR DISK, and then go to [step 9](#) . If a VSS client license does not exist, or VSS is disabled, go to [step 7](#) .
7. Check for the presence of the COM+ Registration Database component. If it is present, select the **SYSTEM STATE** save set for recovery.
8. In the **Recover** window, select the **SYSTEM DB** save set:
 - a. Check for the presence of these components:
 - Disk Quota Database
 - Windows Management Instrumentation Database
 - b. If either of these components is present, select the SYSTEM DB save set for recovery.
9. If you selected any save sets for recovery, click **Start**.

Verifying the NetWorker client recovery

VSS is unavailable during ASR recovery. Once ASR recovery is complete and the system is rebooted, VSS is available for proper recovery of the writers. [Appendix A, “SYSTEM, VSS SYSTEM, and WINDOWS ROLES AND FEATURES Save Sets”](#) has more information about VSS writers.

[Appendix D, “Additional Features of the Microsoft Windows Server”](#) provides information about how the NetWorker software handles the Windows system state and system-protected files.

To verify the NetWorker client recovery:

1. Reboot the NetWorker client computer and verify that the NetWorker Remote Exec and NetWorker Power Monitor services have started.
2. Use the Windows Event Viewer to check for errors such as:
 - Service startup errors related to the Windows system state.
 - Errors regarding the recovery of Windows system-protected files.
3. Verify that any applications (such as Microsoft Office) that were running prior to the disaster have been properly restored.

To verify:

- a. Run each application.
- b. Open a previously saved document.

CHAPTER 25

Windows Bare Metal Recovery

This chapter covers these topics:

- ◆ Overview of Windows Bare Metal Recovery 720
- ◆ Windows BMR Planning 731
- ◆ Windows BMR Backup..... 747
- ◆ Windows Bare Metal Recovery to Physical or Virtual Computers..... 751

Overview of Windows Bare Metal Recovery

The overview of Windows Bare Metal Recovery (BMR) contains the following topics:

- ◆ [“Changes from previous versions of NetWorker” on page 720](#)
- ◆ [“Supported Operating Systems” on page 721](#)
- ◆ [“BMR Support for Windows Features” on page 721](#)
- ◆ [“Offline recovery versus online recovery” on page 723](#)
- ◆ [“Components of the DISASTER_RECOVERY:\ save set” on page 724](#)
- ◆ [“Full versus incremental backups” on page 727](#)
- ◆ [“Synthetic full backups” on page 728](#)
- ◆ [“Online recovery of Active Directory, DFSR, or Cluster services” on page 729](#)
- ◆ [“Terminology” on page 730](#)

Changes from previous versions of NetWorker

This section describes the changes to NetWorker functionality since the release of NetWorker 8.1 that affect Windows BMR.

Windows Server 2012 Cluster Shared Volumes (CSV)

NetWorker 8.1 SP1 supports backup and recovery of Windows Server 2012 File Servers configured with Windows Continuous Availability using Cluster Shared Volumes (CSV).

The topic “Cluster-aware application” in the Cluster Integration Guide NetWorker Cluster Integration Guide provides more information.

Physical Computer to Virtual Machine (P2V) Recovery

NetWorker now supports BMR recoveries to VMware virtual machines and Microsoft Hyper-V virtual machines. In addition to providing you with a method to transition a physical computer to a virtual equivalent, this process allows you to recover a BMR image to a virtual machine in the event that the physical computer is no longer available. The topic [“To perform a BMR from a Physical Computer to a Virtual Machine \(P2V\)” on page 763](#) provides more information.

NOTICE

The NetWorker software performs P2V BMR operations on a best-effort basis. It is not guaranteed that all possible combinations of physical hardware can be successfully converted to a virtual machine”.

Windows Server 2012 Storage Spaces Support

NetWorker 8.1 SP1 Windows BMR does not support the backup and recovery of critical System State data located on Storage Spaces virtual disks. BMR backup skips all critical volume data located on Storage Spaces and it is not added to the BMR critical volume list. Critical volume data on Storage Spaces cannot be recovered during a BMR recovery.

As long as the Storage Pool disks that compose a Storage Spaces virtual disk are not damaged, a recovery to the original computer includes mounting the Storage Pool virtual disks after recovering the computer's critical volumes.

The topic [“Windows Server 2012 Storage Spaces” on page 722](#) provides more information.

Supported Operating Systems

NetWorker 8.1 SP1 supports Windows BMR for NetWorker clients that run on one of the following operating systems:

- ◆ Windows Server 2008 (x86 and x64) and Windows Server 2008 Core
- ◆ Windows Server 2008 R2 SP1 (x64) and Windows Server 2008 R2 Core
- ◆ Windows Storage Server 2008 R2
- ◆ Windows 7 (x86 and x64)
- ◆ Windows 7 SP1 (x86 and x64)
- ◆ Windows 8 (x86 and x64)
- ◆ Windows 8.1 (x86 and x64)
- ◆ Windows Server 2012 (x64)
- ◆ Windows Server 2012 R2 (x64)
- ◆ Windows Storage Server 2012
- ◆ Windows Storage Server 2012 R2

Windows BMR support with a NetWorker 8.1 SP1 client requires the following versions of NetWorker Software:

- ◆ NetWorker server 7.5 SP3 or later
- ◆ NetWorker Management Console (NMC) 7.6 SP1 or later

NetWorker Windows BMR provides an automated BMR solution by using the Windows ASR writer and other Microsoft VSS writers to identify critical volumes that are needed to perform a recovery on a disabled computer.

Windows BMR is performed offline, while the Windows operating system is inactive. This removes the requirement to reinstall Windows manually and prevents problems that can occur when you restore operating system files to a running version of Windows.

To support Windows BMR, NetWorker provides a bootable Windows BMR image that contains NetWorker binaries and a wizard to control the recovery process.

BMR Support for Windows Features

UEFI Partition Support

NetWorker supports a backup and recovery of unmounted UEFI partitions on computers that run a 64-bit version of a supported Windows operating system. The currently supported 64-bit operating systems include computers that run Windows 7, Windows 8, Windows 8.1, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2.

For more information on how to perform a Bare Metal Recovery of a computer with UEFI partitions, refer to [“To perform a Bare Metal Recovery \(BMR\) to a Physical Computer” on page 751](#)

UEFI partitions are backed up with the following properties:

- UEFI partitions do not need to be mounted before backed up
- UEFI partitions are backed up using the following path pattern:
\\<root>\Device\HarddiskVolume#, # is the number of the volume
- UEFI partitions are backed up as part of DISASTER_RECOVERY
- UEFI partitions are always backed up at level FULL regardless the backup level of DISASTER_RECOVERY
- UEFI partitions are not indexed and is not available for regular online restores

NOTICE

Windows BMR supports file system backup and recovery and 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.

WINDOWS ROLES AND FEATURES

NetWorker supports backup and recovery of Windows Roles and Features.

Windows Roles and Features cannot be restored at the same time as files from the regular volume backup. If you need to recover both Windows Roles and Features and a regular volume backup, restore the volume backups first, and then restore Windows Roles and Features.

Windows Server 2012 Cluster Shared Volumes (CSV)

CSV is not supported as a critical volume. If a CSV disk is marked as a NetWorker critical disk, then the Windows BMR backup posts a warning, and continue as if the CSV is not on the critical list. No backup of the CSV occurs because a CSV cannot be 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 8.1 SP1 Windows BMR does not support the backup and recovery of critical System State data located on Storage Spaces virtual disks. BMR backup skips all critical volume data located on Storage Spaces and it is not added to the BMR critical volume list. Critical volume data on Storage Spaces cannot be recovered during a BMR recovery.

As long as the Storage Pool disks that compose a Storage Spaces virtual disk are not damaged, a recovery to the original computer includes mounting the Storage Pool virtual disks after recovering the computer's critical volumes. There are no additional tasks required to recover Storage Spaces data that reside on the same computer as the original backup.

NOTICE

It is recommended that you detach the physical disks that are used by Storage Spaces during recovery of the critical volumes using Windows BMR and then reattach the physical disks after recovery. Attached Storage Spaces disks may be overwritten during Windows BMR recovery.

As long as a Storage Pool disk is not damaged, a recovery to the original computer automatically mounts any Storage Pool disks after recovering the computer's critical volumes. There are no additional tasks required to recover Storage Spaces data that reside on the same computer as the original backup.

For information on how to perform a Windows BMR recovery of Storage Spaces to a new computer, refer to [“Windows Storage Pools considerations” on page 746](#).

NOTICE

NetWorker allows backup and recovery of data on virtual hard disks and volumes created with Storage Spaces in online mode using NetWorker file system backup and recovery.

A BMR backup of a Windows 2012 host creates a new file named OSSR_sysinfo.xml. The file is located at [root]\EMC NetWorker\nsr\tmp. This file captures pertinent information about the configuration of the host that is backed up, for example:

- ◆ Host information (name, boot drive, BIOS or EFI)
- ◆ NIC cards and their parameters
- ◆ Disk information
- ◆ StorageSpaces information

The purpose of this file is to support the manual recreation of the StorageSpaces configuration following a BMR recovery.

Offline recovery versus online recovery

An offline recovery is an operation that does not require the manual installation of an operating system. Windows BMR is an offline operation. Offline recovery is only supported for backups created on computers that run Windows Server 2008, Windows Server 2008 R2, Windows 7, Windows 7 SP1, Windows 8, Windows 8.1, Windows Server 2012 or Windows Server 2012 R2.

You can recover the Windows System State data to the same or similar hardware only through offline recovery.

You cannot select specific files or save sets to recover during an offline recovery. To select specific files or save sets for recovery, you must perform an online recovery. An online recovery is a NetWorker recovery that is performed when you need to recover only specific files or save sets. [Chapter 14, “Recovering Filesystem Data”](#) provides more information about online recoveries.

Components of the DISASTER_RECOVERY:\ save set

The DISASTER_RECOVERY:\ save set comprises a group of component save sets that are required to provide complete Windows BMR capabilities. A full backup of the DISASTER_RECOVERY:\ save set backs up the following:

- ◆ All critical volumes
- ◆ WINDOWS ROLES AND FEATURES
- ◆ System Reserved partition
- ◆ UEFI partition (if available)

The DISASTER_RECOVERY:\ save set is created during a full backup. A full level backup of the DISASTER_RECOVERY:\ save set is required for a Windows BMR operation.

Each volume in the DISASTER_RECOVERY:\ save set is required for recovery. During a backup of the DISASTER_RECOVERY:\ save set, if a backup of a volume in the save set fails, the entire backup follows the retry settings specified for the save group.

All the save sets that are included in a DISASTER_RECOVERY:\ save set can be displayed in the NMC Administration window in either the Log tab of the Monitoring window or in the Save Set tab of the Media window.

Critical volumes

The following volumes are considered critical and are included in a Windows BMR backup:

- ◆ Any volume that contains operating system files or files required by an installed service.
- ◆ Any volume that has a System State writer file installed.
- ◆ 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. For Windows BMR with a NetWorker version prior to version 8.0, this default behavior is re-configurable. [“Save Operations attribute for Windows BMR backups” on page 738](#) provides more information.
- ◆ If one of the volumes on a dynamic disk is critical, all volumes on a dynamic disk are critical. If one disk in a dynamic disk pack is critical, all disks in that pack must be selected.
- ◆ For Windows Server 2008 and Windows Server 2008 R2, a volume is critical if a Windows service is installed on the volume by an application.
- ◆ For Windows 8, a volume is critical if a Windows service is installed on a volume by an application. However, that volume can be configured to not be critical if the HKLM\SYSTEM\CurrentControlSet\Control\SystemWriter\ReportWin32ServicesNonSystemState is set to 1.
- ◆ A volume can be configured to be critical if the registry key HKLM\SYSTEM\CurrentControlSet\Control\SystemWriter\ReportWin32ServicesNonSystemState is set to 0.
- ◆ If a mount point is critical, then the parent volume is also critical.
- ◆ For Windows BMR with a NetWorker version prior to version 8.0, critical volume behavior is re-configurable. [“Save Operations attribute for Windows BMR backups” on page 738](#) 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 are backed up by an application VSS writer, such as Exchange databases. These files must be backed up with an application backup program such as NetWorker Module for Microsoft Applications (NMM).

Table 116 on page 726 provides examples of how critical and non-critical data is handled in backup and recovery operations associated with the DISASTER_RECOVERY:\ save set.

Table 116 Critical and non-critical data in backup and recovery

Volume type	DISASTER_RECOVERY:\ backup examples	DISASTER_RECOVERY:\ recovery examples (1 of 2)
Critical	<p>BCD (Boot Configuration Data)</p> <p>BCD is included in the System State data. The partition is saved based on the Unique Volume GUID.</p> <p>C:\</p> <p>All System State data.</p> <p>The data is not backed up by an application VSS writer.</p>	<p>The System State files for BCD and C:\ are recovered along with any files <i>not</i> associated with an application VSS writer. Files associated with an application VSS writer are not recovered.</p>
Critical	<p>[drive letter]:\</p> <p>Any volume that has a System State file installed on it is critical.</p> <p>Windows Server 2008 or Windows Server 2008 R2 - a volume is critical if an application installs a Windows service on the volume.</p> <p>Windows Server 2012, Windows Server 2012 R2, Windows 8 and Windows 8.1 - a volume is critical if an application installs a Windows services on the volume and if the registry key HKLM\SYSTEM\CurrentControlSet\Control\SystemWriter\ReportWin32ServicesNonSystemState is set to 0.</p> <hr/> <p>Notice: A Windows Server 2012 or Windows 8 volume that has a Windows application service installed is not critical if the registry key HKLM\SYSTEM\CurrentControlSet\Control\SystemWriter\ReportWin32ServicesNonSystemState is set to 1.</p> <p>The default setting for Windows Server 2012 is 1, or not critical.</p> <p>The default setting for Windows 8 is 0, or critical.</p> <p>Refer to Windows Article ID 2792088 for additional information.</p> <hr/> <p>If a mount point is a critical volume, then the parent volume is also a critical volume.</p> <p>If one of the volumes on a dynamic disk is critical, all volumes of the dynamic disk are critical.</p>	<p>Files associated with installed services (System State) are recovered along with any files <i>not</i> associated with an application VSS writer.</p>

Table 116 Critical and non-critical data in backup and recovery

Volume type	DISASTER_RECOVERY:\ backup examples	DISASTER_RECOVERY:\ recovery examples (2 of 2)
Critical	<p>Notice: Files that are associated with the application VSS writer are not backed up as part of the DISASTER_RECOVERY:\ save set and cannot be recovered unless they are backed up by an application backup program, such as NetWorker User for SQL Server, NetWorker Module for Microsoft Exchange Server (NME), or NMM. These application backup programs are designed to back up applications, such as SQL Server, Exchange, and others.</p>	
Non-critical	No data is backed as part of the DISASTER_RECOVERY:\ backup.	No data is recovered as part of a Windows BMR operation.

Listing the critical volumes

You can list the critical volumes for a NetWorker client by running 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 is displayed:

```
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)
```

Full versus incremental backups

The DISASTER_RECOVERY:\ save set is created during full backups when either the All or the DISASTER_RECOVERY:\ save set is specified in the NetWorker client resource.

[Table 117 on page 727](#) shows how the components of a DISASTER_RECOVERY:\ save set are handled during an incremental backup.

Table 117 DISASTER_RECOVERY:\ components in an incremental backup

Save set	Description
DISASTER_RECOVERY:\	An incremental DISASTER_RECOVERY:\ save set is created. A Bare Metal Recovery (BMR) can be performed from these daily incremental backups.
CRITICAL VOLUME save set	Saved at level incremental
WINDOWS ROLES AND FEATURES save set	Saved at level full*
* This behavior is configurable. “Windows BMR limitations and considerations” on page 739 provides more information.	

During incremental backups, all specified non-critical save sets are backed up at the incremental level.

During an incremental backup, the NetWorker client checks both the modification time and the archive bit in determining whether a file needs to be backed up. The archive bit is ignored if the environment variable `Nsr_avoid_archive` is set, resulting in a backup of the file at every incremental backup.

Use the environment variable `Nsr_avoid_archive` with caution. If you use the environment variable `Nsr_avoid_archive`, test your BMR backup image to ensure that you can recover your Windows system state correctly. [“Perform a NetWorker Bare Metal Recovery wizard test before recovery” on page 754](#) provides more information on testing your BMR backup image.

Note: When the save set All is backed up for the first time on a NetWorker client before a `DISASTER_RECOVERY:\` save set has been saved, the save level is forced to Full so that a `DISASTER_RECOVERY:\` backup is assured. This behavior is also enforced when the save set All is backed up for the first time after upgrading a client to NetWorker 8.0 or later.

Synthetic full backups

The synthetic full backup feature is introduced in NetWorker 8.0. Synthetic full backups use the most recent full and incremental backups to create a full backup without transferring any data from the client. All of the work to synthesize a full backup is performed on the backup 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.

For a full description of the synthetic full backup feature, see [Chapter 2 “Synthetic full backups” on page 76](#).

When the `DISASTER_RECOVERY:\` save set is included in a client backup, volumes that are identified as critical are always backed up at the full level. No synthetic full backup is created for these 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 63 Synthetic full backups

The save set All is specified in the client resource and a synthetic full backup is scheduled on Sunday. 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.

For the synthetic full backup on Sunday, the following save sets are created in addition to a level full `DISASTER_RECOVERY:\` save set and the `WINDOWS ROLES AND FEATURES` save set:

- ◆ C:\ — A true level full backup is created.
- ◆ E:\ — A true level full backup is created.
- ◆ F:\ — A synthetic full backup is created.
- ◆ G:\ — A synthetic full backup is created.

If the DISASTER_RECOVERY:\ save set is not included in a backup, then all four volumes are backed up as synthetic fulls on Sunday. In this example, if you specify the four volumes in the client resource without specifying the save set All or the DISASTER_RECOVERY:\ save set, then a synthetic full backup is created for all four volumes.

You can change the default behavior of the ALL save set so that a DISASTER_RECOVERY:\ save set is not automatically created. For more information on how to change the default behavior of the ALL save set, see [“Blending Windows BMR recovery backups with synthetic full backups” on page 739](#).

In this example, when you change the default behavior of the ALL save set, the critical volumes, C:\ and E:\, are backed up as synthetic full. However, when the DISASTER_RECOVERY:\ save set is explicitly specified in the client resource, the critical volumes are still backed up as true full backups, regardless of whether a synthetic full backup is specified.

Online recovery of Active Directory, DFSR, or Cluster services

NetWorker 8.1 and higher support the online recovery of the following Windows services:

- ◆ Active Directory
- ◆ Distributed File System Replication (DFSR)
- ◆ Cluster

The WINDOWS ROLES AND FEATURES save sets are included in a Windows BMR back up and are also recovered in a regular online recovery operation.

NOTICE

Online recovery of WINDOWS ROLES AND FEATURES save sets is not supported except for the Windows services described in this section. For any other Windows server service, online recovery of WINDOWS ROLES AND FEATURES save sets is supported only if the same Windows operating system instance is used. The incorrect online recovery of WINDOWS ROLES AND FEATURES save sets leads to an inconsistent state of the Windows server.

[Table 118 on page 729](#) provides references to additional information for Windows services supported with online recovery.

Table 118 Additional information for Windows services supported with online recovery

Service	Reference
Active Directory	See <i>NetWorker Procedure Generator</i> .
DFSR	Appendix C, “Backing Up and Restoring a Microsoft DFS”
Cluster	See <i>NetWorker Procedure Generator</i> .

Terminology

This chapter uses the following terms to describe NetWorker support for Windows BMR technology:

- ◆ Application data — User data that is created by an application, such as log files or a database. For example, the application data of a SQL server includes databases and log files. You cannot recover application data by using a Windows BMR operation. You must back up and recover application data with a NetWorker module, such as the NetWorker Module for Microsoft (NMM).
- ◆ ASR writer — The Volume Shadow Copy Service (VSS) writer that identifies critical data needed to perform an offline recovery.
- ◆ Bare Metal Recovery (BMR) — The restoration of a computer operating system and its data 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.
- ◆ Boot Configuration Data (BCD) — A data store that contains a description of boot applications and boot application settings to start Microsoft Windows 7 and Microsoft Windows Server 2008 operating systems. You require a backup of this ASR writer component to perform an offline recovery.
- ◆ Critical volume — One of the following:
 - Any volume that contains System State data files or 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 where a critical volume is mounted.

NOTICE

All volumes on all dynamic disks are considered critical if at least one of the volumes is critical.

You require a current backup of all critical volumes to perform a Windows BMR.

- ◆ Recovery — The restoration of a computer operating system and its data after a catastrophic failure, such as a hard disk failure or the corruption of critical operating system components. A recovery might or might not be a Windows BMR.
- ◆ NetWorker Windows BMR image — A bootable image that contains NetWorker binaries and a wizard to control the Windows BMR process.
- ◆ Non-critical volume — A volume that contains files that are not part of System State data or an installed service.
- ◆ Offline recovery — A restore operation performed from the NetWorker Windows BMR boot image. An offline recovery is an automated process that does not require the manual installation of an operating system. A BMR is an offline recovery.
- ◆ Online recovery — A restore operation performed from the regular NetWorker Recover user interface. An online recovery requires that the computer has been booted from an installed operating system.

- ◆ 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 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 data is not backed up or recovered with Windows BMR unless it 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. User data can be recovered online at any time (on demand) or after a Windows BMR operation.
- ◆ Windows Bare Metal Recovery — A BMR of a host, also known as Windows BMR. NetWorker provides an automated BMR solution for Windows.
- ◆ 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.

Windows BMR Planning

This section provides guidelines on how to plan your Windows BMR backups.

Road map for Windows BMR Planning

[Table 119 on page 731](#) provides a roadmap of the typical Windows BMR backup and recovery process. The roadmap indicates which steps need to be performed before you attempt a Windows BMR.

Table 119 Windows BMR backup and recovery roadmap

Action	For more information	Verify before attempting a recovery? (1 of 2)
Plan your backup	“Changes from previous versions of NetWorker” on page 720	Not required
	“Hardware Requirements for Windows BMR Backup and Restore” on page 732	Yes
	“Configuration requirements for Windows BMR backups” on page 733	Yes
	“Save set planning” on page 734	Yes
	“Best Practices for Windows BMR” on page 736	Not required
Backup - scheduled	“Include Windows BMR in scheduled backups” on page 748	Yes
Backup - manual	“Include Windows BMR in manual backups” on page 749	Yes

Table 119 Windows BMR backup and recovery roadmap

Action	For more information	Verify before attempting a recovery? (2 of 2)
Verify backup	“How to verify a valid Windows BMR backup” on page 750	Yes
Test Recovery Process	“Windows Bare Metal Recovery to Physical or Virtual Computers” on page 751	Yes
Recovery	“To perform a Bare Metal Recovery (BMR) to a Physical Computer” on page 751	-na-
Post-recovery	“Troubleshooting Windows BMR” on page 765	-na-
Troubleshooting	“Troubleshooting Windows BMR” on page 765	-na-
Additional Options	“Additional recovery options” on page 769	-na-

Hardware Requirements for Windows BMR Backup and Restore

The BMR recovery process completes a bare metal recovery of the operating system used on the server that was backed up. If the new server has hardware that requires new or additional drivers, after the computer is recovered and rebooted following a BMR recovery procedure, Windows prompts you to install the required drivers for the new hardware.

A BMR 32-bit ISO image can be used to recover an x86 operating system on an x86 or x64 computer. A Windows BMR 64-bit ISO image can only be used to recover a 64-bit operating system running on a 64-bit computer.

For the purposes of BMR, AMD and Intel processors can be treated as equivalent if they follow the same architecture. The backup of AMD x64 computer's OS can be recovered on to an Intel x64 computer and the opposite is true as well.

When performing a recovery to a startup disk, the recovery process restore the backup to the same logical disk number as on the original server. You cannot change to another hard disk to restore the operating system. Windows BMR supports IDE, SATA, or SCSI hard disks. You can make the backup on one kind of hard disk and recover on another kind of hard disk. For example, SAS to SATA is supported.

You should ensure that your RAID setup on the destination computer does not interfere with the disk order of the hard disks.

The replacement hardware used for a BMR recovery process must meet the following requirements:

- The operating system architecture and processor architecture must match.
- The hardware on the host to be recovered (target host) is operational.
- The target host requires a minimum of 512 MB of RAM.
- The startup hard disk capacity should be larger or the same size. If the disk is smaller by a single byte, BMR fails.

- There are at least as many disks on the target host as 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 computer can not interfere with the disk order of the hard disks.
- The disk or RAID drivers used in the old system must be compatible with the disk or RAID controllers in the new system.
- The recovery boot image must be available as a bootable CD volume or from the network boot location.

After a Windows BMR recovery, you must install NIC drivers that match the NIC in the target computer after the computer reboots. All NIC or storage device drivers must not require a reboot to complete the driver installation process.

Configuration requirements for Windows BMR backups

Ensure that you have the following before you start the Windows Bare Metal Recovery wizard:

- ◆ Network or disk drivers that might need to be installed.
- ◆ The network name and IP address of the host to be recovered.
- ◆ The network name and IP address of the NetWorker server that are used to recover the host.
- ◆ The network name and IP address of the NetWorker storage node if it does not reside on the NetWorker server host.
- ◆ If a DNS server is being used to resolve IP addresses, the default gateway and the name of the DNS server:
 - If a DNS server is not available, you can use a local hosts file to resolve the NetWorker server name and its IP address.
 - If the NetWorker storage node does not reside on the NetWorker server host, ensure that its hostname and IP address is also added to the hosts file.
 - The NetWorker media volumes that are required for the recovery.
- ◆ If you are backing up Active Directory, DFSR or Cluster Services, ensure that WINDOWS ROLES AND FEATURES is backed up.

Save set planning

This section describes the attributes of save sets. This information helps you select the correct save set configuration for your computer and operating system.

ALL save set definition

The definition of the NetWorker save set All depends on the Windows operating system and the NetWorker version. [Table 120 on page 734](#) lists the components of the ALL save set for different Windows operating systems and NetWorker versions.

Table 120 Components of the ALL save set

Components of the ALL save set for: Windows Server 2003 with VSS enabled*	Components of the ALL save set for: Windows Server 2008, Windows Server 2008 R2, Windows 7, Windows Server 2012, Windows Server 2012 R2, Windows 8.0 Windows 8.1 in NetWorker 7.6 SP2* and SP3*	Components of the ALL save set for: Windows Server 2008, Windows Server 2008 R2, Windows 7, Windows Server 2012, Windows Server 2012 R2, Windows 8.0 Windows 8.1 in NetWorker 8.0* and NetWorker 8.0 SP1	Components of the ALL save set for: Windows Server 2008, Windows Server 2008 R2, Windows Storage Server 2008 R2 Windows 7, Windows Server 2012, Windows Server 2012 R2, Windows 8.0 Windows 8.1 Windows Storage Server 2012 in NetWorker 8.1 Windows Storage Server 2012 R2 in NetWorker 8.1 Windows Storage Server 2012 in NetWorker 8.1 SP1 Windows Storage Server 2012 R2 in NetWorker 8.1 SP1
VSS SYSTEM BOOT:\ VSS SYSTEM FILESET:\ VSS SYSTEM SERVICES:\ VSS USER DATA:\ VSS OTHER:\ VSS ASR DISK:\ (Windows Server 2003 only) All local physical drives	DISASTER_RECOVERY:\ (included in full backup only) All local physical drives	DISASTER_RECOVERY:\ (included in full backup only) VSS SYSTEM BOOT:\ VSS SYSTEM FILESET:\ VSS SYSTEM SERVICES:\ All local physical drives	DISASTER_RECOVERY:\ WINDOWS ROLES AND FEATURES All local physical drives SRP or EFI, as applicable

In addition to Windows BMR capabilities, regular file system backup and recovery is provided by specifying the save set All.

WINDOWS ROLES AND FEATURES save set definition

WINDOWS ROLES AND FEATURES save sets are backed up as part of the DISASTER_RECOVERY:\ save set backup.

These save sets are recovered as part of a Windows BMR operation and are also available for online recovery. However, the incorrect online recovery of WINDOWS ROLES AND FEATURES save sets could lead to an inconsistent state of the Windows system. Therefore, do not recover WINDOWS ROLES AND FEATURES save sets online except in special cases

such as the recovery of specific writers for Active Directory, DFSR, or Windows Server Failover Cluster. The topic [“Online recovery of Active Directory, DFSR, or Cluster services”](#) on page 729 provides more information.

NOTICE

The NetWorker 8.1 SP1 client can only recover WINDOWS ROLES AND FEATURES save sets. If you attempt to recover a VSS System State save set that was created with a NetWorker 8.0 SP1 client or earlier, then NetWorker 8.1 SP1 does not function correctly.

Restoring VSS System State savesets from NetWorker 8.0 SP1 or earlier should not be performed with a NetWorker 8.1 SP1 client. Instead to restore the system state it is recommended to restore the WINDOWS ROLES AND FEATURES saveset from a NetWorker 8.1 or later backup.

Save set configuration by host type

[Table 121 on page 735](#) lists the save sets to back up, depending on the Windows host to be protected.

Table 121 Save set configuration for a specific host

To back up this host	Specify these save sets in the client resource Save Set attribute	Considerations (1 of 2)
a host or file server with one of the following operating systems: Windows Server 2008 Windows Server 2008 R2 Windows Storage Server 2008 R2 Windows 7 Windows Server 2012 Windows Server 2012 R2 Windows Storage Server 2012 Windows Storage Server 2012 R2 Windows 8 Windows 8.1	<ul style="list-style-type: none"> Specify the save set All in the NetWorker client resource. By default, the save set All includes the DISASTER_RECOVERY:\ save set and all of the local physical drives. 	<ul style="list-style-type: none"> WINDOWS ROLES AND FEATURES must be backed up. WINDOWS ROLES AND FEATURES save sets are recovered in a Windows BMR operation and are also available for online recovery. WINDOWS ROLES AND FEATURES save sets should only be recovered online as part of an Active Directory, DFSR, or Windows Server Failover Cluster online recovery.

Table 121 Save set configuration for a specific host

To back up this host	Specify these save sets in the client resource Save Set attribute	Considerations (2 of 2)
a host or file server with one of the following operating systems: Windows Server 2003	<ul style="list-style-type: none"> Specify the save set All in the NetWorker client resource. By default, the save set All includes the DISASTER_RECOVERY:\ save set and all of the local physical drives. 	
A host with server roles that use a SQL Server database and one of the following operating systems: Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2	<ul style="list-style-type: none"> Specify the ALL save set in the NetWorker client resource. Use an application backup product such as NMM or NetWorker User for SQL Server to back up the SQL Server databases. The application backup product documentation provides details. 	<ul style="list-style-type: none"> Recover the host's System State with NetWorker Windows BMR. Recover the application databases with NMM or NetWorker User for SQL Server.
A host with Microsoft server applications, such as a Microsoft Exchange Server, Microsoft SQL Server, Hyper-V, or Microsoft Office SharePoint Service and one of the following operating systems: Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2	<ul style="list-style-type: none"> Specify the ALL save set in the NetWorker client resource. Use an application backup product such as NMM or NetWorker User for SQL Server to back up the application databases. The application backup product documentation provides details. 	<ul style="list-style-type: none"> Recover the host's System State with NetWorker Windows BMR. Recover the application databases with NMM, NetWorker User for SQL Server, or a third-party application.

Best Practices for Windows BMR

The following sections outline best practices for Windows BMR.

- ◆ [“Performing regular backups” on page 736](#)
- ◆ [“Windows BMR recovery backups” on page 737](#)
- ◆ [“Cloning Windows BMR recovery save sets” on page 737](#)
- ◆ [“Capturing disk configuration changes for Windows BMR” on page 737](#)
- ◆ [“Mixing critical and non-critical volumes on a physical disk” on page 737](#)
- ◆ [“Considerations for NetWorker user defined directives” on page 737](#)
- ◆ [“Optimizing Windows BMR Recovery backups” on page 738](#)

Performing regular backups

Ensure that a full backup that contains the DISASTER_RECOVERY:\ save set is performed regularly and after any significant change to the Windows operating system. The DISASTER_RECOVERY:\ save set is specified automatically when the ALL save set is specified in the Save Set attribute of the NetWorker client resource or when DISASTER_RECOVERY:\ is specified explicitly in the Save Set attribute.

Windows BMR recovery backups

Windows BMR recovery backups should be taken after any system components are installed, removed, or updated. This includes adding, changing, or removing Windows roles and features, or installing Windows updates and service packs.

Cloning Windows BMR recovery save sets

If you have set up a backup group to automatically clone Windows BMR recovery save sets, ensure that NetWorker server and client host clocks are synchronized across the network. Otherwise, some of the save sets might not be cloned.

Capturing disk configuration changes for Windows BMR

Windows BMR uses the Microsoft ASR writer to reconstruct a disk configuration during the recovery. The ASR writer is sensitive to the disk numbers and disk configuration on the original host. This disk information is saved during a Windows BMR backup and is used when the disk configuration is recovered. After any disk reconfiguration on the protected host, reboot and do a Windows BMR backup to ensure that the new disk configuration is captured. Examples of a disk reconfiguration are the adding or removing of a disk or partition on the protected host.

Mixing critical and non-critical volumes on a physical disk

Windows allows a physical disk to be partitioned into multiple volumes. These volumes can be either critical or non-critical, depending on the data they contain. During a recovery, the ASR writer can re-create partitions and perform formatting of those partitions. In some cases, such as a disk replacement scenario, this could include non-critical partitions. If this happens, data on the non-critical partitions must be recovered with an online recovery. [“Post-recovery tasks to recover file system data” on page 761](#) provides more information on 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

Employ user defined directives, such as `nsr.dir`, with caution. Using such directives in directories where system state files reside will lead to an incomplete BMR backup image and potentially render your BMR backup image unusable. If you create user defined directives, test your BMR backup image to ensure that you can recover your Windows system state correctly. [“Perform a NetWorker Bare Metal Recovery wizard test before recovery” on page 754](#) provides more information about testing your BMR backup image.

Optimizing Windows BMR Recovery backups

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.

Save Operations attribute for Windows BMR backups

[Table 122 on page 738](#) summarizes the values of the Save Operations attribute you can specify for Windows BMR backups.

Specify these attributes in the Save Operations attribute of the client resource to reduce the size and improve the speed of recovery backups.

NOTICE

When you use the attribute values in [Table 122 on page 738](#), there is a potential that some backed-up data might not be available for recovery. To prevent this situation, follow the important notice in the table.

Table 122 Save Operations attribute for Window BMR

Objective	Value specified in Save Operations attribute	Result
Do not embed the DISASTER_RECOVERY:\ save set within the ALL save set.	VSS:DISASTER_RECOVERY=off Default value: No value, which means that the ALL save set includes the DISASTER_RECOVERY:\ save set.	The DISASTER_RECOVERY:\ save set is not included in the ALL save set. The save set All consists of the following save sets (consistent with pre-7.6 SP2 NetWorker): <ul style="list-style-type: none"> • WINDOWS ROLES AND FEATURES • All local physical drives
	Notice: If you specify this attribute value, create a second client resource for the NetWorker client host and set up a scheduled back up of the DISASTER_RECOVERY:\ save set for the second client resource. This ensures that you can perform a Windows BMR for the host. “Blending Windows BMR recovery backups with synthetic full backups” on page 739 provides an example.	
Disable the backup of non-critical volumes that have mount points for critical volumes.	VSS:SAVE_NONCRITICAL_MNTPARENT_DISKS=off Default: No value, which means that a non-critical parent volume that has a critical volume mounted to it will be saved with the DISASTER_RECOVERY:\ save set.	If at least one critical volume is mounted on a non-critical volume, do not back up the non-critical volume with the DISASTER_RECOVERY:\ save set.
	If this option is used, ensure that all non-critical volumes are backed up. Otherwise, user data may not be available for an online recovery. You can back up all non-critical volumes by specifying the save set All in the client resource’s Save Set attribute.	

Blending Windows BMR recovery backups with synthetic full backups

When the DISASTER_RECOVERY:\ save set is included in a client backup, volumes that are identified as critical are always backed up at the full level. No synthetic full backup is created for these critical volumes. The DISASTER_RECOVERY:\ save set is included in full backups when either the All or DISASTER_RECOVERY:\ save set is specified in the NetWorker client resource.

Example 64 Performing Windows BMR recovery backups less frequently than regular backups

Performing a Windows BMR recovery backup every two weeks might be adequate. However, you might want to take advantage of the synthetic full feature to back up the data, including critical volumes and Windows Roles and Features data, more frequently.

To set up different schedules for your Windows BMR recovery backups and regular backups, do the following:

1. Set up two NetWorker client resources for the host that you back up.
2. In the first client resource, set the following attributes:
 - Save Operations — **VSS:DISASTER_RECOVERY=off**
 - Save Set — **ALL**
 - Schedule — Select a schedule that performs a synthetic full backup weekly.
 - Group — Select a different group than for the second client resource. The following data undergoes a synthetic full backup every week:
 - WINDOWS ROLES AND FEATURES
 - All local physical drives including critical drives
3. In the second client resource, set the following attributes:
 - Save Operations — No value
 - Save Set — **DISASTER_RECOVERY:**
 - Schedule — Select a schedule that performs a full backup every two weeks.
 - Group — Select a different group than for the first client resource. The DISASTER_RECOVERY:\ save set is backed up every two weeks.

Windows BMR limitations and considerations

This section includes the following Windows BMR limitations and special considerations to be aware of before performing Windows BMR recovery backup and recovery operations.

- ◆ [“Disk configuration considerations” on page 740](#)
- ◆ [“Optimized deduplication considerations” on page 740](#)
- ◆ [“Save set considerations” on page 741](#)
- ◆ [“Security considerations” on page 741](#)
- ◆ [“Server role considerations” on page 742](#)
- ◆ [“Windows Server application considerations” on page 743](#)
- ◆ [“Windows services considerations” on page 746](#)

- ◆ [“Windows Storage Pools considerations” on page 746](#)
- ◆ [“WinPE configuration for SAN boot devices” on page 747](#)
- ◆ [“VMware considerations” on page 747](#)

Disk configuration considerations

This sections describes disk configurations limitations in Windows BMR.

- ◆ [“Dynamic Disks” on page 740](#)
- ◆ [“Only NTFS and ReFS file systems are recognized as critical volumes” on page 740](#)

Dynamic Disks

Following an OSSR BMR procedure and a WINDOWS ROLES AND FEATURES save set recovery, Dynamic disk volumes are not online. After you perform the BMR procedure and WINDOWS ROLES AND FEATURES recoveries, bring the dynamic disks back online using Windows Disk Manager.

Only NTFS and ReFS file systems are recognized as critical volumes

Windows BMR 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 with the following error message, which is logged in the savegrp.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.
```

Although the backup of the DISASTER_RECOVERY:\ save set fails, the contents of the partition are backed up and are available for online recovery only.

To ensure that the DISASTER_RECOVERY:\ save set is backed up properly, find the service or application that is installed on the volume, remove it, and reinstall it on an NTFS volume.

Windows BMR does not support FAT and FAT32 file systems as critical volumes.

Optimized deduplication considerations

NetWorker supports a full volume saveset restore to the original volume for a Windows Server configured for optimized deduplication created with a level Full saveset.

The following restores are not supported for Windows servers configured for or optimized deduplication:

- Restore of a list of files from a level FULL or INCREMENTAL save set.
- Full volume restore of a non-full level saveset.

If an optimized deduplication backup is recovered to a deduplicated volume and the recovery aborts or fails for some reason, the volume is left in an unusable state. You must repeat the recovery process and the recovery must complete successfully or the volume is corrupt.

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 section describes limitations and considerations related to save sets.

- ◆ [“Checkpoint restart backup for Windows DISASTER_RECOVERY:\ save set is not supported” on page 741](#)
- ◆ [“Including DISASTER_RECOVERY:\ in multiple save sets” on page 741](#)
- ◆ [“Full backups are required for Windows BMR” on page 741](#)

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. If a client with a DISASTER_RECOVERY:\ save set is enabled for checkpoint restart, the backup fails.

[“Save sets” on page 66](#) provides more information.

Including DISASTER_RECOVERY:\ in multiple save sets

If multiple save sets specified in the command line interface include the DISASTER_RECOVERY:\ save set, list the save sets in this order:

```
save.exe -s server -N "DISASTER_RECOVERY:\"" save_set1 save_set2 ...
"DISASTER_RECOVERY:\""
```

where *save_set1* or *save_set2* is a save set name, such as a drive letter (f:\) or mount point (n:\mountpoint) and DISASTER_RECOVERY:\ must be first and last save set specified.

Full backups are required for Windows BMR

The DISASTER_RECOVERY:\ save set is backed up only during a full backup. The DISASTER_RECOVERY:\ save set is not backed up during an incremental backup.

Monitoring save operations

When monitoring Windows BMR save operations, for example, when viewing the NetWorker **Administration** > **Monitoring** > **Sessions** window, you might notice that the number of save sessions is different than the number of save sets listed in the client resource. This is because Windows BMR backups are optimized to generate the correct number of Windows BMR backup sessions and save sets.

Security considerations

This section describes security issues related to Windows BMR planning.

NetWorker Strong Authentication and Windows BMR

Using NetWorker strong authentication (nsrauth) for a Windows client can result in extra steps. The server authentication requirements must be relaxed for the WinPE ISO so that the server does not refuse the recovery.

NetWorker strong authentication uses the Secure Sockets Layer protocol. [“NetWorker authentication” on page 612](#) provides information on how to use and relax authentication. If nsrauth is exclusively used for a NetWorker client, then during Windows

BMR, the WinPE client image does not have the nsrauth credentials file that was on the original client. This causes communication with the NetWorker server to fail. To address this issue, do one of the following:

- ◆ Follow the procedures for creating new credentials for the WinPE client. Treat the WinPE system like a NetWorker client with corrupted credentials. After the recovery, and after the recovered system has rebooted, repeat the steps for creating new credentials for this client. Repeating the steps is required because, although the credentials file for this client was recovered, the credentials are now old and do not match the credentials on the server.
- ◆ Opt out of nsrauth for the WinPE client during the WinPE recovery process. Use oldauth until the recovery is complete. After the recovery is complete and the system is rebooted, opt back into nsrauth for this client. Opting back into nsrauth creates a new set of credentials.

Server role considerations

This section describes considerations for Windows Server Roles in Windows BMR.

- ◆ [“Protecting Windows server roles” on page 742](#)
- ◆ [“Backup and recovery workflow for server roles that use WID” on page 743](#)
- ◆ [“Backup and recovery workflow for server roles that use SQL Server” on page 743](#)

Protecting Windows server roles

Several server role components of Windows Server 2008 and Windows Server 2008 R2 need a database for their data storage. Examples of Windows server roles with databases include:

- ◆ Active Directory Rights Management Services (ADRMS)
- ◆ Windows System Resource Manager (WSRM)
- ◆ Universal Description, Discovery, and Integrations (UDDI) Services
- ◆ Windows Server Update Services (WSUS)

When these Windows server roles are installed, you can select either an existing SQL Server installation or the Windows Internal Database (WID). WID is a variant of SQL Server Express 2005 and is a feature component of Windows Server 2008. The VSS SQL Server writer is used to protect the WID databases.

NetWorker Windows BMR protects role databases stored in WID but does not protect role databases stored in a SQL Server outside of WID.

If System State data, such as roles, are stored in a SQL Server database outside of WID, then the data must be protected by using NMM, NetWorker User for SQL Server, or a third-party SQL backup product. SQL Server system databases (master, model, and msdb) are not recovered as part of NetWorker Windows BMR. To perform a SQL Server recovery, you must:

1. Perform a SQL Server recovery to rebuild the SQL system databases.
2. Use the SQL application backup tool, such as NMM, to recover all the SQL databases, as required.

Backup and recovery workflow for server roles that use WID

To use NetWorker Windows BMR to save and recover the WID database:

1. Perform a NetWorker Windows BMR backup. All SQL writer components for WID are included in the backup.
2. Perform a NetWorker Windows BMR operation. All WID components are recovered.

After the NetWorker Windows BMR system reboot, the WID service is available and Windows server roles have access to their databases.

Backup and recovery workflow for server roles that use SQL Server

To use NMM, NetWorker User for SQL Server, or a third-party application to recover SQL Server components:

1. Protect the host System State and user data with the NetWorker Windows client.
2. Protect the SQL Server application with NMM, NetWorker User for SQL Server, or a third-party application.
3. Perform a NetWorker Windows BMR operation.

After the recovery reboot, the system is recovered but the SQL Server service does not run. Any roles with databases in the SQL Server do not work.

4. On node A, rebuild the SQL server by running the following Setup command. 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:

```
C:\> Setup /QUIET /ACTION=REBUILDDATABASE
/INSTANCENAME=Instance_name
/SQLSYSADMINACCOUNTS=domain_name\administrator
```

The following Microsoft URL provides more information:

[http://msdn.microsoft.com/en-us/library/ms189302\(SQL.105\).aspx](http://msdn.microsoft.com/en-us/library/ms189302(SQL.105).aspx)

5. Bring the SQL sever services online.
6. Recover the SQL system databases (master, model, msdb) with NetWorker User for SQL Server, or a third-party application.
7. Recover the role databases with NetWorker User for SQL Server, or a third-party application.
8. Restart the services that require the recovered role databases.

Windows Server application considerations

This section describes considerations and limitations to backing up Windows Server applications.

- ◆ “Protecting Microsoft server applications” on page 744
- ◆ “Disk Quota database considerations” on page 744
- ◆ “SQL cluster with NetWorker User for SQL Server as backup software” on page 744

Protecting Microsoft server applications

Microsoft server applications, such as Microsoft Exchange Server, Microsoft SQL Server, Hyper-V, and Microsoft Office SharePoint Service, should be protected with both a NetWorker module and the regular NetWorker client. The NetWorker module protects the application data, such as databases and log files. The NetWorker client protects the host's critical disks for the purposes of Windows BMR.

The workflow of a complete Windows BMR recovery operation is as follows:

1. Back up critical and non-critical disks as part of the regular NetWorker file system backups.
2. Back up application data, such as Microsoft SQL Server, by using a NetWorker module, such as NMM or NetWorker User for SQL Server.
3. Perform a Windows BMR of the host's critical volumes.
4. Recover any non-critical disks by using the NetWorker User Program or, if NMM is used, the NMM recovery interface. The NMM documentation provides details on the NMM operations and recovery interface.
5. Recover application data by using a NetWorker module, such as NMM.

The NetWorker module documentation provides more information about recovering application data.

Disk Quota database considerations

Disk Quota databases are not restored during a Windows BMR. However, you can restore the Disk Quota databases after a Windows BMR by performing a regular online recovery operation and selecting **Disk Quota Databases** from the WINDOWS ROLES AND FEATURES save set.

SQL cluster with NetWorker User for SQL Server as backup software

If you use NetWorker User for SQL Server to protect a SQL server that is hosted in a Microsoft cluster, there are special considerations when using NetWorker Windows BMR. This section provides an example and the steps to follow in order to use NetWorker Windows BMR with a clustered Microsoft SQL server that is protected with NetWorker User for SQL Server.

Example 65 Protecting a clustered SQL server with NetWorker User for SQL Server and Windows BMR

This example assumes the following conditions:

- ◆ Microsoft SQL Server 2008 R2 is hosted on a Microsoft Failover Cluster on Windows Server 2008 SP2.
- ◆ NetWorker User for SQL Server is used to protect the SQL server.
- ◆ **Node & File Share Majority** is the quorum setting.
- ◆ Cluster node A has SQL installed in C:\Program Files\Microsoft SQL Server.
- ◆ Cluster node B has SQL installed in C:\Program Files\Microsoft SQL Server.
- ◆ Nodes A and B can access the shared disks that include the SQL databases.

To back up the clustered SQL server in this example:

1. Install the NetWorker client and NetWorker User for SQL Server 5.2 SP2 on both nodes A and B.
2. Perform an instance-level backup of the virtual SQL server by using NetWorker User for SQL Server. The EMC NetWorker *Module for Microsoft SQL Server Administration Guide* provides details.
3. Perform a Windows BMR backup for node A. “Windows BMR Backup” on page 747 provides more information.
4. Perform a Windows BMR backup for node B.

To recover the clustered SQL server in this example:

1. Perform a Windows BMR operation on node A first and then on node B. “Windows BMR Backup” on page 747 provides more information.

The Microsoft Failover Cluster should be running.

NOTICE

If the host is a virtual machine, first create a new virtual machine with the same configuration as before the disaster, and then perform the Windows BMR operation.

2. Export the required shared disks. Ensure that the disks have the same properties as the original ones, and that the disks are shared between node A and node B.
3. From the **Cluster Management** interface, delete any disk entries that are listed as failed disks.
4. Add the newly exported shared disks to the SQL group and assign the same drive letters as before the disaster.
5. On node A, rebuild the SQL server by running the following **Setup** command. 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:

```
Setup /QUIET /ACTION=REBUILDDATABASE /INSTANCENAME=Instance_name /SQLSYSADMINACCOUNTS=domain_name\administrator
```

The system databases are generated on the new shared disk. The following links provide more information on rebuilding SQL Server:

- SQL Server 2005
<http://msdn.microsoft.com/en-us/library/ms144259%28v=sql.90%29.aspx>
- SQL Server 2008
[http://msdn.microsoft.com/en-us/library/ms144259\(v=SQL.100\).aspx](http://msdn.microsoft.com/en-us/library/ms144259(v=SQL.100).aspx)

6. After the SQL server is rebuilt, bring the SQL group online.
7. Try to create a new database.

If the new database is not successfully created, check the disk dependencies on the SQL Server resource and correct them, if necessary.

8. Open the NetWorker User for SQL Server interface on node A and recover the backed-up database.

Windows services considerations

This section describes limitations and considerations related to Windows services.

- ◆ [“Active Directory considerations” on page 746](#)
- ◆ [“DFSRR considerations” on page 746](#)
- ◆ [“MSCS considerations” on page 746](#)

Active Directory considerations

Windows BMR of a Domain Controller is non-authoritative by default. If you need to perform an authoritative recovery, you must boot into DSRM mode directly from the NetWorker Bare Metal Recovery wizard. [“Post-recovery tasks for Active Directory services” on page 761](#) provides more information.

DFSRR considerations

The DFSRR namespaces are junction mount points and are not backed up with the DISASTER_RECOVERY:\ or ALL save set even if the DFSRR shares reside on a critical volume. To backup DFSRR Shares, either use the new save set ALL-DFSRR or provide the full DFSRR Share path as the save set name.

MSCS considerations

You might experience a recovery error if the shared disks are attached during a BMR recovery. To correct this error, detach the shared disks and perform the recover. Once the cluster client boots, attach the shared disk before performing the online recovery.

During an authoritative restore, after the restore completes, cluster services are not brought online on the remote nodes. You must bring the services online manually.

Windows Storage Pools considerations

In the event of system failure where Storage Pool disks are damaged, you must follow the manual steps described below to perform a Windows BMR recovery to a new computer. These steps follow recommendations provided by Microsoft. In the case of complete system failure, there may not be a preexisting Storage Pool on the target computer. There may only be physical disks. Some of these disks are required to create Storage Pools.

To recover Storage Spaces to a new computer, perform the following manual steps:

1. Before beginning Windows Bare Metal Recovery wizard, physically remove from the target recovery computer any physical disks reserved for storage pools. This manual step is required because the Windows Bare Metal Recovery wizard does not have any option to exclude the disks.
2. Boot the computer using the Windows Bare Metal Recovery wizard.
3. Recover only the computer's critical volumes.
4. Reboot the computer to the recovered operating system.
5. Attach physical disks that are reserved for Storage Pools.
6. Configure Storage Pools through Windows Server Manager or using Powershell Cmdlets.

7. Perform a volume or file recovery for Storage Spaces volumes.
8. Perform a volume or file recovery for other volumes on physical disks.

WinPE considerations

This section describes Windows BMR limitations that are specific to WinPE.

- ◆ [“WinPE configuration for SAN boot devices” on page 747](#)

WinPE configuration for SAN boot devices

When a system uses a SAN boot device and a recovery must be done, the WinPE environment requires that all but one path to the boot device be temporarily disabled. Once the operating system has been rebooted, the remaining paths can be re-enabled.

VMware considerations

This section describes Windows BMR limitations that are specific to VMware virtual machines.

- ◆ [“Network card driver limitations” on page 747](#)
- ◆ [“Virtual machines using VMware drivers” on page 747](#)

Network card driver limitations

The Windows BMR image does not contain a driver for any of the VMware VMXNET NIC models. However, the Windows BMR image does contain a driver for the e1000 NIC. If you have at least one e1000 NIC configured for the virtual machine, you can perform the Windows BMR recovery. Alternatively, you can add custom NIC drivers when you run the NetWorker Bare Metal Recovery wizard.

Virtual machines using VMware drivers

Recovery of a VMware Windows guest that uses VMware specific drivers requires special considerations. For example, if a guest is built by using the VMXNET or VMXNET3 NIC driver or the VMware Paravirtual SCSI driver, the WinPE environment cannot see the NIC or any hard disks. These drivers must be added by using the NetWorker Bare Metal Recovery wizard.

The drivers are part of the VMware Tools installation and are located in the Program Files\VMware\VMware Tools\Drivers folder on the virtual machine’s system drive.

Windows BMR Backup

You can run Windows BMR backups on a scheduled basis (full backups only) or manually (on demand). This section describes how to plan and set up Windows BMR backups.

- ◆ [“Include Windows BMR in scheduled backups” on page 748](#)
- ◆ [“Include Windows BMR in manual backups” on page 749](#)
- ◆ [“How to verify a valid Windows BMR backup” on page 750](#)

Include Windows BMR in scheduled backups

Although you can back up the DISASTER_RECOVERY:\ save set separately, the easiest way to schedule Windows BMR backups is to specify the save set All. Specifying the save set All ensures that all System State data and user data can be recovered in the event of a disaster.

To set up a Windows Server 2008, Windows Server 2008 R2, Windows 7, Windows Server 2012, Windows 8, Windows Server 2012 R2, or Windows 8.1 client for scheduled Windows BMR backups, you can use either the NetWorker Client Configuration wizard or the NetWorker client properties window.

- ◆ [“How to configure Windows BMR backups by using the wizard” on page 748](#)
- ◆ [“How to configure Windows BMR backups from the properties window” on page 748](#)

How to configure Windows BMR backups by using the wizard

To configure a Windows BMR backup with the NetWorker Client Configuration wizard:

1. In the NMC **Enterprise** view, select the NetWorker server name, and then double-click the **NetWorker** managed application to launch it.
2. In the **Configuration** view, under the NetWorker server name, right-click **Clients**, and then select **Client Backup Configuration > New**. Alternatively, if you are modifying an existing client, right-click the client name, and then select **Client Backup Configuration > Modify**.
3. Follow the instructions on the wizard pages.

On the **Select Files to Back Up** page, select all files for backup. This is the default selection.

- For NetWorker 7.6 SP2 clients, the option **Perform disaster recovery for this client** is selected by default. Clearing this option prevents the DISASTER_RECOVERY:\ save set from being backed up, and causes the VSS savesets to be backed up.
 - For NetWorker 8.x clients, the option **Perform disaster recovery for this client** does not appear. Select only the DISASTER_RECOVERY:\ save set and the WINDOWS ROLES AND FEATURES save sets for Windows BMR protection.
4. Complete the wizard, as instructed.

How to configure Windows BMR backups from the properties window

To configure a Windows BMR backup with the client properties window:

1. In the NMC **Enterprise** view, select the NetWorker server name and double-click the NetWorker application to launch it.
2. In the **NetWorker Administration** window, click **Configuration**.
3. In the browser tree, select **Clients**, and then perform one of the following steps:
 - To create a new client resource, select the **Clients** icon, and then from the **File** menu, select **New**.
 - To edit an existing client resource, select the client name from the list in the right panel, and then from the **File** menu, select **Properties**.
4. In the **Name** attribute, type the hostname for the client.

5. Optionally, type a comment in the **Comment** attribute.
6. Select values for **Browse Policy** and **Retention Policy**:
 - The browse policy determines how long the details of individual backed-up files are maintained in a browsable index for quick recovery through the GUI or command line.
 - The retention policy determines how long backed-up data is protected and available for recovery, even though the browse policy has lapsed. Recovery might require rebuilding an index.
7. Select the **Scheduled Backups** checkbox.
8. In the **Save Sets** attribute of the client resource, type **All**. Alternatively, type **DISASTER_RECOVERY:** to create a back up for recovery purposes only.
9. In the **Group** attribute, select a backup group. Ensure that you do not select a snapshot group.
10. In the **Pool** attribute, select a pool that targets the NetWorker devices you want to use. The pool selected in this attribute overrides any other pool that might be configured for the client or its save sets.

NOTICE

Do not use this setting for NetWorker clients earlier than version 7.6 SP1.

11. In the **Schedule** attribute, select a backup schedule.
12. The schedule selected in this attribute overrides any other schedules that might be configured for the client or its save sets.
13. When you have completed the client configuration, click **OK**.
14. In the **NetWorker Administration** window, the configured client shows a checkmark in the **Scheduled backup** column to indicate that scheduled backup is enabled.

[Chapter 2, “Backing Up Data,”](#) provides more information about setting up a scheduled backup.

Include Windows BMR in manual backups

Backing up data by selecting *Computer* is the preferred method of performing manual Windows BMR backups. This method ensures that all data is backed up.

If you select just the `DISASTER_RECOVERY:\` save set, then the NetWorker User program automatically selects the critical volumes and `WINDOWS ROLES AND FEATURES` save sets.

Manual backups launched in the NetWorker client properties window or in the command line interface are performed in a single backup stream. Backups performed with save groups on the server use multi-backup stream technology.

To perform a recovery backup manually:

1. In the NetWorker **User** program, click **Backup**, and then select **Computer** to save all data.
2. At a minimum, select the `DISASTER_RECOVERY:\` save set, all critical volumes, and the `WINDOWS ROLES AND FEATURES` save sets.

[Chapter 2, “Backing Up Data”](#) provides more information about manual backups.

How to verify a valid Windows BMR backup

After you perform a Windows BMR backup, verify that the backup exists. The save sets that make up the DISASTER_RECOVERY:\ save set correspond to each critical volume.

You can verify that the backup exists by using the NMC console, the NetWorker **User** program, or **nsrinfo** program.

If any of the save sets required for a Windows BMR backup fail, no DISASTER_RECOVERY:\ save set is created. However, some items such as the WINDOWS ROLES AND FEATURES save sets or critical volumes might have been backed up successfully and are available for online recovery.

To verify that a valid backup exists by using the NMC console:

1. In the NMC **Enterprise** view, select the NetWorker server name and double-click the NetWorker application to launch it.
2. In the **NetWorker Administration** window, click **Media**.
3. In the left pane, click **Save Sets**.
4. In the **Query Save Set** tab in the right pane, specify 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 display a list of save sets that meet the search criteria.

To verify that a valid DISASTER_RECOVERY:\ save set exists by using the NetWorker User Program:

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** to open the **Source Client** dialog box.
3. Select the client whose DISASTER_RECOVERY:\ save set you are verifying and click **OK**.
4. Select a destination client and click **OK**.
5. In the **Recover** window, browse and locate the save set named DISASTER_RECOVERY:\.

By default, the most recent backup is listed. You might have to adjust the browse time if you are verifying older files. You can adjust the browse time by selecting the **View > Change Browse Time** menu option.

To verify that a valid DISASTER_RECOVERY:\ save set exists by using the **nsrinfo** program, type the following command at the command prompt:

```
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.

This command must be run on a host that has the NetWorker client version 7.6 SP2 or later. Earlier versions of **nsrinfo** are not capable of displaying information about Windows BMR backups.

NOTICE

When performing a backup of the DISASTER_RECOVERY:\ save set, the backup may fail with an error 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.
```

This occurs because the BCD partition is offline.

Windows Bare Metal Recovery to Physical or Virtual Computers

This section describes how to use the NetWorker Windows BMR image to perform a Bare Metal Recovery on protected hosts and VMware virtual machines.

- ◆ [“To perform a Bare Metal Recovery \(BMR\) to a Physical Computer” on page 751](#)
- ◆ [“To perform a BMR from a Physical Computer to a Virtual Machine \(P2V\)” on page 763](#)

This section also includes the following topics to help you complete a Windows BMR operation:

- ◆ [“Post-recovery tasks when using NMM” on page 760](#)
- ◆ [“Troubleshooting Windows BMR” on page 765](#)
- ◆ [“Additional recovery options” on page 769](#)

To perform a Bare Metal Recovery (BMR) to a Physical Computer

Before you perform a bare metal recovery on a host, ensure that you meet the minimum requirements of having backed up the DISASTER_RECOVERY:\ save set for that host and completed the following tasks:

- ◆ [“Verify the Windows BMR Requirements” on page 752](#)
- ◆ [“Prepare to create the NetWorker Windows BMR bootable media” on page 752](#)
- ◆ [“Obtain the Windows Bare Metal Recovery wizard” on page 752](#)
- ◆ [“Create a Windows BMR bootable image” on page 753](#)
- ◆ [“Perform a NetWorker Bare Metal Recovery wizard test before recovery” on page 754](#)

Verify the Windows BMR Requirements

In preparation for a recovery, verify that the new computer meets the “[Hardware Requirements for Windows BMR Backup and Restore](#)” on page 732 and the “[Configuration requirements for Windows BMR backups](#)” on page 733.

Prepare to create the NetWorker Windows BMR bootable media

NetWorker provides a Windows BMR image that you can use to create a bootable CD or deploy for a network boot operation. To begin the recovery process, boot the Windows WinPE operating system from the bootable CD or network boot location. The recovery starts a NetWorker Bare Metal Recovery wizard that guides you through the recovery process.

Note: The Microsoft Windows(R) Preinstallation Environment software included with this computer or software may be used for boot, diagnostic, setup, restoration, installation, configuration, test, or disaster recovery purposes only. NOTE: THIS SOFTWARE CONTAINS A SECURITY FEATURE THAT WILL CAUSE YOUR COMPUTER SYSTEM TO REBOOT WITHOUT PRIOR NOTIFICATION AFTER 72 HOURS OF CONTINUOUS USE.

NOTICE

FOR AUSTRALIA ONLY. References to “Limited Warranty” are references to the warranty provided by the manufacturer or installer. This warranty is given in addition to other rights and remedies you may have under law, including your rights and remedies in accordance with the statutory guarantees under the Australian Consumer Law. Our goods come with guarantees that cannot be excluded under the Australian Consumer Law. You are entitled to a replacement or refund for a major failure and compensation for any other reasonably foreseeable loss or damage. You are also entitled to have the goods repaired or replaced if the goods fail to be of acceptable quality and the failure does not amount to a major failure.

Goods presented for repair may be replaced by refurbished goods of the same type rather than being replaced. Refurbished parts may be used to repair the goods.

For further information regarding this warranty and to claim expenses in relation to the warranty (if applicable), please contact the manufacturer or installer; see the contact information provided in the system packaging.

Obtain the Windows Bare Metal Recovery wizard

The Windows Bare Metal Recovery wizard files are located on the EMC Online Support website at:

<https://support.emc.com/>

Select the file that is appropriate for the computer that you recover. The choices are:

- ◆ NetWorker 8.1 SP1 Windows BMR Wizard (x86)
 - For computers that run 32-bit versions of Windows 7, Windows 8, or Windows Server 2008
- ◆ NetWorker 8.1 SP1 Windows BMR Wizard (x64)

- For computers that run 64-bit versions of Windows 7, Windows 8, Windows 8.1, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012 or Windows Server 2012 R2
- Supports computers that use UEFI volumes. The topic [“UEFI Partition Support” on page 721](#) provides additional information.

NOTICE

WinPE is only available in English. There are no localized versions of the Windows Bare Metal Recovery wizard.

NOTICE

The build number for the release version of NetWorker 8.1 SP1 is located in the NetWorker 8.1 SP1 Release Notes. This documentation refers to the build number as *xxx*.

To download the recovery boot image:

1. On the EMC Online Support website, search for “NetWorker Wizard ISO” at EMC Online Support and narrow the search results by selecting items associated with the NetWorker 8.1 SP1 release.
2. On the **NetWorker Software Downloads** page, locate the section labeled **NetWorker 8.1 SP1 - Build *xxx***, where *xxx* is the build number of the released version, and then select the appropriate link to download a Windows BMR ISO recovery file:
 - Select **NW 8.1 SP1 Windows BMR Wizard x86** to download the file `NetWorker_8.1.1.XXX_Windows_BMR_Wizard_x86_WinPE_50.iso`
 - Select **NW 8.1 SP1 Windows BMR Wizard x64** to download the file `NetWorker_8.1.1.XXX_Windows_BMR_Wizard_x64_WinPE_50.iso`

The download process begins.

Create a Windows BMR bootable image

Create either a Windows BMR bootable CD or a network boot location from the downloaded ISO image.

Creating a Windows BMR bootable CD

To create a bootable CD:

1. Open your CD creation software and select an option to burn an ISO image.
2. Browse to the location of the downloaded NetWorker Windows BMR image and complete the steps required to create a bootable CD with the image.

Enabling a host to boot from a CD

To ensure that the protected host can boot from a CD:

1. Start the host and enter the BIOS setup program, typically by pressing **F2**.

Note: If you are restoring a virtual host such as a VMware virtual machine, you can set up options such as the host boot location from within vSphere. The VMware documentation provides specific steps.

2. Select the boot options menu and ensure that the CD boot option is at the top of the list of locations from which to boot.
3. Save your changes and exit the BIOS program.

Creating a Windows BMR recovery network boot location

Ensure that you meet the following requirements for using the network boot option:

- ◆ The NetWorker clients that you are protecting must be enabled to boot from the network with a Pre-Boot Execution Environment (PXE).
- ◆ A Deployment Services server must be configured and available.
- ◆ The NetWorker Windows System Recover boot image must be added to the Deployment Services server so that a client host on the network can boot from it.

For example, the following link provides details on how to configure Windows Deployment Services in Windows Server 2008:

[http://technet.microsoft.com/en-us/library/cc771670\(ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc771670(ws.10).aspx)

Enabling a host to boot from the network

To enable a host to boot from the network (PXE enabled):

1. Start the host and enter the BIOS setup program, typically by pressing **F2**.

Note: If you are restoring a virtual host such as a VMware virtual machine, you can set up options such as the host boot location from within vSphere. The VMware documentation provides specific steps.

2. Select the BIOS options necessary so that the network boot option is enabled. The BIOS documentation provides more information.
3. Save your changes and exit the BIOS program.

Your host should obtain an IP address from your WDS server and prompt you if you want to do a network boot. Typically, a network boot is activated by pressing the F12 key.

Perform a NetWorker Bare Metal Recovery wizard test before recovery

Before you need to perform a Windows BMR, test the wizard to ensure that you can run it and that you have the required drivers. This task is especially important for 64-bit hosts that might require additional drivers to complete the wizard. For both 64-bit and 32-bit hosts, the wizard must use drivers that do not require a reboot.

After you test the wizard, you can safely exit the wizard before completing the entire recovery process.

To test the wizard:

1. Complete steps 1 to 10 in [“Perform a Windows BMR recovery using the wizard” on page 755](#).

Verify the following as you step through the wizard screens:

- If DNS is not available, 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.

Perform a Windows BMR recovery using the wizard

To perform a Windows BMR using the wizard:

1. Boot the host to be recovered from the location of the NetWorker Windows BMR image, either a bootable CD or network boot location.
If the computer to be restored has UEFI volumes, unmount any UEFI volumes. Only the computer is recovered.
2. Click **Next** when the **Welcome** screen of NetWorker **Bare Metal Recovery wizard** appears.
3. Perform the following steps only if there is no DNS server available on the network. If a DNS server is not available, you can manually add and edit a hosts file to add information about the NetWorker server.
 - a. Exit the NetWorker **Bare Metal Recovery wizard** but do not reboot the host.
You are be returned to the WinPE command line.
 - b. Edit the hosts file, for example, X:\Windows\System32\Drivers\etc\hosts, and add the IP address and hostname for the NetWorker server, NetWorker storage node, and Avamar deduplication node (if one is being used).
 - c. 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
```
 - d. When the wizard appears, click **Next** to continue.
4. In the **Select Network Interface** screen, select the NIC driver so that the host can communicate with the NetWorker server during the recovery process.

If the required NIC driver is not in the list, click **Load Driver** to browse to a location, such as a CD or USB drive, and locate the required driver.

The selected driver cannot require a reboot operation because the WinPE environment is loaded in memory only and changes are not persistent across a reboot operation. Although some drivers prompt for a reboot operation, most modern NIC drivers are generally plug-and-play, and ignoring the reboot prompt might actually work.
5. Click **Next**.

6. Complete the fields on the **Configure Hostname and Network** screen:
 - a. Type the hostname of the computer that you are recovering in the **Hostname** field.
 - b. Type the name of the domain in which the host resides in the **DNS domain** field. If the host resides in a workgroup instead of a domain, you can leave this field blank.
 - c. Select a tab under the **Configure desired IP Settings** field. Choose the tab for the Network Protocol deployed on your network, either IPv4 or IPv6.
 - d. Under the **TCP/IP Address** settings, select either **Obtain an IP address automatically (DHCP)** or **Use the following IP Address**.

If the host to be recovered has an assigned IP address, type the IP address in the **IP address** field. If applicable, type the subnet mask in the **Subnet mask** field and the default gateway in **Default gateway** field.

- e. Under **DNS Server**, select either **Obtain DNS server address automatically** or **Use the following DNS server address**.

If the host to be recovered uses a DNS server with a static IP 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.

Note: You can ignore the DNS Server fields if you added the NetWorker server hostname and IP address to the X:\Windows\System32\Drivers\etc\hosts file in [step 3](#).

- f. Click **Next**.

All local disks that have been detected are displayed in the **Available Disks** screen.
7. If the wizard has failed to detect a disk, click **Load Driver** to browse to a location, such as a CD or USB drive, and locate the correct driver for the disk. After loading the required disk driver, click **Refresh** to update the list of disks that have been detected.
8. Click **Next**.
9. Complete the fields on the **Select NetWorker Server** screen:
 - a. In the **Server** field, specify the NetWorker server to which the host was backed up by double-clicking the appropriate NetWorker server from the list or typing the fully-qualified domain name (FQDN). You have to click **Search** to update the list of NetWorker servers. The Search function locates only those NetWorker servers on the local subnet.
 - b. In the **Client** field, ensure that the client name matches the client resource name on the NetWorker server. For example, if the client resource on the NetWorker server uses a FQDN, then use the client's FQDN in the **Client** field.

This field is automatically populated with the values that were typed in to the **Hostname** and **DNS Domain** fields on the **Configure Hostname and Network** screen of the wizard.

You can modify the **Client** field if you want to recover a backup that was created by a different computer. However, the hardware configuration of the target computer must be similar to the original computer. You must also satisfy the following requirements for performing a directed recovery:

- The NetWorker server must have a client resource for both the source computer and the target computer.
- The **Remote Access** attribute of the client resource for the source computer must allow access to the user or host computer that is performing the directed recovery operation. Do this by adding **SYSTEM@target_client** to the source client resource's **Remote Access** attribute.
- Add "**user=system,host=target_client**" to the **Users** attribute of the NetWorker server's preconfigured Administrators user group.

NOTICE

If a different client (target) is specified, the recovered computer uses the same hostname and IP settings as the source computer. This can cause hostname and IP address conflicts if the source computer is running on the same network.

e. Click **Next**.

10. Select the system backup that you wish to recover to the host in the **Select System Recovery** screen. System backups are listed in descending order from most recent to oldest.

11. Click **Next**.

The **Save Sets to Restore** screen lists the volumes to be recovered to the host, as shown in [Figure 50 on page 757](#).

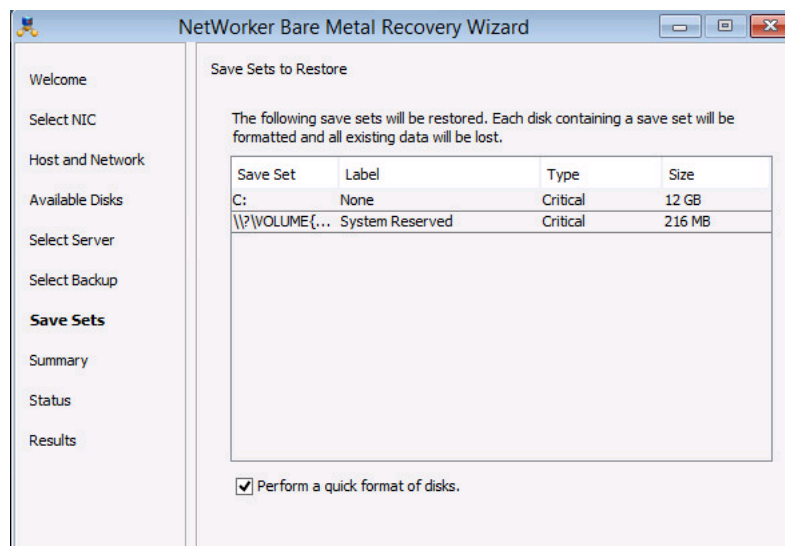


Figure 50 Save Sets to Restore

Beginning with NetWorker 8.0, VSS save sets are listed with the critical volumes. If you are recovering a Windows Server 2008 R2 or Windows 7 computer, B volume save sets are listed as critical volumes because the boot data is held on a partition separate from the operating system partition.

During the recovery process, critical volumes are reformatted. Non-critical volumes are reformatted only if the disk signature is different, for instance if the disk was replaced.

12. To perform a quick format instead of a full format operation, select **Perform a quick format of disks**. This is the default selection. Although quick formatting is much faster than full formatting, a quick format, unlike a full format, does not verify each sector on the volume.

The recovery process does not recover non-critical volume data even if the non-critical volume is reformatted. Non-critical volumes can be recovered, if necessary, by using the NetWorker **User** program after the wizard has completed and the host has been rebooted.

13. Click **Next**.

The **System Recovery Summary** screen lists the selected recovery options.

14. If you need to specify any non-default recovery options, click **Options** to display the **Non-Default Recover Options** screen, shown in [Figure 51 on page 758](#).

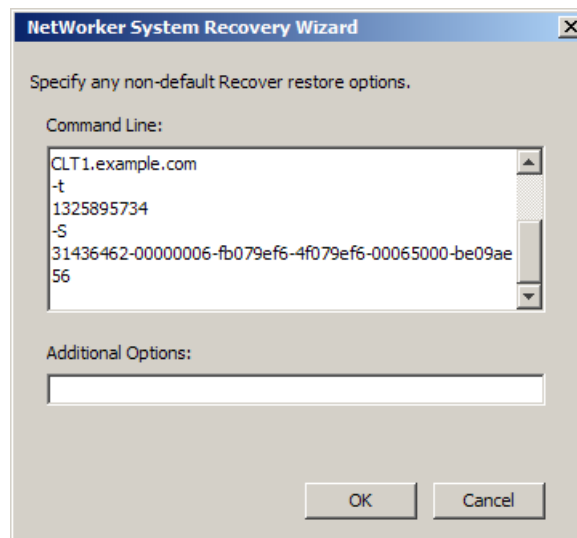


Figure 51 Non-Default Recover Options

On the **Non-Default Recover Options** screen:

- a. Type any required non-default options with their corresponding values in the **Additional Options** field. Non-default options are primarily used for troubleshooting purposes. [“Additional recovery options” on page 769](#) provides more information.
 - b. Click **OK** to save and close the **Non-Default Recover Options** screen and return to the **System Recovery Summary** screen.
15. Click **Restore** to begin the recovery process.
 16. In the confirmation screen that appears, confirm your choice to begin the recovery process and then click **OK**.

NOTICE

All data is lost on volumes that are being reformatted.

At the end of the recovery process, log files are backed up to the NetWorker server. These files can be used to troubleshoot a failed recovery. Ensure that the NetWorker server has a mounted writable backup volume available for these log files. Otherwise, the recovery is not complete because it waits for a writable volume. If this occurs, you can cancel the log file backup without affecting the recovery operation. By default, recovery log files are written to the default backup pool. If desired, you can set up a special backup pool for log files. To do so, ensure that the backup pool accepts manual save sets that are named Offline Restore Logs. [“Recovering and viewing Windows BMR log files” on page 766](#) provides more information about accessing log files.

The **System Recovery Results** screen is displayed as shown in [Figure 52 on page 759](#) when the recovery process has completed.




Figure 52 System Recovery Results

17. Click **Reboot** or **Exit** as appropriate:

- Click **Reboot** to reboot the system and restore the application data. If you are recovering an Active Directory domain controller, it is recovered in non-authoritative mode by default.
- If you need to recover the domain controller in authoritative mode, click **Exit** to return to the WinPE command prompt and boot into Directory Services Restore Mode (DSRM). [“Post-recovery tasks for Active Directory services” on page 761](#) provides more information.

Post-recovery tasks

The following sections provide information about recovering data that was not recovered in the Windows BMR operation:

- ◆ [“Post-recovery tasks when using NMM” on page 760](#)
- ◆ [“Post-recovery tasks when using an application backup tool other than NMM” on page 761](#)
- ◆ [“Post-recovery tasks to recover file system data” on page 761](#)
- ◆ [“Post-recovery tasks for Active Directory services” on page 761](#)
- ◆ [“Post-recovery tasks for hosts with Windows server roles that use SQL Server” on page 763](#)
- ◆ [“Post-recovery tasks for a Microsoft Hyper-V virtual machine” on page 763](#)

Post-recovery tasks when using NMM

If the recovered host has applications 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.

Before proceeding to the NMM documentation, ensure that the following conditions are met:

- ◆ 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. In this case, you should 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 this correct drive letter assignments do not occur, manually assign drive letters to non-critical disks as needed. Non-critical volumes accessed by mount points might have similar issues.
- ◆ To completely recover the host, you might have to perform additional online recovery steps 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 is a Microsoft limitation.
- ◆ Windows BMR can back up critical volumes that are BitLocker encrypted. However, the recovered volumes are not encrypted. You can use the BitLocker icon in the Control Panel to reapply the volume encryption. This is a Microsoft limitation.

Post-recovery tasks when using an application backup tool other than NMM

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 [“Post-recovery tasks to recover file system data” on page 761](#).
- ◆ 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 comes with your application backup tool.

Post-recovery tasks to recover file system data

Perform an online recovery of any required user data on non-critical volumes. In some cases, user data on non-critical volumes needs to be recovered, for instance, when disk hardware was replaced due to a disaster prior to the Windows BMR operation.

To perform an online recovery:

1. Manually remount any non-critical volumes as needed.
2. Start the NetWorker User program by using the **winworkr** command with the **-s** option to connect to the NetWorker server that backed up the source client data:

```
winworkr -s server_name
```

If the **-s** option is not used 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, allowing you to choose the server.

3. Click **Recover** to open the **Source Client** dialog box.
4. Select the source client, and then click **OK**.
5. Select the destination client for the recovered data, and click **OK**.
6. In the **Recover** window, select the files to recover.
7. Click **Start** to begin the directed recovery.

The following sources provide more information:

- [Chapter 14, “Recovering Filesystem Data”](#) provides more information about recovery options.
- The *NetWorker Procedure Generator* provides more information about an Active Directory online recovery and Windows Server Failover Cluster online recovery.
- [Appendix C, “Backing Up and Restoring a Microsoft DFS”](#) provides more information about a DFSR online recovery.

Post-recovery tasks for Active Directory services

The offline recovery of the DISASTER_RECOVERY:\ save set in the case of a domain controller is non-authoritative. If a non-authoritative recovery is desired, then no additional steps are required. However, if you need to perform an authoritative recovery, follow these steps.

To perform an authoritative recovery:

1. On the last screen of the **NetWorker Bare Metal Recovery wizard**, titled **System Recovery Results**, do not select **Reboot**. Instead, select **Exit** to exit the wizard so that you can boot into Directory Services Restore Mode (DSRM).

NOTICE

Failure to boot into DSRM mode results in a non-authoritative recovery. If this occurs, run Windows BMR again and ensure that you boot into DSRM mode.

The WinPE command prompt appears.

2. At a command prompt, type the following **bcdedit** commands to add a boot loader entry to force the system to boot into DSRM:

- a. Type the following command to add a boot loader entry:

```
X:\>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. Using the generated GUID, type the following command to set the safeboot option for the boot loader entry in the BCD store:

```
X:\> bcdedit /set {GUID_value} safeboot dsrepair
```

where *GUID_value* is the GUID displayed by the previous **bcdedit** command.

- c. Exit the command prompt to reboot the system.

Failure to boot into DSRM results in a non-authoritative recovery.

3. (Optional) If you have WINDOWS ROLES AND FEATURES save set backups and they are more recent than the DISASTER_RECOVERY:\ save set, you can recover them in DSRM by using the NetWorker User program.

4. Run the Windows **ntdsutil** command-line 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 command prompt so that the host does not boot into DSRM mode upon reboot.

```
C:\> bcdedit /deletevalue safeboot
```

9. Restart the domain controller in normal mode and then log in and verify that the authoritative changes are replicated to the Active Directory replication partners.

Post-recovery tasks for hosts with Windows server roles that use SQL Server

To recover hosts with Windows server roles that use SQL server:

1. On node A, rebuild the SQL server by running the following Setup command. 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:

```
C:\> Setup /QUIET /ACTION=REBUILDDATABASE
/INSTANCENAME=Instance_name
/SQLSYSADMINACCOUNTS=domain_name\administrator
```

The following Microsoft URL provides more information:

[http://msdn.microsoft.com/en-us/library/ms189302\(SQL.105\).aspx](http://msdn.microsoft.com/en-us/library/ms189302(SQL.105).aspx)

2. Bring the SQL sever services online.
3. Recover the SQL system databases (master, model, msdb) with NetWorker User for SQL Server, or a third-party application.

Post-recovery tasks for a Microsoft Hyper-V virtual machine

Use NMM to restore your Hyper-V virtual machines.

To perform 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).

P2V is supported for physical computers running the following operating systems:

- Microsoft Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

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
1. Perform a backup of the physical computer as described in [“Windows BMR Backup” on page 747](#).
 2. On the computer that runs your hypervisor, create a target virtual machine (VM).
For VMware hypervisors, use a Windows Server 2008 R2 template as the guest operating system when you create the VM.
 - a. Configure the VM to use a virtual network adapter.
 - b. On the **VM configuration** page, select the LSI Logic SAS SCSI controller. Configure the disks on the VM to match the original physical computer configuration.
 - Create the same number of physical disks. Extra disks can be added after the P2V recovery.
 - The SCSI disk numbers must match the original disk numbers.
 - The VM disk sizes must match, or exceed, the original disk sizes.
 3. On the VM, boot the WinPE ISO which starts the BMR Recovery wizard. Refer to [“Create a Windows BMR bootable image” on page 753](#) for more information.
 4. On the VM, use the BMR Recovery wizard to configure the hostname and network configuration.

On the **Select Server** page, ensure that you specify the name of the physical computer as the NetWorker client.

5. On the **Select Backup** page, select the backup to restore. Backups are listed in chronological order with the most recent backup first.
6. On the **Summary** page, if the save set was created with NetWorker 8.1 or earlier, select the Restore physical computer to virtual machine (P2V) checkbox.

Note: If the Restore physical computer to virtual machine (P2V) checkbox is not marked, the VM might not boot successfully after the restore is complete.

7. Select **Restore** to start the restore.
8. Reboot the VM when the P2V BMR is complete.

Post-P2V tasks

The following section provides information about additional tasks that are required after a P2V recovery.

1. If you are running VMware, install VMware tools.
2. Use Device Manager to remove disabled NIC devices. This is required because the original network adapter is no longer available. It can also be required for other devices.

To remove disabled devices:

- a. From Device Manager, select the Show Hidden Devices option.
 - b. Select the hidden NIC device.
 - c. Select uninstall.
3. Configure the virtual network adapter to restore network connectivity.

Troubleshooting Windows BMR

The following topics provide information to help troubleshoot Windows BMR operations.

- ◆ [“Manually uninstall and reconfigure a NIC on Windows 2012” on page 765](#)
- ◆ [“Recovering and viewing Windows BMR log files” on page 766](#)
- ◆ [“BMR backup fails when System Reserved Partition is offline” on page 767](#)
- ◆ [“Errors during backup of Windows Server 2003 Enterprise Edition or Windows Server 2003 R2” on page 768](#)
- ◆ [“Wizard cannot locate the NetWorker server or DNS server” on page 768](#)
- ◆ [“Multiple NICs cause errors in locating the NetWorker server” on page 768](#)
- ◆ [“Network configuration values might not be retained after reboot” on page 769](#)
- ◆ [“VSS backups fail because a critical disk is offline” on page 769](#)
- ◆ [“Jobquery fails to establish a connection with large scale jobs” on page 769](#)

Manually uninstall and reconfigure a NIC on Windows 2012

If the guest operating system is Windows 2008 R2, the NIC settings are retained by the P2V BMR.

However, if the guest operating system is Windows 2012 then Windows performs some Plug-N-Play configuration during the post-BMR reboot. This activity disables the original NIC and creates a new NIC.

Perform the following to configure the new NIC:

1. In the device manager select Display disabled devices > Uninstall the disabled NIC.
2. Configure the new NIC with the desired network settings.

Recovering and viewing 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 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 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 related to critical volume data recovery.
- ◆ **WinPE_Wizard.log** — Contains information about the work flow 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 related to Microsoft native API calls.

If the Windows BMR fails, you can recover the log files by using FTP on the recovery host or by using a directed recovery. 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.

Accessing the log files

To access the log files by using FTP:

1. Access the WinPE command line on the recovery host.

You might have to exit the Windows Bare Metal Recovery wizard to access the WinPE command line. If you exit the wizard, do not reboot.

2. Disable the Windows firewall. For example:

```
X:\Program Files\EMC Networker\nsr\wizard>wpeutil DisableFirewall
```

By default, the Windows firewall is enabled on WinPE, and this blocks the FTP port from transferring files.

3. Move to the following directory that contains the log files:

```
X:\Program Files\EMC Networker\nsr\logs
```

4. Use the FTP utility to move the log files to another NetWorker host.

To access the log files by using a directed recovery operation:

1. Start the NetWorker User program by using the **winworkr** command with the **-s** option to connect to the NetWorker server that backed up the source client data:

```
winworkr -s server_name
```

If the **-s** option is not included, 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, allowing you to choose the server.

2. Click **Recover** to open the **Source Client** dialog box.
3. Select the source client, which is the client that was being recovered, and click **OK**.
4. Select the destination client for the recovered data, and 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. Click **Start** to begin the directed recovery.

[Chapter 14, "Recovering Filesystem Data,"](#) provides more information about the permissions required for directed recoveries.

Viewing the log files

To view the log files:

- ◆ Use a text editor to view the following log files:
 - recover.log
 - WinPE_Wizard.log
- ◆ Use the **nsr_render_log** program to view the following log files:
 - Ossr_director.raw
 - winpe_nw_support.raw

For example, type the following command at a command prompt to display the **Ossr_director.raw** file:

```
c:\> 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:

```
c:\> nsr_render_log "C:\logs\Client-bv1\Ossr_director.raw" > mylog.txt
```

BMR backup fails when System Reserved Partition is offline

Windows Server 2008 R2 has a 100MB as System Reserved Partition. When backing up system state, VSS includes the System Reserved Partition but the backup fails because the System Reserved Partition is offline.

Automount must be enabled for a BMR backup to succeed. Refer to the NetWorker 8.1 SP1 Release Notes for more information.

Errors during backup of Windows Server 2003 Enterprise Edition or Windows Server 2003 R2

When you create a backup on either a Windows Server 2003 Enterprise Edition or Windows Server 2003 R2 with any writer or service that requires that a database and log files be installed at a location other than the default boot drive NetWorker will report errors.

To prevent these errors, directory must to be manually created and available at the new installation location for the database and log files before you install the service to these locations.

If the database and log files are installed at the root level of a volume, NetWorker backup will report errors.

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), try 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 and Avamar deduplication node (if one is being used) was entered correctly.
- ◆ If you are using a DNS server, verify that the values entered in the **Configure Hostname and Network** screen, were entered 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 included in the WinPE environment:

1. Exit the **NetWorker Bare Metal Recovery wizard** but do not reboot the host.

You are returned to the WinPE command line.

2. Use the **ping** utility to locate and verify hostnames and IP addresses. For example:

```
X:\Program Files\EMC Networker\nsr\wizard>ping -a hostname
```

3. Restart the wizard. For example:

```
X:\Program Files\EMC Networker\nsr\wizard> 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 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

In some cases, a host does not retain its network configuration data after a Windows BMR operation and after the host boots. If the recovered host is experiencing network connectivity issues, confirm that network properties for the local connections are correct. If necessary, 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 might be able to remedy the problem by following the steps outlined in the Microsoft Knowledgebase (KB) article 980794, which can be found at:

<http://support.microsoft.com/kb/980794>.

The patch mentioned in this KB article is most likely on your Windows system if it has been kept up-to-date. In this case, you only have to create and populate the Registry keys as described in the article.

This issue is most often encountered when backing up a passive node in a MSCS cluster and a critical volume is not located on the physical host of the passive node but is instead located on the physical host of the active node.

Jobquery fails to establish a connection with large scale jobs

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]
```

Additional recovery options

You can specify non-default recovery options on the WinPE command line or in the **Additional Options** field of the **NetWorker Bare Metal Recovery wizard**, shown in [Figure 53 on page 769](#).




Figure 53 Additional Options

Table 123 on page 770 describes the additional recovery options that can be used with a Windows BMR operation.

Table 123 Additional recovery options

Entry	Result
<p>-D <i>n</i> where <i>n</i> is a number from 1 to 9, with 9 providing the most debug information and 1 providing the least.</p>	<p>Additional debug information is included in the Windows BMR log files.</p>
<p>-v</p>	<p>Additional information on the progress of the recovery displays in the wizard's System Recovery Status window.</p>
<p>-p</p>	<p>By default, the Windows BMR recovery skips the formatting of non-critical disks.</p> <p>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 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 -p option might resolve those errors.</p> <p>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 rebooted.</p>
<p>recover -s <NetWorker server> -U -N "WINDOWS ROLES AND FEATURES\Cluster Database"</p>	<p>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.</p> <p>While the recovery is in progress, observe that the status of the groups change 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.</p> <p>Recover the shared drive data through Winworkr on the cluster node with its current active node. Select source client as virtual client and destination client as the current active node.</p>

CHAPTER 26

Volume Shadow Copy Service

This chapter covers these topics:

- ◆ Overview of VSS 772
- ◆ VSS and the backup process 772
- ◆ Controlling VSS from NetWorker software 775

Overview of VSS

If the NetWorker Module for Microsoft is installed on the client computer, information in this chapter may be superseded by information in the NetWorker Module for Microsoft documentation. The *EMC 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.

In NetWorker software releases 7.2 and later, NetWorker software 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 Snapshot 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 774](#) 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.

[Figure 54 on page 773](#) provides a graphical representation of the VSS backup process.

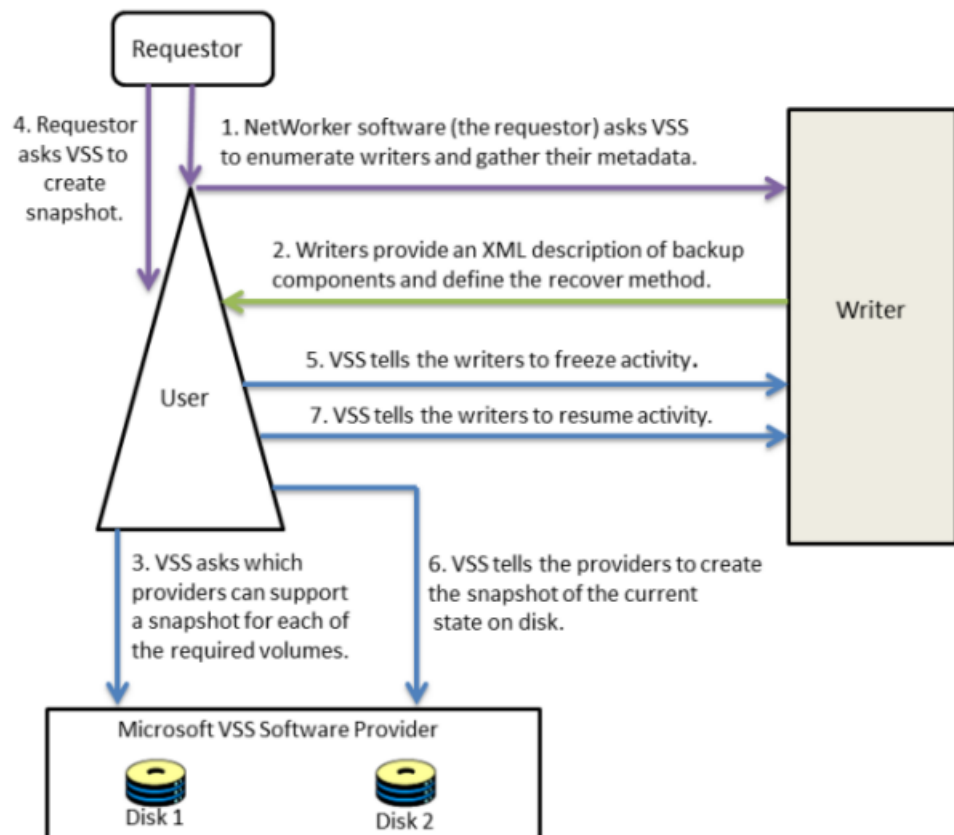


Figure 54 VSS backup process

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 775](#) provides more information about specifying VSS commands for a NetWorker client. [“VSS commands for Windows 2008, 7, and higher” on page 780](#) 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 774](#) 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 properly backing up data. They provide metadata information about what data to back up, and specific methods for properly handling components and applications during backup and restore. They also identify the type of application or service that is being backed up, for example System Boot or System Services. 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 will not be available. Consequently, a writer can appear or disappear from backup to backup.

In addition, NetWorker software maintains a list of supported writers in the NSRLA database of the client machine. When backing up data, the software checks to ensure that these conditions exist:

- ◆ The writer 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 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 *EMC NetWorker Software Compatibility Guide* 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. [Appendix A, “SYSTEM, VSS SYSTEM, and WINDOWS ROLES AND FEATURES Save Sets”](#) provides details. For the file system, this means the software attempts 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 given 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 need to disable VSS. You can control VSS from the Administration window, the NetWorker User program, or the command prompt.

Note: Microsoft’s newer versions of Windows (Windows Vista and above and Windows Server 2008 and above) recommend VSS for backups and requires it for the backup of the operating system state data.

These sections provide more information:

- ◆ [“Controlling VSS from the Administration window” on page 776](#) provides details on how to control VSS from the Administration window.
- ◆ [“Controlling VSS from the NetWorker client” on page 776](#) provides details on how to control VSS from the NetWorker User program.
- ◆ [“Control VSS from the command-prompt” on page 777](#) provides details on how to control VSS from the command prompt.
- ◆ [“Globally disabling VSS” on page 778](#) provides details on how to disable VSS globally.

Controlling VSS from the Administration window

To control VSS from the Administration window:

1. From the **Administration** window, click **Configuration**.
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 appropriate command according to [Table 124 on page 779](#), 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.

Controlling VSS from the NetWorker client

NOTICE

This section does not apply to NetWorker clients on Windows Vista and higher or on Windows Server 2008 and higher.

The Local Save Operations dialog box is read-only once the Backup window is opened. To modify local save operations after opening the Backup window, exit and restart the NetWorker User program.

To modify Local Save Operations, run only one instance of the NetWorker User program at a time.

To control VSS from the NetWorker client:

1. Start the NetWorker **User** program.
2. From the **Options** menu, select **Save Operations**.
3. In the **Local Save Operations** dialog box, type the appropriate command according to [Table 124 on page 779](#), then click **OK**. Separate multiple commands with a semicolon (;).

If the **Local Save Operations** dialog box is left empty, NetWorker software backs up data by using VSS:

- Local Save Operations 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. If you enter a VSS command in the **Local Save Operations** dialog box on the client, the command runs only if the backup is started by using the NetWorker User program on the client.
- When typing a writer name in the **Local Save Operations** dialog box, the name must match the name the writer uses to identify itself. If the name does not match, the command is ignored.
 - To confirm the name, browse the corresponding save set in the **Backup** window.
 - Remember that the list of writers that appear under each save set is dynamically determined at runtime.
- A writer might not appear under its corresponding save set if:
 - NetWorker software does not support the writer.
 - The writer's associated application, service, or database is not currently running on the system.
 - The writer has already been disabled using the NetWorker User program.

Control VSS from the command-prompt

VSS can be controlled from the command-prompt on a NetWorker client or the Console server by using the **-o** option and the Save Operations commands in [Table 124 on page 779](#), 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 *EMC 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.

To disable VSS:

1. Log in as root or as Windows Administrator on the NetWorker server.

To disable VSS for all NetWorker clients:

- a. Create an input file for the **nsradmin** command. The input file will eliminate interactive prompting as each client gets updated. For example, create a text file named `disable-vss.txt` and type the following into the file:

```
show name; client OS type; Save operations
print type: NSR client
update Save operations: "VSS\:*=off"
print
```

To disable VSS only for clients on a particular Windows operating system such as Windows NT:

- a. Create an input text file, for example, create a file named `disable-vss-nt.txt` and 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 either of 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.

VSS commands

This section lists the commands and syntax used to control VSS. [“Controlling VSS from NetWorker software” on page 775](#) describes how to enter these commands in NetWorker.

Commands for the following Windows operating systems are provided:

- ◆ [“VSS commands for Windows XP and 2003” on page 779](#)
- ◆ [“VSS commands for Windows 2008, 7, and higher” on page 780](#)

VSS commands for Windows XP and 2003

Table 124 VSS commands on Windows XP and 2003

Task	Entry	Result on Windows XP or 2003
Enable VSS.	This attribute should be left empty.	Leaving the attribute empty results in NetWorker software automatically using VSS.
Completely disable VSS.	VSS:*=off	The file system and the system components are backed up by using the legacy method, which means that the backup is performed without taking a snapshot.
Instruct NetWorker software to back up the file system by using VSS only.	VSS: <i>root drive path</i> =only To indicate all drives, enter: VSS:*:=only For example: VSS:C:\=only	All VSS SYSTEM save sets, writers, and file systems are backed up by using VSS. However, if VSS fails, instead of backing up the file system by using the legacy method, the specified drive is not backed up at all.
Disable VSS for a particular drive.	VSS: <i>drive</i> :\=off For example: VSS:c:\=off	The specified drive is backed up by using the legacy method.
Disable an individual writer.	VSS: <i>writer</i> =off where <i>writer</i> is the name of the writer to disable. For example: VSS:WINS Writer=off	The application data served by <i>writer</i> is not saved, unless a NetWorker Module exists for that application and is installed and configured on the system. When a writer is disabled, NetWorker still processes the writer so its files are skipped during the file system backup.

Notes:

When typing a writer name in the **Save Operations** attribute, the name must match the name that the writer uses to identify itself. If the name does not match, the command is ignored.

- To confirm the name, open the NetWorker User program and browse the corresponding save set.
- Remember that the list of writers appearing under each save set is dynamically determined at runtime.
- A writer might not appear under its corresponding save set under the following conditions:
 - NetWorker software does not support the writer.
 - The writer’s associated application, service, or database is not currently running on the system.
 - The writer has already been disabled by using the NetWorker User program.

VSS commands for Windows 2008, 7, and higher

Table 125 VSS Commands on Windows 7, 2008, and higher (1 of 2)

Task	Entry	Result on Windows 7, 2008 or higher
Enable VSS.	This attribute should be left empty.	Leaving the attribute empty will result in NetWorker software automatically using VSS.
Completely disable VSS.	VSS:*=off	<p>VSS backups will not occur and backing up the following save sets for a NetWorker client resource will yield these results:</p> <ul style="list-style-type: none"> • DISASTER_RECOVERY:\ save set Backup will fail at the beginning of backup operation. • VSS SYSTEM STATE save sets Backup will fail. • VSS SYSTEM STATE save sets and volume save sets Volumes will back up but VSS save sets will fail. • All save set Backups will fail.
Use a hardware provider or a specific software provider for a NetWorker client backup.	VSS:VSS_ALLOW_DEFAULT_PROVIDER=yes	<p>A backup provider is selected based on the following default Microsoft criteria:</p> <ol style="list-style-type: none"> 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. <p>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 719 provides more information about Windows Bare Metal recovery backups.</p>

Table 125 VSS Commands on Windows 7, 2008, and higher (2 of 2)

Task	Entry	Result on Windows 7, 2008 or higher
Do not embed the DISASTER_RECOVERY:\ save set within the All save set definition.	VSS:DISASTER_RECOVERY=off Default: No value, which means that the All save set will include the DISASTER_RECOVERY:\ save set. “Windows Bare Metal Recovery” on page 719 contains more information about using VSS commands with the Windows Bare Metal Recovery feature.	The DISASTER_RECOVERY:\ save set is not included in the All save set. Instead, the All save set is consistent with pre-7.6 SP2 NetWorker. This means that the save set All will consist of the following save sets: <ul style="list-style-type: none"> • VSS SYSTEM BOOT • VSS SYSTEM FILESET • VSS SYSTEM SERVICES • VSS USER DATA • All local physical drives “Blending Windows BMR recovery backups with synthetic full backups” on page 739 provides an example of when you may want to use this command.
Limiting the frequency of VSS SYSTEM STATE backups.	VSS:VSS_SYSTEM_SAVESETS =off Default: No value, which means that VSS SYSTEM STATE save sets are backed up at level full when the DISASTER_RECOVERY:\ save set is included in an incremental level backup.	This entry prevents the DISASTER_RECOVERY:\ save set from triggering a full backup of the VSS SYSTEM STATE save sets when the backup level is incremental. The VSS SYSTEM STATE save sets will still be saved during a level full backup. “Windows BMR limitations and considerations” on page 739 provides an example of when to use this attribute.
Do not mark volumes with Windows Server 2012 application services as critical.	VSS:NSR_SYS_WRITER_WIN32_SERVCOMP_USER=yes Default: No value, which means all application services binaries are treated as part of the DISASTER_RECOVERY:\ save set. For Windows Server 2012, Microsoft has separated System Writer Win32 Services Files by adding a new component for all Windows application services. By default NetWorker will back up the Win32 Services Files as part of the VSS SYSTEM FILESET save set.	Application services binaries listed in the System Writer Win32 Services Files component are saved and recovered as part of a volume back up and not as part of Windows BMR or the VSS SYSTEM FILESET save set. In a disaster recovery scenario, these services must be recovered online after the Windows BMR operation is complete.

NetWorker Support for Microsoft Applications

The NetWorker client supports VSS backups of SQL Server databases that are associated with Windows server roles on computers installed with Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012. For details on this support, see [Chapter 25, “Windows Bare Metal Recovery,”](#) in this guide.

NetWorker Management Console (NMC) does not support VSS backups of Microsoft applications such as Microsoft Exchange, SQL Server, SharePoint Server, Windows Hyper-V, and so on. To fully protect these applications, you must use a product such as the NetWorker Module for Microsoft or another third-party product.

Authoritative restores of the Active Directory Application Mode (ADAM) writers

Authoritative restores of the ADAM writer must be performed from the command line. Restores from the NetWorker User GUI are nonauthoritative.

To perform an authoritative restore of the ADAM writer, use the -U option for the recover command. To recover the ADAM writer, type the following command:

```
recover -s server -U -N "VSS SYSTEM SERVICES:\ADAM (Address Book) Writer"
```


CHAPTER 27

Networking and connectivity

This chapter includes the following sections:

- ◆ Name resolution and connectivity 784
- ◆ Troubleshooting name resolution and connectivity errors 784
- ◆ Using multihomed systems 793
- ◆ NIC Teaming..... 799
- ◆ Using DHCP clients..... 799

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. 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, EDNS0 queries increase the size of the DNS UDP packet and some firewalls block UDP packets larger than 512 bytes. EMC recommends that you disable EDNSprobes on hosts that operate in a firewalled 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.

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 savegroup completion 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 EMC Technical Support.
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 785](#) 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 788](#) 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 791](#) 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 the 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 examples try to connect 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.

In the following example, the source host `mnd.emc.com` is a Linux host with the IP address `10.1.1.10`. The destination host `pwd.emc.com` is a Windows host with the IP address `10.1.1.20`.

1. On the `pwd.emc.com` host, run the following **pathping** commands:

```
pathping pwd.emc.com
pathping pwd
pathping 10.1.1.20
pathping mnd.emc.com
pathping mnd
pathping 10.1.1.10
```

A successful **pathping** command displays the following information:

```
C:>pathping mnd.emc.com

Tracing route to mnd.emc.com [10.1.1.10]
over a maximum of 30 hops:
  0  pwd.emc.com [10.1.1.20]
  1  mnd.emc.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             0/ 100 = 0%      0/ 100 = 0%      pwd.emc.com [10.1.1.20]
  1    0ms    0/ 100 = 0%      0/ 100 = 0%      mnd.emc.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.emc.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             0/ 100 = 0%      0/ 100 = 0%      pwd.emc.com [10.10.10.20]

Trace complete.
```

2. Complete the following steps on the `mnd.emc.com` host:
 - a. Run the following **ping** commands:

```
ping pwd.emc.com
ping pwd
ping 10.1.1.20
ping mnd.emc.com
ping mnd
ping 10.1.1.10
```

b. Run the following **tracert** commands:

```
tracert pdw.emc.com
tracert pdw
tracert 10.1.1.20
tracert mnd.emc.com
tracert mnd
tracert 10.1.1.10
```

Ensure that each **ping** and **tracert** command succeeds. Lost packets can indicate a slow connection between hosts. If any attempt 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 **rpcinfo** to verify the portmapper session establishment

Use the **rpcinfo** 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, start the **portmapper** process that NetWorker uses. “[Fallback to RPC portmapper service on port 111](#)” on page 873 provides more information about how NetWorker uses **portmapper**.

Type the following commands on the source and destination host:

```
rpcinfo -p shortname_of_NetWorker_server
rpcinfo -p FQDN_of_NetWorker_server
rpcinfo -p IP_address_of_NetWorker_server
rpcinfo -p shortname_of_destination_host
rpcinfo -p FQDN_of_destination_host
rpcinfo -p IP_address_of_the_destination_host
```

Note: On Windows, the *NetWorker_installation_dir\bin* contains the **rpcinfo** program.

When the **rpcinfo** command runs successfully, the output displays a list of port numbers and names. For example:

```
rpcinfo for mnd.emc.com
program vers proto port
100000 2 tcp 7938
100000 2 udp 7938
390103 2 tcp 760
390103 2 udp 764
390109 2 udp 764
390107 4 tcp 819
390107 5 tcp 819
```

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. “[Firewall Support](#)” on page 853 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 host tables and the DNS until a cache flush occurs.

A cache flush occurs for the following hosts:

- ◆ For 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.
- ◆ For 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 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 variables 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 operating systems use the following search order: WINS, network broadcast, LMhosts file, hosts file, then DNS. Windows Server 2008 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 returns to NetWorker the first entry in the hosts file that matches the name resolution requirement. Additional instances of an IP address, FQDN, or shortname present in the hosts file for a host are ignored when attempting 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 is operating 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 shortname, FQDN, and IP address.

Perform the following steps on the source host and destination host.

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. Type the following command from a command prompt: **nslookup**
 - 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: EMC recommends 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.emc.com
Address:  10.5.5.10

> server 10.5.5.11
Default Server:  dmd.emc.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: EMC recommends 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 increases the hostname resolution speed by removing the reliance on checking name resolution services for each name resolution request. 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 flush 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 **nsd** 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 that allows you to define the names by which a host is known. When either option contains incorrect hostnames, NetWorker operations can fail when name resolution is correct and there is an established connection 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.
- ◆ 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 NetWorker client start the **save** process to begin a backup. The NetWorker 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 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 shortname 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 services on the host.

Confirming the validity of Aliases attribute

Each NetWorker Client resource contains an Aliases attribute that defines a list of known names associated with the client. The NetWorker server generates this list the when you create a NetWorker client.

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 shortnames 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:

```
hostname is not a registered client
```

Clearing the NetWorker name cache

NetWorker maintains an internal name resolution cache that does not reflect changes that you make to name resolution services until the NetWorker services restart. 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.

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.

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 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

[Table 126 on page 794](#) summarizes how to configure the NetWorker environment to use a multihomed NetWorker server, storage node, and client.

Table 126 Configuring multihomed hosts in NetWorker

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.	<p>The servers 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.</p> <p>Notice: Each instance of the Client resource must have the same value for the Server NetWorker Interface attribute.</p> <p>The Alias field for the NetWorker server Client resource must contain an entry for the shortname and FQDN of each NIC.</p>
	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 nsrmmdbd process on the NetWorker server.	<p>The Server network interface attribute of each Storage Node resource must contain the FQDN of the NetWorker server NIC.</p> <p>The Aliases attribute of the NetWorker server Client resource must contain an entry for the shortname and FQDN of each NIC.</p>
	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 nsrmmgd process.	<p>The Server network interface attribute of Library resource must contain the FQDN of the NetWorker server NIC.</p> <p>The Aliases attribute of the NetWorker server Client resource must contain an entry for the shortname and FQDN of each NIC.</p>
Storage node	The client sends backup data to a NetWorker storage node over a specific NIC.	<p>The Storage Nodes attribute of each Client resource must contain the FQDN of the storage node NIC.</p> <p>Notice: This also applies when the NetWork server is the storage node.</p> <p>The Aliases attribute in the Client resource for the storage node must contain an entry for the shortname and FQDN of each NIC.</p>
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.

Example 66 Configuring NetWorker 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. [Figure 55 on page 795](#) provides a graphical representation of the environment.




Figure 55 Multihomed environment

Complete the following steps to configure the multihomed environment:

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. [Figure 56 on page 796](#) shows the values in the Aliases attribute.

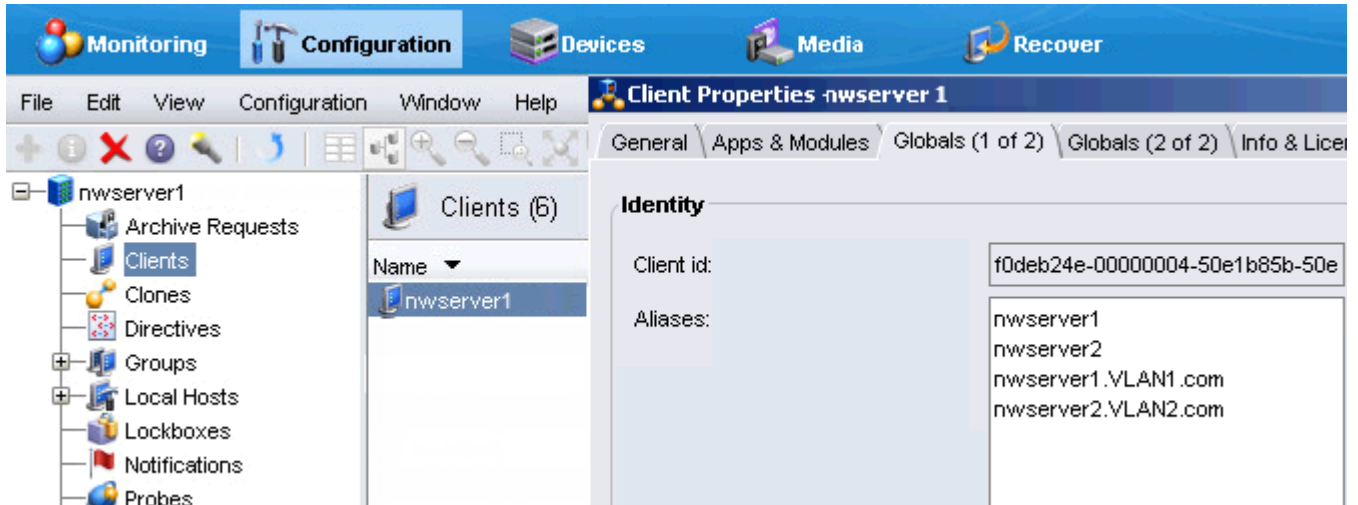


Figure 56 Configuring the Aliases attribute for NetWorker server Client resource

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. [Figure 57 on page 796](#) shows the values in the Aliases attribute.

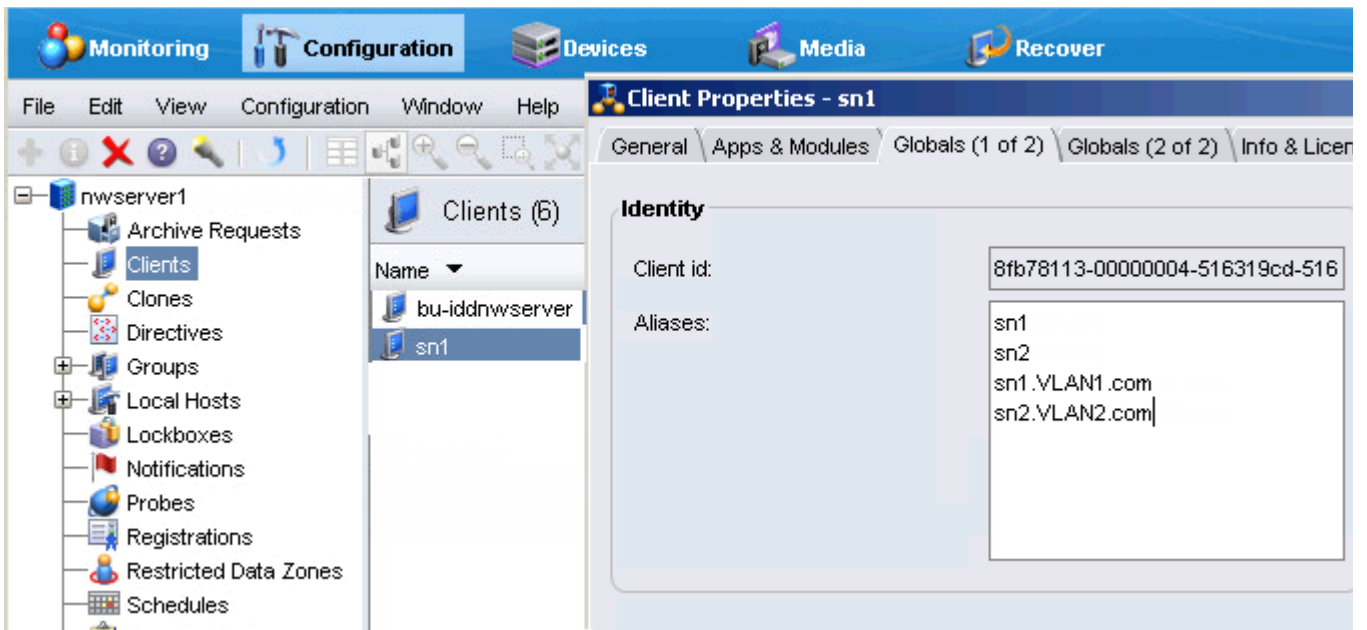


Figure 57 Configuring the Aliases attribute for NetWorker storage node Client resource

- 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. [Figure 58 on page 797](#) shows the values in the Storage node attribute.

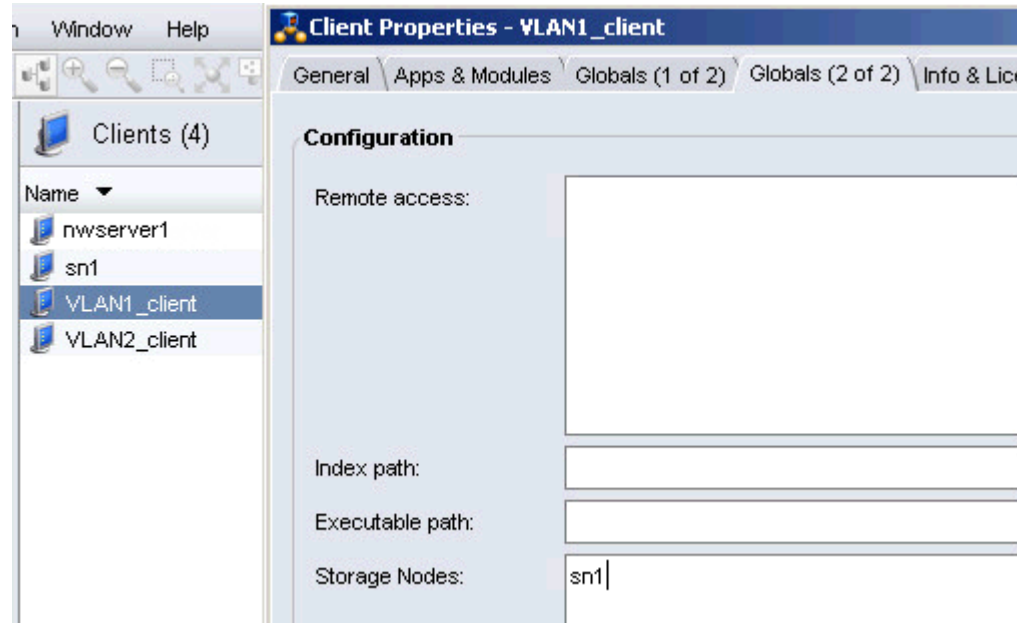


Figure 58 Storage nodes attribute for clients in VLAN1

- 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. [Figure 59 on page 797](#) shows the values in the Aliases and Server network interface attributes.

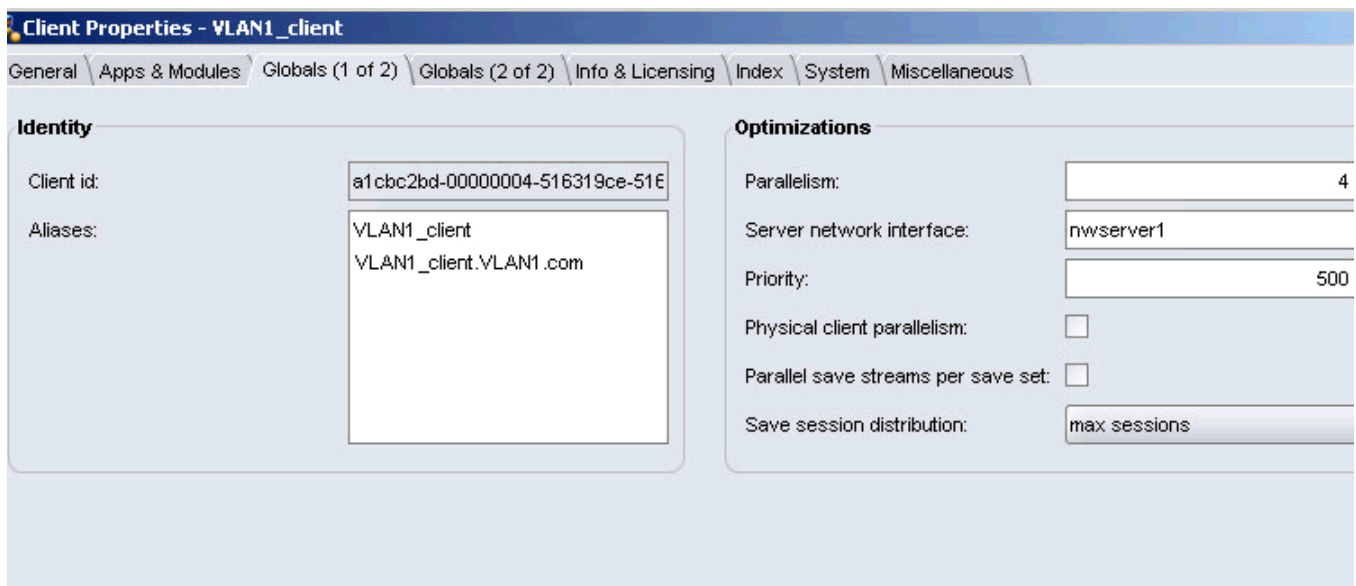


Figure 59 Aliases and Server network interface attributes for VLAN1 clients

- 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. [Figure 60 on page 798](#) shows the values in the Storage node attribute.

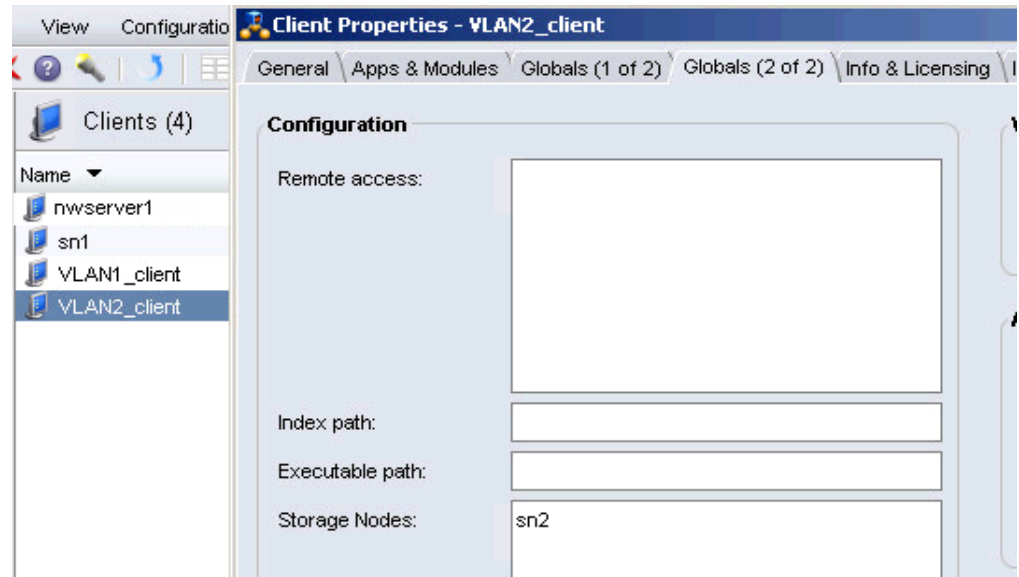


Figure 60 Storage node attribute for clients in VLAN2

- 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. [Figure 61 on page 798](#) shows the values in the Aliases and Server network interface attributes.

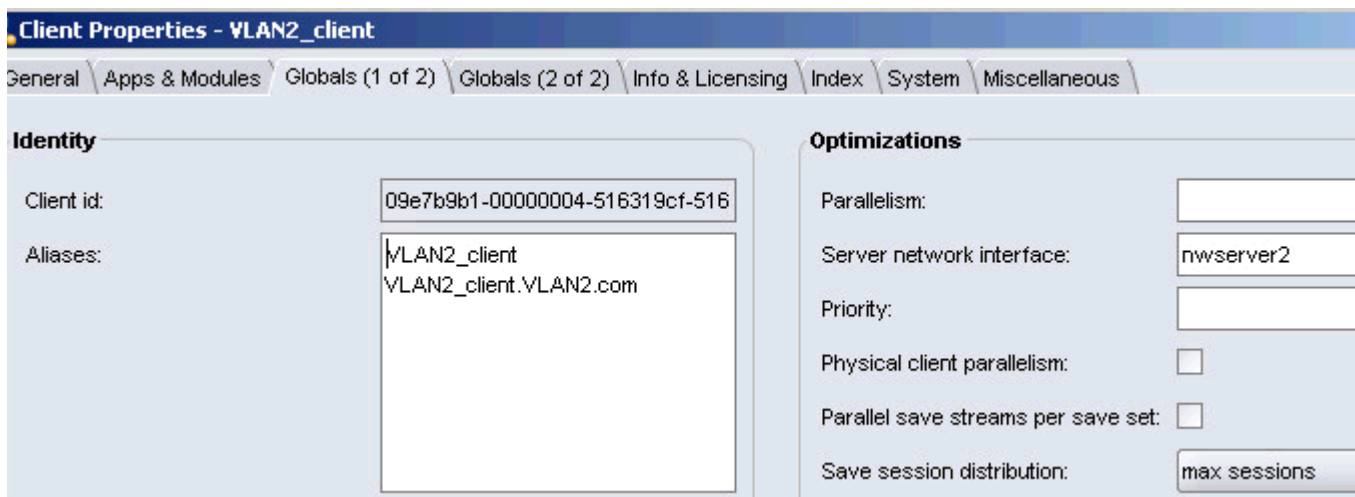


Figure 61 Aliases and Server network interface attributes for VLAN2 clients

- 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, executing 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 resolution for communication between NetWorker hosts. When an IP address changes due to DHCP allocation, 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

EMC recommends 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.

CHAPTER 28

Troubleshooting

This chapter covers these topics:

◆ Before contacting technical support	802
◆ Viewing log files	803
◆ NetWorker functionality issues	807
◆ Devices and Autochangers	820
◆ NetWorker locale and code set support	828
◆ Resource database notes	828
◆ Enabling service mode for NetWorker	829
◆ Network and server communication errors	830
◆ UNIX communication issues	832
◆ Microsoft Windows issues	833
◆ NetWorker archiving and retrieval	834
◆ Storage nodes	835
◆ Console error messages and corrective actions	836
◆ Console log files	838
◆ Console troubleshooting notes and tips	839

Consult the *NetWorker Error Message Guide* for common error messages and possible resolutions.

Before contacting technical support

If the solutions in this chapter do not solve the problem, go to the EMC online support for technical assistance. Provide this information:

- ◆ The software version of the NetWorker component.
- ◆ The 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.
- ◆ The hardware configuration.
- ◆ Information about devices and other SCSI IDs.
To determine this information, use the following commands:
 - For AIX, Linux, and Solaris, enter the **/usr/sbin/inquire** command.
 - For HP-UX, enter the **/etc/ioscan** command.
- ◆ If you are using an autochanger, the type of connection (SCSI or RS-232). Also, provide the version of the autochanger driver you are using:
 - For Solaris, enter the **pkginfo -x** command:


```
# pkginfo LGTOdrv
```
 - For AIX, enter the **lspp -l | grep EMC** command.
- ◆ Be able to supply this information:
 - How to reproduce the problem.
 - The exact error messages.
 - Number of times you have seen the problem.
 - Whether the NetWorker command was successful before you made any changes and, if so, the changes 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

To determine the software version by using the NMC:

1. Connect to the NetWorker server.
2. In the **Configuration** window, select **Clients** from the left navigation pane.
3. Right-click the client and select **Modify client properties**.
4. 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

To determine the software version by using the NetWorker User program on a Windows host:

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.

Note: [Chapter 1, “Overview”](#) provides more information about opening the NetWorker User program.

Determining the client software version by using nsradmin

To determine the client version by using **nsradmin**:

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 version of NetWorker software running on each client is displayed.

Displaying diagnostic mode attributes

NetWorker resources such as clients and devices contain diagnostic attributes that are hidden by default.

To display diagnostic attributes:

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.

Viewing log files

The **nsr_render_log** command renders internationalized NetWorker log files into the current locale of the user who is executing the program. All other log files, as well as messages displayed in the NetWorker Console, use the locale of the service that is generating the log message.

The **nsr_render_log** program is non-interactive. You must specify the log file at the command line when the **nsr_render_log** program is executed. The output of the command is printed to stdout, and can be redirected to a file to save the output. A number of command line options are available with the **nsr_render_log** program as well.

Log files that can be localized using the **nsr_render_log** (UNIX/Linux) or **nsr_render_log.exe** (Microsoft Windows) command include the:

- ◆ Daemon log file — **daemon.raw**
- ◆ NMC server log file — **gstd.raw**
- ◆ NetWorker User log file — **networkr.raw** (Microsoft Windows only)
- ◆ Application Administrators log file — **audit_server_name_sec_audit.raw**
- ◆ Client push log file — **nsrcpd.raw**

Rendering log files in the current locale at runtime

You can also instruct the NetWorker software to render log files into the current locale at runtime, in addition to creating locale-independent log files. This allows you to view log files by using a text viewer.

To instruct the NetWorker software to render logs in the current locale of the machine hosting the file, set the runtime rendered log file in the NSRLA database to the full path of the location for the rendered log file. This must be a valid path.

For backward compatibility with previous releases of NetWorker software, runtime rendered log files do not display all of the fields that are displayed using the **nsr_render_log** program. The runtime rendered log files will contain the message ID followed by the date and time the message was logged, and then the rendered message.

How to render log files in the current locale at runtime

To instruct the NetWorker software to render log files into the current locale at runtime:

1. Log in as root or as Windows administrator on the NetWorker client.

Note: You must have security administrator privileges to view Application Administrators logs.

2. Type this at the command prompt:

```
nsradmin -p nsrexec
```

The **nsradmin** prompt appears.

3. To display a list of all available log file resources:

- a. Type the following at the **nsradmin** prompt:

```
. type: NSR log
```

- b. Next, type the following:

```
print
```

A list of all available log file resources will be displayed.

4. Select the appropriate log file resource for editing by typing the following at the **nsradmin** prompt:

```
. 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. Set the path for the **Runtime rendered log** attribute by typing the following at the **nsradmin** prompt:

```
update runtime rendered log: log_file_location
```

For example, to set the location of the rendered daemon file to the NetWorker log file direction on Microsoft Windows, type the following:

```
update runtime rendered log:
"⟨NetWorker_install_path⟩\nsr\logs\daemon.log"
```

How to view log files with the **nsr_render_log** program

To view log files with the **nsr_render_log** program, execute the following at the command line:

```
nsr_render_log log_file_name
```

If there are spaces in the log file path name, the path and filename should be enclosed in double quotes. For example:

```
nsr_render_log "C:\Program Files\EMC NetWorker\nsr\logs\daemon.raw"
```

Note: The **nsr_render_log** program is located in the bin directory of the NetWorker installation. If the bin directory is not in your search path, you must include the location of the program when executing it from the command line.

The *EMC NetWorker Command Reference Guide* or the UNIX man page provide a complete usage information for the **nsr_render_log** program.

How to redirect **nsr_render_log** output to a file

To redirect **nsr_render_log** output to a file, use the **>** character:

```
nsr_render_log "⟨NetWorker_install_path⟩\nsr\logs\daemon.raw" > mylog.txt
```

You can also save the log file by using a special separator character for export to another program, such as a spreadsheet. To do this, use the **-x exportspec** option, where *exportspec* is a *c* followed by the separator character.

For example, to create a comma-separated list:

```
nsr_render_log -x c, "⟨NetWorker_install_path⟩\nsr\logs\daemon.raw" >
mylog.csv
```

Viewing log files from remote host machines

The **nsr_render_log** program allows you to view log files from remote NetWorker hosts, by using the **-R hostname** option:

```
nsr_render_log -R hostname log_file_name
```

When the **-R** option is used, the log file will be rendered in the locale of the user executing the **nsr_render_log** program, regardless of the locale that is running on the remote host.

Log files from previous releases of NetWorker

For log files generated by clients that are running releases prior to release 7.4, or for preexisting log files that were created before upgrading to from a release prior to release 7.4, do not use the **nsr_render_log** program to view the log files. These log files, which will use the previous naming convention of **.log*, should be viewed by using a standard text editor such as notepad.exe or vi.

Filtering log file information displayed by nsr_render_log

A number of command line options are available for the **nsr_render_log** program to narrow the information output by the program. For example, to view only log file messages referencing a specific device, use the **-D** *devicename* option.

The *EMC NetWorker Command Reference Guide* or the **nsr_render_log** man page provide a complete list of available options.

Viewing only the most recently logged messages

To view only the most recently logged messages in the log file, use the **-B** *beginning_line* option. If *beginning_line* is specified using a negative number, this will instruct the **nsr_render_log** program to display only the specified number of lines from the end of the file.

For example, to display only the last 100 lines from the log file, run the following command:

```
nsr_render_log -B -100 " <NetWorker_install_path>\nsr\logs\daemon.raw" >
  mylog.txt
```

Locating savegroup job logs

If the **Savegroup log by job id** attribute on the NetWorker server resource is selected, you can use the **jobsquery** command to locate logs for child jobs of a savegroup operation. This command takes a query, or query file, and searches the jobs database on the NetWorker server.

Example 67 Using the jobsquery program

This example shows how you could use the **jobsquery** program to locate the child jobs of a savegroup job:

1. Open the **jobsquery** program and use the **show** option to specify which job attributes to display.

```
# jobsquery
show type; command; completion status; start time; end time; job id;
parent job id; job log file
```

2. Use the **print** option to specify that only savegroup jobs will be displayed.

```
print type: savegroup job
```

The output shows that a savegroup with a job ID of 128000 completed successfully:

```
type: savegroup job;
command: ;
completion status: succeeded;
end time: 1228409390;
job id: 128008;
job log file: ;
parent job id: 0;
start time: 1228409364;
```

3. Use the **print** option to display all jobs whose parent job ID is 128008.

```
print parent job id: 128008

type: savefs job;
command: \
savefs -s daphne.lego.com -c daphne.lego.com -g Default -p -l full \
-R -v -F /usr/share/man/man1 /usr/share/man/man3;
completion status: succeeded;
end time: 1228409365;
job id: 128009;
job log file: /nsr/logs/sg/Default/128009;
parent job id: 128008;
start time: 1228409365;

type: index save job;
command: \
"save -s daphne.lego.com -S -g Default -LL -f - -m daphne.lego.com \
-V -l full -LL -W 78 -N
index:c177b9a2-00000004-4936d6d0-4936d6cf-0001c000-69\
7aa04f /nsr/index/daphne.lego.com";
completion status: succeeded;
end time: 1228409388;
job id: 128012;
job log file: /nsr/logs/sg/Default/128012;
parent job id: 128008;
start time: 1228409388;
```

Notice that the job log file attribute in the previous display shows the location of the job logs for two child job IDs: 128009 and 128012.

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information about the **jobsquery** command.

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 properly. 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, enter one of these commands:

```
ps -ef | grep nsr
```

```
ps -ax | grep nsr
```

You should receive a response similar to this:

```
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/nsrmmmd -n 1
12234 ?          S   0:00 /usr/sbin/nsr/nsrmmmd -n 2
12410 pts/8      S   0:00 grep nsr
```

If daemons are not present, start the NetWorker daemons.

Restarting a failed save set

Failed save sets can be restarted without requiring that the entire save group be re-run. You can initiate a restart from the **nsradmin** command line utility or from the **savegrp** program.

NOTICE

Bare Metal Recovery (BMR) enabled clients do not support the restart a single save set. This is because BMR workflows report all save sets within a save group as failure or success.

The following limitations apply to restarting individual save sets:

- Unable to accept requests if the restart window has passed.
- Unable to accept requests for clients with defined pre or post commands.
- Unable to accept requests if backup is in progress.
- Unable to restart a save set that completed successfully.
- Unable to restart a save set that is in progress.

nsradmin

You can use **nsradmin** to restart failed save sets within a previously run save group.

From the command line type:

```
# nsradmin

nsradmin> . type: Nsr group; name: GroupName|Default

nsradmin> update client subset: client1:ss1,ss2,
client2:ss3,ss4; autorestart: restart now
```

The response should be similar to:

```
update client subset: client1:ss1,ss2;client2:ss3,ss4; autorestart:
restart now
```


savegrp

The **savegrp** program enables you to restart failed clients or save sets while a group is running.

From the command line type:

```
savegrp -R -c "client1:ss1,ss2;client2:ss3,ss4" GroupName
```

Known limitations when using the **savegrp** program to restart individual save sets:

- **savegrp** does not accept requests if it is already saving the bootstrap.

Client wizard issues

Improper font size for the Client Wizard with Netscape on Solaris

When using the Netscape browser on Solaris, the font size of the Client Wizard may be too small.

To change the font type and size:

1. Open the `/usr/bin/nwwiz` script 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.

Backups fail to start when daylight savings time change occurs

If backups are scheduled to occur during the hour in which the operating system moves the clock ahead or behind by one hour, the backup operation will be skipped. For example, suppose that 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 clock is moved forward to 3:00 A.M. All times from 2:01 to 2:59 are skipped and scheduled is not initiated.

To avoid this situation, set the backup time to occur at least one minute before the time change occurs.

Note: Using the `mminfo` query to get a weekly save set usage summary during the change to daylight savings time does not display any information for the day of the change.

Shut down NetWorker services prior to any significant changes to system date

If a significant change needs to be made to the system clock/date (for example, a change of more than a day), ensure that NetWorker services have been shut down prior to making the change. NetWorker services are heavily dependent on the system clock for operations such as active sessions, mounting and unmounting of volumes, expiration of save sets and licenses, and so on.

Clone ID timestamp does not reflect the time the clone was created

To guarantee that cloned save sets created 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 the clone was created.

Backups fail to stop

Attempting to stop the backup process by clicking Stop in the Group Control window should stop the process for all clients in the selected group. However, sometimes a client is missed and messages appear indicating that the server is still busy.

To resolve the problem:

1. From the **Administration** window, click **Monitoring**.
2. Select the **Groups** tab and determine which group is currently being backed up by looking at the messages that display.

If the group status shows that the **save** processes are running, but the associated **savegrp** process is not running, perform *one* of these:

- ◆ Stop the conflicting group from running by clicking **Stop** in the **Group Control** window. [“Stopping a group” on page 466](#) provides more information.
- ◆ Shut down and restart the NetWorker services. [“Stopping and starting a NetWorker server, client, or storage node” on page 55](#) provides more information.

Memory usage when browsing large save sets

Browsing or recovering from a large save set, such as a save set with one million or more files, may consume all of the host’s memory. The workaround is to perform a save set recovery instead. [“Recovering the data” on page 373](#) provides information on save set recovery.

The **recover** command 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, **nsrmmmd** must reposition the volume to read previously written data. It does not always succeed on the first attempt. These 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 media verification fails:

- ◆ Reset the device.
- ◆ Verify the configuration of the device.
- ◆ Verify that the media can be recognized.
- ◆ Verify that the device is functioning properly.

PACKET RECEIVE BUFFER and NO ECB counters increase

When the server is waiting for a tape to be mounted or is in the process of changing an autochanger volume, the PACKET RECEIVE BUFFER and NO ECB counters increase on a NetWare client.

To resolve this problem, shut down and restart the NetWorker server.

For servers that run on HP-UX, edit the `/sbin/init.d/networker` file. Add this line before the line that starts **nsrd**:

```
NSR_NO_PING=ok; export NSR_NO_PING
```

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 the last save set on the backup volume from being overwritten.

To write to the media without marking it read-only, use the **nsrmm -o** command:

```
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, this message may appear:

```
Please enter record size for this volume ('q' to quit)
```

If this message appears, enter the block size. The block size must be an integer equal to or greater than 32.

Limitations for groups containing a bootstrap

Backups for a group that generates a bootstrap file can be written to a storage node only when a tape from the default pool is already labeled and mounted on a local drive attached to the NetWorker server.

Index recovery to a different location fails

If you attempt to recover indexes to a directory other than the one where they were originally located, this error message 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.
```

Recover indexes to their original location before moving them to another directory. To move the indexes, log in as root and invoke this command from within the /nsr/index directory:

```
uasm -s -i "client_index_directory_name" | (cd target_directory; uasm -r)
```

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.

Illegal characters in configurations

When naming label templates, directives, groups, policies, and schedules, these characters are not allowed:

```
/ \ * [ ] ( ) $ ! ^ ' " ? ; ` ~ < > & | { }
```

Error backing up large number of clients

Backing up a large number of clients may cause this CMD.exe application error message to appear on the NetWorker server:

```
The application failed to initialize properly (0xc0000142). Click on
OK to terminate the application.
```

If this problem occurs, increase the desktop heap allocation by editing the following Windows registry key on the NetWorker server:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\
  Session Manager\SubSystems\Windows
```

In the following example, the desktop heap allocation has been changed from a value of 512 KB to 1023 KB.

Previous version, with a desktop heap allocation of 512 KB:

```
%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
```

Updated version, with a desktop heap allocation of 1024 KB:

```
%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
```

The Microsoft Knowledge Base article 18480 on the Microsoft website provides more information.

Hostname aliases

Savegroups fail when aliases are improperly defined. 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, every client should have both its hostname and its fully qualified domain name listed in its NetWorker aliases field.

If you encounter any of these situations, a client alias problem might be the cause:

- ◆ This error message appears:

```
No Client resource for client_name
```

- ◆ A client machine always performs full backups, regardless of the level of the scheduled backup.
- ◆ Automatic index management, as set up in the browse and retention policies, does not work.
- ◆ In the `/nsr/index` directory, which contains the indexes, there are two directories for the same client that use two different client names.

A client alias change is needed in the following situations:

- ◆ Machines have two or more network interfaces.
- ◆ Sites mix short and fully qualified hostnames for the same machines, for example, `mars` and `mars.jupiter.com`.
- ◆ Sites use both (Network Information Services (NIS) and DNS).

Add all network names for the host to the Aliases attribute in the Client resource.

NOTICE

Do not include aliases that are shared by other hosts in the Aliases attribute.

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 pathnames with fewer than 128 characters.

Backup of a new client defaults to level full

The first time you back up a new client, this message appears:

```
client: save point: There are no save sets in the media database;  
performing a full backup
```

This message indicates that the specified save set has not been previously backed up.

Before you can perform an incremental or level backup on a save set, perform a full backup of the save set.

If the save set was previously backed up, the reasons this message appears include:

- ◆ The clocks on the client and server are not synchronized.

- ◆ The **savegrp** session begins before midnight and ends after midnight.

Non-full backup of Solaris files with modified extended attributes

When the extended attributes for a Solaris file are changed, but the file is not otherwise modified, the change time (ctime) for the file is not updated. 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 70](#) provides more information.

Renamed clients cannot recover old backups

The NetWorker server maintains a client file index for every client it backs up. If you change the name of the client, the index for that client is not associated with the client’s new name and you cannot recover files backed up under the old client name.

To recover data that was backed up by using the old client name, perform a directed recovery by directing data saved under the old client name to the new client. [“Directed recoveries” on page 369](#) provides information about performing directed recoveries.

Client file index errors

These issues are related to client file indexes:

- ◆ [“Missing client file indexes” on page 814](#)
- ◆ [“Check failure of client file indexes” on page 814](#)
- ◆ [“No notification of client file index size growth” on page 815](#)

Missing client file indexes

The **scanner** program must have a client file index to rebuild from before it can proceed. If you attempt to recover a client file index with the **scanner -i** command without first using **nsrck -L2** to create a new client file index, a message similar to the following could appear:

```
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
```

Check failure of client file indexes

Each time the NetWorker server starts, it uses **nsrck -ML1** 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 an index might be corrupt, run a higher level check on the index, for example:

```
nsrck -L5
```

If the index is still corrupt, refer to [“Recovering expired save sets” on page 398](#) 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. [“Reducing client file index size” on page 591](#) provides information on how to manage the NetWorker client file indexes.

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide more information on the **nsrls**, **nsrck**, and **nsrim** commands.

Cannot use the Console interface to stop the savegrp command

If you start the **savegrp** command at the command prompt, and then attempt to stop the backup from the Console window, this message appears:

```
Only automatically started groups that are currently running can be
stopped
```

Manually stop the **savegrp** process.

Aborting a recovery

When you stop a recovery in progress on a client, the following could occur:

- ◆ The recovery might stop immediately.
- ◆ The files that still need to be recovered are listed.
- ◆ Messages similar to this might appear:

```
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
```

The messages indicate that a recovery was not stopped cleanly.

RPC error

If NetWorker has trouble backing up a directory path, a message similar to this appears, which notes the path:

```
* jupiter:E:\ save: xdr of win32 attributes failed for 'E:\PROGRAMS\'
```

The rest of the save set completes successfully.

To solve this problem, perform another backup of the directory.

Error message when relocating data

If you attempt to relocate data to a directory that does not exist, this error message appears:

```
Cannot create directory directory
```

Ignore this message. The recovery creates the new directory and completes successfully.

Desktop heap size limitation

Microsoft Windows XP has a set desktop heap size limitation that might produce the following error message when exceeded:

```
The application failed to initialize properly
```

The Microsoft Knowledge Base article 142676 on the Microsoft website provides information about this problem and how to correct it.

Other failures can also cause the desktop heap size to be exceeded. If this occurs, it is usually the result of numerous NetWorker processes that are running simultaneously. To determine the number of NetWorker processes that are running on the server, use the Windows Task Manager.

These conditions can cause the NetWorker server to exceed the desktop heap size limit:

- ◆ Large number of system services are running as Local System.
- ◆ More than 30 backup devices exist on the NetWorker server.
- ◆ NetWorker parallelism is set above 30.
- ◆ Performing more than 15 simultaneous recoveries.

If any of these failures occur, or if the “Application failed to initialize properly” message appears, increase the desktop heap size for system services. To do this, modify the third parameter to SharedSection as outlined in Microsoft Knowledge Base article 142676 on the Microsoft website. Increasing the size by 3072 usually corrects the problem. For example:

```
%SystemRoot%\system32\csrss.exe ObjectDirectory=\Windows
SharedSection=1024,3072,3072 Windows=On SubSystemType=Windows
ServerDll=basesrv,1 ServerDll=winsrv:UserServerDllInitialization,3
ServerDll=winsrv:ConServerDllInitialization,2 ProfileControl=Off
MaxRequestThreads=16
```

The All save set and duplicate drive serial numbers

The All save set, which backs up all locally mounted drives as well as the SYSTEM and VSS SYSTEM save sets, uses the serial numbers assigned to drives as part of its logic to determine if a drive should be backed up. It is possible for local drives to use the same serial number. In this case, the All save set results in only one of the drives being backed up.

If you encounter this problem, there are two possible 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 and SYSTEM or VSS SYSTEM save set separately. [“Scheduling predefined save sets for backup” on page 66](#) provides more information about the All save set.

Disk label errors

If a nonoptical device was configured as an optical device, this error message appears:

No disk label

Verify that the Media Type attribute in the Device resource matches the expected media for the device, and correct if necessary.

Cannot print bootstrap information

If the server bootstraps do not print, enter the printer's name when configuring the Group resource:

1. In the **Administration** window, right-click the group and select **Properties**.
2. In the **Printer** attribute of the **Setup** tab, enter the name of the printer where the bootstrap is to print.

Server index not forced

If the NetWorker server belongs to a group that is disabled or if it does not belong to any group, the **savegrp** program does not back up the NetWorker server.

The information to recover server indexes is stored in the media database on the NetWorker server.

Copy violation

If NetWorker software is installed on multiple servers and the same NetWorker enabler code was used for them all, messages similar to this appear in the save group completion email:

```
--- 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 successfully rerun the backup:

1. Issue the **nsr_shutdown** command on each server.
2. Remove the NetWorker software from the extra servers.
3. Restart the NetWorker services on the server where the backups are to go.

Converting sparse files to fully allocated files

The NetWorker server determines that 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, the file is considered to be sparse and is saved such that long strings of zeroes are replaced with "holes" in the recovered file.

Some files that were not sparse when saved might be recovered as sparse. Oracle databases are susceptible to this problem because they are zero-filled, fully allocated files and are not sparse.

To workaroud this issue, use the **cp** command to copy the file after recovery:

```
cp recovered_filename zero_filled_filename
```

This converts a sparse file to a fully allocated file.

NOTICE

Ensure that you have enough free disk space to accommodate a duplicate of each sparse file that is copied.

Backing up large sparse files

To conserve backup media, sparse files are compressed before being written to tape. During this time, the backup job may stop with this message:

```
savegrp: Aborting inactive job (633).
```

This can occur because no data is being written to the backup media while the sparse file is processed. 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 required to complete a full save of the file system.
3. Use this time as the inactivity timeout limit. [“How to edit a group” on page 253](#) provides information about setting the inactivity timeout attribute for the group.

The **mminfo -N** command is case-sensitive regarding save set names

When querying the media database by using the **mminfo** command, the **-N name** option is case-sensitive. The save set name the **-N** option references must match the case of the save set name entered in the Client resource.

However, when backing up drive partitions on Microsoft 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 entered in the Client resource in lowercase, you must query by using uppercase:

```
mminfo -N C:\
```

Renamed directories and incremental backups

If the name of a directory is changed after a full backup, but no files or subfolders in the directory were changed, the renamed directory is not included in subsequent incremental backups.

To avoid this issue, select the Backup renamed directories attribute on the Client resource.

Resolvable names for multiple network interface cards

If any component of NetWorker (client, storage node, server) has multiple network interface cards (NICs) with unique IPs and hostnames, all NICs must be configured and must be resolvable names, even if one or more NICs are not being used. Failure to have all NICs resolvable may cause problems with host connectivity to the NetWorker server.

Follow these steps to configure NetWorker so that the appropriate hostname is used for the associated IP, and to ensure the hosts file and routing table on the machine are configured properly:

- ◆ Set up DNS so that a separate name is associated with each IP
- ◆ Configure the hosts file and routing table on each machine that has multiple interfaces with the appropriate IP
- ◆ Configure NetWorker to use the names configured in steps 1 and 2.

Example for configuring multiple NICs

In the following example, a dual-interface client connects to the NetWorker Server and Storage Node over **interface1** having IP **1.1.1.1** and has a dedicated connection to the Storage Node over **interface2** having IP **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 hostnames for IPs **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 hostnames 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 add the two IPs to the local hosts file.
3. In NetWorker, enter **node-2** in the Storage Node Affinity List of the client.

Libraries entering ready state

When starting NetWorker or after configuring a library, it may take a short amount of time for the library to enter the Ready state within NetWorker. This is normal behavior.

Improper font size for the Client Wizard with Netscape on Solaris

When using the Netscape browser on Solaris, the font size of the Client Wizard may be too small.

To change the font type and size:

1. Open the `/usr/bin/nwwiz` script 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.

Successful save sets listed as failed in the Group Backup Details window

Certain backup operations, such as the direct SCSI feature and some NetWorker modules, create multiple sessions during a single backup job. If one of these sessions fails, the Console will report the entire backup job as having failed.

To determine the status of each session, click the **Show Messages** button in the Failed table of the **Savegroup Completion** dialog. This information is also available 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 if the **RPC ping via UDP when connecting to NetWorker** check box is selected in the NetWorker Console **Setup > Systems Options** dialog box is checked and the NetWorker server window does not appear:

```
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.

Devices and Autochangers

This section explains how to resolve problems with devices and autochangers.

NOTICE

Device files and directories should not be edited. Editing these files can cause 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 options that the **nsrjb** program uses. These are hidden attributes. [“Displaying diagnostic mode attributes” on page 803](#) provides information about displaying hidden attributes.

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide 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**, for diagnosing problems on tape devices and autochangers.

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

Autodetected SCSI jukebox option causes server to stop responding

If an autodetected SCSI jukebox is installed using **jbconfig**, and the server stops responding:

1. Select the **jbconfig** option that installs an SJI jukebox.
2. Enter the number that corresponds to the type of jukebox you are installing.
3. Proceed with **jbconfig** until this message appears:

```
Jukebox has been added successfully.
```

Autochanger inventory problems

These situations cause the autochanger inventory to become outdated:

- ◆ The media is manually ejected from the autochanger drive.
- ◆ The media is removed from the autochanger.
- ◆ The autochanger door is opened.

An outdated inventory means that the NetWorker software cannot use the autochanger.

To make the autochanger usable again:

1. Verify that the media cartridge 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 by typing this command:

```
nsrjb -Hv
```

4. Perform an inventory by typing this command:

```
nsrjb -Iv
```

The NetWorker server can use the autochanger after the inventory operation completes.

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide complete details on the **nsrjb** command.

Destination component full messages

If a manual operation is performed on an autochanger, for example unloading the tape drive by using the buttons on the autochanger rather than by using the NetWorker server, this error message may appear:

```
Destination component full
```

To resolve the problem, use the **nsrjb -H** command to reset the autochanger.

Tapes do not fill to capacity

Tapes may not always be filled to capacity. For example, a tape with an advertised capacity of 4,000 MB can be marked full by the NetWorker server after only 3,000 MB of data have been written to it.

To use the tape to its fullest capacity, select the highest density device driver for the device. 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:
 - The device driver is properly configured.
 - Any necessary switch settings on the tape drive are set to the manufacturer's specifications.
 - All cables are secure.
 - Other potential SCSI problems have been addressed.
- ◆ NetWorker filemarks consume space on the tape.
 - The NetWorker server periodically writes filemarks to facilitate rapid recovery of data. These filemarks consume varying amounts of tape depending on the type of tape drive.
 - The number of filemarks the server writes to tape depends on how many save sets are on the tape. Many small save sets require more filemarks than a few larger ones.
- ◆ Tape capacities vary.

Two apparently identical tapes from the same vendor can vary significantly in capacity. This can cause problems if 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 noncompressing drive.
 - The capacity could vary depending on the type of data being backed up. For example, if a noncompressing 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.
- ◆ Length of tape. Verify tape lengths. A 120-meter DAT tape holds more data than a 90-meter DAT tape.

Tapes get stuck in drive when labelling on Linux Red Hat platform

When labeling a tape in a DDS configuration using a NetWorker server that is running Linux Red Hat, the tape may become stuck in the drive and display the following error message:

```
unload failure-retrying 30 seconds
```

To prevent a tape from being stuck in the drive, set the `auto_lock` setting to “0” (Off) in the `/etc/stinit.def` file for these 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 may take more than 30 minutes before it fails under these conditions:

- ◆ Automeia management is enabled and a backup is initiated, and
- ◆ The NetWorker software encounters a corrupted tape during label operations.

The NetWorker software keeps a record of the location of the corrupted tape only for the current backup operation, so a corrupted tape could be used again for the next backup operation if the operator does not remove it.

To increase the value of the Save Mount Time-out attribute to 60 minutes from the default 30 minutes:

1. In the Administrator program, select **Devices** from the Media menu to open the **Devices** window.
2. From the **View** menu, select **Details** to display the hidden attributes.
3. Set the **Save Mount Time-out** attribute to 60 minutes.

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 whether the control port is properly connected.

If you cannot determine whether the control port is working, contact the autochanger vendor for assistance.

Modifying the control port

A change in the control port of the robotic arm of a library is characterized by the inability to perform library operations, such as labeling, mounting and unmounting, and inventorying. You may see the error "no such file or directory."

To update the NetWorker to use the new control port:

1. Run the **inquire** command to determine the SCSI device address of the library arm and to confirm that a serial number is reported. If a serial number is not reported go to [step 5](#).

NOTICE

Use the inquire command with caution. Running inquire sends the SCSI inquiry command to all devices detected on the SCSI bus. Using inquire during normal operations may cause unforeseen errors and possible data loss may result.

2. If the serial number of the arm is reported, follow the procedure at [“Scanning for libraries and devices” on page 142](#) to scan the library for devices.
3. Click **Monitoring**, then click to **Logs** tab and locate the message:


```
media info: The control port of the disabled library 'library_name'
has been changed to 'scsidev@b.t.l' on storage node
'storage_node_name'!
```
4. Enable the library
 - a. In the **Administration** window, click **Devices**.
 - b. Expand the Libraries folder and right-click the library and select **Enabled/Disable**.
5. If the serial number was not reported in [step 1](#), or if scanning for devices 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 client.
 - b. Enter the **nsradmin** command at the command-prompt. The **nsradmin** prompt appears.
 - c. Disable the library by typing the following at the **nsradmin** prompt:
 - type: **NSR jukebox**
 - update enabled: **no**
 - When prompted to update the resource, type **yes**.
 - d. Update the control port by typing the following at the **nsradmin** prompt:


```
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).

When prompted to update the resource, enter **yes**.

- e. Reenable the library:
 - update enabled: **yes**
 - When prompted to update the resource, enter **yes**.
- f. To verify that the control port was changed and the library is now enabled, enter **print** at the **nsradmin** prompt.

Nonrewinding device requirement

Use a nonrewinding device for NetWorker backups. The NetWorker server writes a filemark on a volume at the end of each backup. When the next backup occurs, the server appends data to the volume based on the position of the filemark. If the device automatically rewinds the data, the filemark position is lost and the data is overwritten by the next backup.

Scanner command behaves differently with `adv_file` type device

The **scanner** command behaves differently when used with an advanced file type device. When both primary and `_AF_readonly` `adv_file` type devices are unmounted, the following command results in the `_AF_readonly` device being mounted:

```
scanner -m -S ssid primary_device_name
```

This is expected behavior.

Sleep times required for TZ89 drive types

If you are unloading a TZ89 drive and receive the following error, your drives require changes to the sleep attributes in the Autochanger resource.

```
nsrd: media info: unload retry for jukebox `COMPAQTL895' failed - will
  retry again.
```

To change the sleep attributes:

1. Shut down NetWorker services.
2. Shut down and restart the autochanger with the TZ89 drives.
3. When the autochanger is back online, restart NetWorker services. This resets NetWorker so that it stops trying to unload the drive.
4. Use these settings for the sleep time attributes:
 - Eject Sleep: **18** secs
 - Unload Sleep: **40** secs
 - Load Sleep: **40** secs

“[Additional attributes in the Autochanger resource](#)” on page 820 provides information about setting the sleep attributes.
5. Attempt 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 disk FTD

If the CDI feature is enabled while using an NDMP tape device or file type device (FTD), a message similar to this appears in the NetWorker message log:

```
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.

Verifying firmware for switches and routers

If switches or routers are used, make sure that any switch or router firmware on the network was manufactured after August 1995 to ensure that RPC traffic is handled properly. 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

Commands may fail when issued to a NetWorker server or storage node that has multiple network interface cards (NIC).

To prevent this failure, add the domain name of each additional NIC to the **Aliases** attribute in the **Client** resource that is set up for the NetWorker server or storage node. [“Editing a client” on page 606](#) provides information about editing a Client resource.

SCSI reserve/release with dynamic drive sharing

When the NetWorker software uses Dynamic Drive Sharing (DDS) there is a possibility that the operating system's tape driver might use the SCSI reserve/release feature in a manner that interferes with the proper operations of the NetWorker software. This may require that reserve/release be disabled.

To disable the reserve/release feature for the various operating systems:

Solaris

SCSI reserve/release is configurable as a bit setting in the `st.conf` file for each device type in use. 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:

- ◆ DDS is used with the NetWorker software.
- ◆ A tape drive is used that is not supported directly by a Solaris `st` tape driver.

To determine if the tape drive is supported directly by a Solaris `st` tape driver, load a tape in the drive and enter the `mt` command. For example, with the tape device file `0cbn`, type the following:

```
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 is using generic settings for that drive and it is *not* natively supported:

```
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
```

If this configuration is used with the NetWorker software, the process might appear to work, but there might be problems recovering any saved data.

If the output of the **mt** command appears similar to the following, the **st** tape driver recognizes the drive and is using correct internal settings:

```
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
```

The only reason to edit the **st.conf** file is if the drive is being used in a DDS configuration.

AIX

To reset the reserve/release setting on an AIX operating system:

1. Through the **SMIT** interface, select **Tapes** from the **Devices** menu.
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:

1. Change the **st_ats_enable** kernel variable to a value other than zero.
2. (Optional) Restart the computer to ensure the change was implemented.

Note: The reserve/release is a fixed setting in HP-UX 10.

Device ordering issues

[“Device ordering” on page 199](#) provides information about issues related to drive ordering, including how to determine if drive reordering has happened, and procedures to correct the problem.

Recovery of save sets from a VTL

Loading a tape to recover save sets after moving the tape from one VTL of a NetWorker server to a VTL of another NetWorker server

When it is required to load a tape to recover save sets after moving the tape from a VTL of the source NetWorker server to a VTL of a different NetWorker server, the following procedure provides information about how to load the tape prior to running the **scanner** command to restore the backed-up save sets to another NetWorker server, without requiring a NetWorker mount operation:

1. Ensure the destination VTL is the same model, has the same drive names and the same number of drives as the original VTL.
2. Check the inventory of the VTL in the destination NetWorker node
3. Run the **inquire** command to get the Control port of the VTL in the destination NetWorker node.
4. Run the **sjimm** command to load the tape to the corresponding drive of the destination NetWorker server.
5. Run the following command to determine the tape status:

```
mt -f <device> status
```

After ensuring that the tape has been moved to another VTL, run the **scanner** command to restore the backed up save sets.

NetWorker locale and code set support

NetWorker software does not support locales (defined by the operating system) or code sets that remap characters having special meaning for file systems. Depending on the file system, these special characters may include the slash (/), the backslash(\), the colon (:), or the period(.). De_DE.646 is an example of one unsupported locale.

NetWorker software might not function normally if the locale is changed. The previously existing indexes can become invalid.

Resource database notes

The NetWorker resource information resides in directories on the host with the following structure:

```
<NetWorker_install_path>\res\nsrdb\00
```

```
<NetWorker_install_path>\res\nsrdb\09
```

NetWorker stores each resource in a separate numbered file. As new resources are created (for example, Client, Group, or Pool resources), new files are added in these directories.

Note: Because Client resources are generally small, the NetWorker client (**nsrexecd**) continues to use the <NetWorker_install_path>\res\nsrla.res file.

Viewing resources

You can view all NetWorker resources through the Administration window.

Although you can view the contents of the new NetWorker resource files with a text editor, direct user edits are not supported.

The only supported access to the resource database is through either of the following:

- ◆ `nsradmin -s server`
- ◆ `nsradmin -d NetWorker_install_path\res\nsrd`

If you inadvertently specify the wrong path with the `nsradmin -d` command, empty resource directories are created. If this occurs, delete the incorrect directories.

Repairing resource database corruption

Corruption of NetWorker resource database files can be caused by a power outage, operating system crash, or manual editing of the database. If the NetWorker server is unable to read the resource files upon startup, messages similar to these are written to the daemon log file:

```
nsrd: WARNING: NSR configuration database detected invalid
resource ...\00019803aa14713c89456b41
nsrd: Invalid resource saved at ...\00019803aa14713c89456b41
```

[“Viewing log files” on page 803](#) provides information about viewing log files.

The NetWorker server removes any invalid resource files from the nsrdb directory structure and places them in the dbg directory. The dbg directory is created only if resource database file corruption has occurred. If you encounter this problem, open the corrupt file with a text editor to determine which resource is corrupted. You can then re-create the resource using either the Console window or the `nsradmin` command.

After you inspect a corrupt resource file, delete it.

Note: If you do not know the cause of the resource file corruption, go to the EMC online support for technical assistance.

Enabling service mode for NetWorker

Two attributes, **Accept new sessions** and **Accept new recover sessions**, are available in NMC for enabling and disabling access to the NetWorker server. Unselecting these attributes prevents the server from accepting new backup and recovery sessions.

By restricting NetWorker server access, you can take all storage nodes offline, effectively putting NetWorker into a service mode operational state where you can stop any external client requests from being accepted, or stop groups from starting automatically. Putting the server into this state provides a maintenance period where you can diagnose and troubleshoot issues before returning the server to normal operation. The section [“Restrict backup and recover access to the NetWorker server” on page 568](#) provides more information on these attributes.

You can also enable/disable specific storage nodes or devices to prevent use and allow for service operations. The section [“Configuring storage nodes” on page 133](#) provides information on how to enable/disable specific storage nodes. The section [“Re-enable a device” on page 236](#) describes how to enable/disable a specific device.

Network and server communication errors

This section provides general, UNIX and Windows network and communication issues that you may encounter in a NetWorker environment.

To help ensure successful communication between NetWorker clients and servers, each host configured in NetWorker must *not* have any invalid or inactive IP addresses stored in the hostname resolution service used (DNS, NIS, Active Directory, hosts file, and so on). Each address mapped to a host must have a configured network interface (NIC).

General issues

This section provides information that may be relevant for multiple platforms.

Unapproved server error

If an unapproved server attempts to contact a client to initiate a backup, this message appears:

```
client_name: server_name cannot request command execution
```

1. After installation, if the client is to accept backup requests from other NetWorker servers, add the NetWorker server names to the *servers* file.
2. Ensure the servers file on a client contains both the short name and the long name of the server to use to back up that client's data. For example, the servers file on a NetWorker client should contain these names for a NetWorker server named *mars* in the *jupiter.com* domain:

```
mars
mars.jupiter.com
```

3. In the Alias attribute of the Client resource, list both the short name and the long name, plus any other applicable aliases for each client.
4. The preferred method of editing the *servers* file is to run the NetWorker Setup program in maintenance mode and edit the Allowed Servers list. The *NetWorker Installation Guide* provides details.

Unapproved server error during client setup

If you add a Windows client to a UNIX NetWorker server, and the UNIX server hostname is not included in the Windows client's *servers* file, you might receive this error message:

```
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 eliminate the message, add the UNIX server hostname to the servers file on the client after you finish adding the client to the UNIX server.

Server copy violation

Add all server aliases that are related to any additional network interfaces to the alias list of the NetWorker server.

If aliases are not recognized, the server may be disabled with this error:

```
nsrd: registration info event: server is disabled copy violation
```

Remote recover access rights

You can control client recover access through the Client resource. The Remote Access attribute displays the users that have recover access to the client's save sets. Add or remove usernames depending on the level of security the files require.

Note: If you enter a hostname or `host=hostname` in the Remote Access attribute, any user on that host is allowed to recover the client's files. To enter a username without specifying the host, enter `user=name`.

These users have permission to recover any files on any client, regardless of the users listed in the Remote Access attribute:

- ◆ 'root' user on a Unix host
- ◆ Member of the 'Administrators' local group on a MS 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 recover only files for which they have read permission, based on file permissions at the time that the file was backed up. Files recovered by a user other than root, operator, or the operator group are owned by that user.

Authentication fails due to duplicate hostnames

Authentication with the NetWorker server may fail if multiple NetWorker hosts share the same short name. For example, suppose hosts from two domains, *accounting.company.com* and *marketing.company.com*, are configured for backup on the same NetWorker server. However, each domain has a host named *jupiter*. In this case, authentication may fail when the second host named *jupiter* attempts to contact the NetWorker server.

To enable the host to authenticate to the NetWorker server:

1. On the NetWorker host that cannot authenticate, stop the NetWorker client service. ["Stopping and starting a NetWorker server, client, or storage node" on page 55](#) provides more information.
2. Delete the **nsrladb** database, which is located in *NetWorker_install_path*\res\nsrladb.

NetWorker server takes a long time to restart

If the NetWorker media management database is very large, the NetWorker server may take a long time to establish client connections when it is restarted. The reason is that a consistency check of the media management database is triggered when the server is restarted.

To reduce the size of the media management database, run the `nsrim -C` command. 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.

[“Reducing media database size” on page 591](#) provides more information about reducing the media management database.

Changing the NetWorker server address

When the TCP/IP address changes on the NetWorker server, the NetWorker host ID also changes, which invalidates the authorization code. In this case, reregister the software. To reregister the software, print out the registration form with the new host ID and return it to EMC Customer Service. If you do not reregister the software within 14 days, the NetWorker software stops working.

If you are using DHCP, use a static IP address for the NetWorker server.

UNIX communication issues

This section covers communication issues on UNIX networks.

Binding to server errors

NetWorker architecture follows the client/server model, where servers provide services to the client through the RPC. 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 a NetWorker service is requested, this messages appear in the savegroup completion email:

```
Server not available
RPC error, no remote program registered
```

These messages indicate that the NetWorker services `nsrd`, `nsrexecd`, `nsrindexd`, `nsrmmmd`, and `nsrmmdbd` might not be running.

Table 127 NetWorker Startup commands

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

Saving remote file systems

This error messages might appear in the save group completion email when a backup for a remote client fails:

```
Host hostname cannot request command execution
hostname: Permission denied
```

The first message means that the **nsrexecd** service on the client is not configured to allow the server to back up its files. The second message means that the **nsrexecd** service is not currently running on the client.

To resolve these problems, ensure that the **nsrexecd** service is running on the client and that the server's hostname is listed in the boot-time file. The boot-time file lists all the servers, in order of precedence, that can contact a client for backups.

[Table 128 on page 833](#) lists the location for the boot-time file. The **nsrexecd** man page provides information about **nsrexecd**.

Table 128 Boot-time file locations

Operating system	Boot-time file
AIX	/etc/inittab /etc/rpc /etc/syslog.conf
HP-UX	/sbin/init.d/networker
Linux	/etc/init.d/networker /etc/rc3.d/S95networker /etc/rc5.d/S95networker /etc/rc0.d/K05networker
Solaris	/etc/init.d/networker

Microsoft Windows issues

This section covers issues on Windows.

Certified protocols

This release of NetWorker software has been tested and is certified to work on Microsoft TCP/IP. Other protocols, such as Novell IPX/SPX, Microsoft IPX/SPX, and Microsoft Proxy Client, are not certified to work with NetWorker software at this time.

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 NetWorker software cannot work around. New.Net software is included in products such as Go!Zilla, BearShare, Mp3.com, iMesh, Babylon, Cydoor, Webshots, and gDivx.

If you suspect that the New.Net DLL is the cause of problems, search for the newdotnet.dll file on the system drive. If you find this file, uninstall the New.Net software.

NOTICE

Do not manually delete the `newdotnet.dll` file. Doing so renders the system unusable.

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 cannot be performed from the NetWorker server, the archive client's username (for example, `root`) might not be listed in that client's Archive Users attribute in the Client resource.

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`, they end up in a single archive save set. To retrieve archives separately, archive them separately.

Wrong archive pool is selected

If multiple archive pools exist, the last one created is selected for archive.

Second archive request does not execute

If you create two archive requests with the same name, only the first request is executed. Do not create two archive requests with the same name.

The nsrarchive program does not start immediately

If `nsrarchive` is run from the command-prompt, the archive will not start immediately. Wait a short time until the archive starts. Do not press `[Ctrl]+[D]` multiple times.

Archive request succeeds but generates error when nsrexecd is not running

During an archive request operation, an error is generated when `nsrexecd` is not running on a remote client. The archive operations succeeds, but the following error message is logged to the daemon log file:

```
Failed to get port range from local nsrexecd: Service not available.
```

[“Viewing log files” on page 803](#) provides information about viewing log files.

Empty annotations in retrieve list

Older releases of the NetWorker Archive application software installed on DOS, Windows, and NetWare lack an annotation feature. As a consequence, the annotations for save sets archived with the older software are empty strings in the retrieve list.

Storage nodes

This section provides troubleshooting information about storage nodes.

Storage node affinity errors

A storage node affinity problem may exist if a backup fails with this error message:

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

Possible reasons include:

- ◆ No devices are enabled for 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.

Fix the problem and restart the backup. Do one of the following:

- ◆ Enable devices on one of the storage nodes.
- ◆ Correct the pool restrictions for the devices listed in the Storage Nodes attribute.
- ◆ Add another storage node to the Storage Nodes attribute that has devices that are enabled and that meet the pool restrictions.
- ◆ Set one of the devices to read/write.
- ◆ Adjust the Save Mount Timeout and Save Lockout attributes for the storage node's Device resource.

Storage node timeout errors

If **nsrd** initializes on the server and detects that a setting for **NSR_MMDCONTROL** exists, this message appears:

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

If you receive this message:

1. Shut down the NetWorker services.
2. Remove the environment setting for **NSR_MMDCONTROL**.
3. Restart the NetWorker services.
4. Start the Console server.
5. Adjust the value of the **nsrmmd** Control Timeout attribute to the value previously assigned to the **NSR_MMDCONTROL** variable, or one that best meets the current requirements. [“Modifying the timeout attribute for storage node operations” on page 135](#) provides more information.

Console error messages and corrective actions

Table 129 on page 836 provides a list of Console error messages or symptoms and corrective actions to take.

Table 129 Error messages or symptoms (1 of 3)

Error message or symptom	Possible cause	Corrective action
If the Console server fails to load and instead displays a Save As... dialog box.	In Internet Explorer: Either the web browser's security level is set to High (disabling JavaScript, which is needed to launch the product), or JavaScript has been disabled by some other means.	In Internet Explorer: Lower the web browser's security setting or enable Active Scripting.
Authorization code not accepted.	NetWorker software temporary enabler code has already expired.	Log out, then stop and restart the Console server.
Application window is unresponsive.	Insufficient disk space on the file system where the Console database is installed.	<ul style="list-style-type: none"> Ensure that the Console server is running. “Console troubleshooting notes and tips” on page 839 provides details. If it is not, close all application windows and check the gstd log file for errors. “Viewing log files” on page 803 provides information about viewing log file. Back up and move the Console database, if necessary. 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 Console window or Administration window.	Close any open dialog boxes or error messages.
Connection refused: no further information. or Problem contacting server <i>server_name</i> :	Console server is in the process of crashing or has already crashed.	Check to see if the Console server is running. <ul style="list-style-type: none"> If it is running, stop and restart the Console server. If it is not, close all application windows and check the gstd log file for errors. “Viewing log files” on page 803 provides information about viewing log files.
	Console server has been started within the previous few minutes.	Wait a couple of minutes and retry.
Failed to bind to port XXXX message in the gstd.raw log file.	The gstd service port (default 9001) is being used by some other process or is in a timeout (TIME_WAIT/FIN_WAIT) state.	Close any running NMC GUIs or any processes that may be using the gstd 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.	Console database is corrupt.	Recover the database. “Recovering the NMC server database” on page 417 provides details.

Table 129 Error messages or symptoms (2 of 3)

Error message or symptom	Possible cause	Corrective action
Display problem:	Console server is not running.	Restart the Console server.
In Internet Explorer: The page cannot be displayed.	Browser is not pointing to the correct URL.	Check the install log file to determine the HTTP port used by the Console server. “The install log” on page 839 provides details.
	Network connection is down.	Ping the Console server to confirm the network connection. If it is available, contact the system administrator.
Enabler code not accepted.	Temporary enabler code has expired.	<ol style="list-style-type: none"> 1. Close the Console server and log in again. 2. Repeat the procedure of entering the enabler code. If the enabler code is still not accepted, log out, then stop and restart the Console 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.	<ul style="list-style-type: none"> • Another user is trying to add a folder to the same location in the Enterprise at the same time. • An object was added with the same name as an existing object. 	<ul style="list-style-type: none"> • 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.	<p>If you are using the License Manager and a hostname has not been assigned:</p> <ol style="list-style-type: none"> 1. Select the Software Administration task. 2. Click Licensing. 3. Click Software Administration on the menu bar. 4. Click Change LLM Server. 5. Enter the new License Manager hostname. 6. Click OK. 7. If License Manager is installed, but not running, start it. <p>The <i>NetWorker License Manager Installation and Administrator’s Guide</i> provides details.</p>
	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 <i>before</i> License Manager is started. If the services are not started in the correct order, an error condition occurs.	<ol style="list-style-type: none"> 1. Stop the NetWorker software. 2. Stop License Manager if it is running. 3. Restart License Manager. 4. Restart the NetWorker software.
License allocation failed.	Temporary license for NetWorker software is expired.	Enter enabler codes and register the product.

Table 129 Error messages or symptoms (3 of 3)

Error message or symptom	Possible cause	Corrective action
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 will be deleted within 24 hours.
Logging of debug messages has stopped. alloc /opt: file system full.	Disk space on the /opt file system is nearly depleted.	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".	<p>Java Web Start preferences are set to something that is incompatible with the rest of the environment.</p> <p>(For example, a proxy server has been set up that stops Java Web Start from downloading the Console client software from the Console web server.)</p> <p>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.</p>	<p>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.</p> <p>(In the proxy server example, go to the General tab of the Preferences dialog box and select None, for Proxies.)</p> <p>If the Java Web Start cache path contains non-English characters, change the path to contain no non-English characters.</p>
<i>gstd.log</i> file error: internal error: could not end transaction	When the system time is moved ahead, a time out event is initiated and the database client connection for the gstd process is closed.	None

Console log files

The Console server produces these log files:

- ◆ install.log
- ◆ gstd.raw
- ◆ db_output.log
- ◆ dbstop_output.log
- ◆ dbstop_output.log
- ◆ web_output

The install log

Refer to install log files when doing one of the following:

- ◆ Troubleshooting a problem with the Console server.
- ◆ Tracking decisions made during installation, such as the HTTP service port chosen for the web interface.

By default, the install log files are located in /opt/lgtonmc/logs (UNIX) or C:\Program Files\EMC NetWorker\Management\GST\logs (Microsoft Windows).

The gstd log

The gstd log file contains messages from the Console server. Whenever the Console server is restarted, the size of the gstd log file is checked. If the gstd log file has reached its maximum size, the Console server starts a new gstd log file.

The gstd log file is placed in these default locations:

- ◆ On UNIX/Linux: /opt/lgtonmc/logs
- ◆ On Windows: %SystemDrive%\Program Files\EMC NetWorker\Management\GST\logs

“[Viewing log files](#)” on page 803 provides information about viewing log files.

[Table 130 on page 839](#) lists the variables that control the gstd.log file.

Table 130 Environment variables for the GSTD log

Variable name	Description
GST_MAXLOGSIZE	Sets the maximum size of the gstd log file before it is renamed on GST restart.
GST_MAXLOGVERS	Sets the maximum value of nnn in gstd.nnn.
GST_DEBUG	Sets the level of verbosity of the gstd log file. Can also be set from the System Options dialog box.

“[Setting system options](#)” on page 531 provides more information.

Console troubleshooting notes and tips

This section provides general troubleshooting tips for the Console server.

Making sure the Console server is running

If the Console 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 started 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 instance of the Console server to run each operation.

- ◆ Are the following processes running?

- GST server (**gstd**)
- Database server (**dbsrv12**)
- Web server (**httpd**)

These processes must be running to support the Console server.

- ◆ Is the **ntpdate** command synchronizing at midnight?

In some cases, having a cron job that has **ntpdate** synchronize at exactly midnight can cause the Console server to lose connection to the database. If such a situation occurs, modify the cron job to have **ntpdate** synchronize at some time other than midnight (12:00 A.M.) or have **ntp** run as a service and synchronize continuously.

How to determine 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.

How to determine if the Console server is running on a Solaris system

- ◆ To check whether the **gst** server process is running, enter this:

```
/usr/bin/ps -ef | grep gstd
```

If the **gst** process is running, a result similar to this appears:

```
root 6140 1 0 12:54:10 ?0:03 /opt/lgtonmc/bin/gstd
```

- ◆ To check whether the database process is running, enter this:

```
/usr/bin/ps -ef | grep dbsrv
```

If the database server is running, a result similar to this appears:

```
LGTOnmc root6140 1 0 12:54:10 ?0:03  
/opt/lgtonmc/sybase/bin/dbsrv12
```

- ◆ To check whether the web server process is running, enter this:

```
/usr/bin/ps -ef | grep httpd
```

If the web process is running, a result similar to this appears:

```
LGTOnmc root6140 1 0 12:54:10 ?0:03 /opt/lgtonmc/bin/httpd
```

Enabling Java script

If JavaScript becomes disabled, the Console server will not launch. Check the web browser's settings and reenable JavaScript if necessary.

Note: The procedure for enabling a given browser's version of JavaScript might differ from the instructions shown here. If it does, consult the browser's Help application for information about enabling JavaScript on the browser.

Java Web Start jnlp file caching issue after upgrading the NetWorker Console

After the NetWorker Console is upgraded or a client locale is changed, the **gconsole.jnlp** file will be different than the original **gconsole.jnlp** file in the Java Web Start cache. The NetWorker console may fail to launch.

Workaround

Remove the NetWorker Management Console Application and Language Pack (if applicable) from the Java Cache Viewer:

1. Run the Java Cache Viewer. From the command line, use the **javaws -viewer** command to launch the application.

Two different windows are displayed on the screen.

2. In Java Cache Viewer window, select **Applications** in the **Show** drop-down list. Remove all instances of **NetWorker Management Console** from the table below.
3. In the **Show** drop-down list, select **Resources**. Remove all URL entries in the table that start with the text, *Error! Hyperlink reference not valid*.
4. Close the Java Cache Viewer window.
5. In the Java Control Panel window click **Settings**.
6. Click **Delete Files** and click **OK**.

Querying large numbers of save sets in the NetWorker user interface may cause a Java heap space error

Querying large numbers of save sets in the NetWorker user interface may fail with a Java heap space error.

Workaround

Increase the Java heap size used by the NMC application:

1. On the Console server host, open the gconsole.jnlp file in a text editor. The gconsole.jnlp file is located in:

```
Console_install_dir\web
```

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.

“Unable to connect to host” error in the Client Backup Configuration wizard

The following message may appear when attempting to complete tasks that use the remote agent:

```
Unable to connect to host: Please check Security setting and daemon
logs on the Networker client and Console server for more details
```

This message may appear when performing one of the following:

- ◆ Client Configuration wizard tasks
- ◆ Device Configuration wizard tasks
- ◆ Save set browsing when adding or modifying a client resource

Check for one of the following when you receive this error:

1. 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
```

Note: After the deletion is complete, it is not mandatory to restart the NetWorker or Console services.

2. The Client cannot resolve hostname of NMC Server or NW Server. Sometimes, NMC can resolve the client hostname, but, client cannot resolve 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.

3. Ensure NetWorker users have at least the "Operate NetWorker" privilege to launch the Client Wizard. To resolve this issue, add the user to the appropriate user_group in the NetWorker Server.

4. 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.
5. The NMC Server, NetWorker Server and NetWorker client hosts must only use nsrauth authentication.

Username/password validation fails when using NMC New Device wizard to configure an AFTD if storage node is UNIX

When using the NMC New Device Wizard to configure an AFTD, username/password validation for browsing the file system may fail if the storage node is a UNIX host. This failure occurs if the system is missing the Pluggable Authentication Modules (PAM) library, or when the rule in the pam.conf file (/etc/pam.conf) for **OTHER service** is set to **deny**.

Perform the following if validation fails when using the New Device Wizard on a UNIX storage node:

1. Install the PAM package appropriate to your environment if it is not already installed.
2. Modify the pam.conf file so that the rule for **OTHER service** is not set to **deny**.

The operating system's documentation provides more information.

NMC user interface exits unexpectedly

If the NMC GUI loses its connection to the **gstd** service because the **gstd** service was shutdown or failed, the GUI will give a warning and exit after 10 seconds. This is normal behavior. [“Console error messages and corrective actions” on page 836](#) provides more troubleshooting information.

APPENDIX A

SYSTEM, VSS SYSTEM, and WINDOWS ROLES AND FEATURES Save Sets

This appendix covers these topics:

- ◆ [SYSTEM save sets](#) 846
- ◆ [VSS SYSTEM save sets](#) 849
- ◆ [WINDOWS ROLES AND FEATURES save sets](#) 851

SYSTEM save sets

The SYSTEM save sets discussed in this section work with these operating systems:

- ◆ Windows XP Professional
- ◆ Windows Server 2003 32 and 64 bit (with no VSS client license, or with VSS disabled)

The SYSTEM save sets includes the system state, system files, and system DB.

Components of the SYSTEM STATE save set

As part of the SYSTEM STATE save set, NetWorker software backs up all Windows system state components except the SFP component. [“Components of the SYSTEM FILES save set” on page 847](#) provides information about how NetWorker software treats the SFP component.

Certain basic system state components are present with every Windows Server 2003, or Windows XP Professional installation. These components are present on every system, and are always part of the SYSTEM STATE save set. [Table 131 on page 846](#) lists these basic components and references to their backup and recovery procedures.

Table 131 SYSTEM STATE save set basic components

SYSTEM STATE basic component	Special backup and restore considerations
COM+ Database	To complete the backup and recovery operation of the COM+ database, ensure that a valid temporary directory is set with the TEMP environment variable.
Internet Information Server (IIS)	“Internet Information Server” on page 888 provides details.
Registry	“Windows registry” on page 888 provides details.
Performance Counters	None

Optional components of the SYSTEM STATE save set

In Windows XP Professional and Windows Server 2003, the SYSTEM STATE save set can include optional components, under these conditions:

- ◆ The optional components have been installed.
- ◆ The components’ corresponding services have been started.

These SYSTEM STATE optional components and their procedures for backup and recovery are listed in [Table 132 on page 846](#).

Table 132 SYSTEM STATE save set optional components

SYSTEM STATE optional component	Special backup and restore considerations
Active Directory (AD)	“Active Directory” on page 886 provides details.
Certificate Server	None
Cluster Server	The NetWorker <i>Cluster Integration Guide</i> provides details.
File Replication Service (FRS, also called SYSVOL)	None

Components of the SYSTEM FILES save set

In Windows XP Professional and Windows Server 2003, the SFP feature prevents overwriting of certain essential system files (most commonly, dynamic link libraries and executables) by application installations. These files are called *system-protected files*. By preventing the replacement of these critical system files, file version mismatches are avoided that might otherwise cause application errors or system crashes. The SFP component includes SFP catalog files, system protected files, and system boot files (ntldr, ntddetect.com, and boot.ini). The SYSTEM FILES save set also includes the IA-64 EFI FAT partition component.

The system state must always be backed up at level full (partial backups are not allowed). However, the SFP component typically consists of more than 200 MB of data in over 1,500 files. If the SFP component were part of the NetWorker SYSTEM STATE save set, backing up or recovering the system state would be extremely resource-intensive. Therefore, NetWorker software backs up the SFP component in its own save set called SYSTEM FILES.

System File Protection (SFP)

For a full backup operation, specifying the SYSTEM FILES save set or save set *All* results in a full backup of the system-protected files. However, on an incremental or level 1-9 backup of the SYSTEM FILES save set or save set *All*, if any system-protected files have changed since the specified time, all system-protected files are backed up. If no system-protected files have changed, none will be backed up and no corresponding save set entry is made in the server's media index.

To ensure a proper recovery of the Windows Server 2003 (with no VSS license, or VSS disabled) or Windows XP Professional state of the computer, it is safest to restore all three SYSTEM save sets in the same operation. In the NetWorker User program, if only the SYSTEM FILES or SYSTEM STATE save set is marked for recovery, a dialog box displays a warning that both of these save sets should be recovered together. No such warning is provided for recoveries performed at the command-prompt.

Not all SFP files necessarily exist on a computer at any given time. Additional SFP files are sometimes installed automatically when a new Windows system component is installed. When this occurs, the new files have a creation date that corresponds to the system component installation date. But the files have a modification date that corresponds to the creation date of the Windows distribution. (The modification date of a new SFP file is usually the same modification date of the already existing SFP files.)

When a backup of save set *All* or the SYSTEM FILES save set detects new SFP files, it checks for the more recent of the file creation and file modification dates. If the more recent date is after the *as of time*, NetWorker software backs up *all* of the system-protected files.

Components of the SYSTEM DB save set

The NetWorker SYSTEM DB save set is used to back up the Windows Server 2003 (32 and 64 bit) or Windows XP Professional system databases that are installed and started.

[Table 133 on page 848](#) lists the system databases that are present by default and references to their backup and recovery procedures.

Table 133 SYSTEM DB save set basic components

SYSTEM DB basic component	Special backup and restore considerations
Content Index Server (CIS)	“Backing up the Windows Content Index Server” on page 114 and “Restoring Windows Content Index Server on Windows” on page 398 provide details.
Disk Quota Database	“Granting full permissions for backup of Disk Quota database” on page 117 provides details.
Removable Storage Database	Removable Storage database backup and recovery is not supported.
Windows Management Instrumentation	None

Note: The CIS appears in the SYSTEM DB save set only if Indexing Service is started.

Windows databases not included in the SYSTEM DB save set

[Table 134 on page 848](#) lists the additional Windows Server 2003 and Windows XP Professional databases that NetWorker software supports. Each of these databases may be installed optionally. These databases are backed up as part of the file system, and *not* as part of any SYSTEM save set.

Table 134 Windows databases not in SYSTEM DB save set

Database	Backup and restore procedures
Distributed File System (DFS)	Appendix C, “Backing Up and Restoring a Microsoft DFS” provides details
Encrypting File System (EFS)	“Encrypting file system” on page 886 provides details.
Event logs	“Event logs” on page 887 provides details.
Sparse files	“Sparse files” on page 888 provides details.

Components of the SHAREPOINT save set

During a Microsoft SharePoint Portal Server (SPS) backup, a single save set called SHAREPOINT is saved. The SHAREPOINT save set can be restored only in its entirety.

Note: Microsoft Windows Server 2003 does not support SPS 2001.

The SHAREPOINT save set contains:

- ◆ Web Storage System files, including database, log, and backup patch files.
- ◆ Microsoft Search resources, including the property and subscription stores, full-text index files, and propagated indexes.
- ◆ Server configuration information for the Web Storage System, content sources, server properties, and access accounts.
- ◆ The Applications folder, which contains a subfolder for each workspace on the server. Each subfolder can include searchable applications designed for the Web Storage System.
 - All application-specific data stored in the Web Storage System is included.
 - All application-specific data stored outside the Web Storage System (such as registry settings) is not included.
- ◆ Any shortcuts or content sources that reference the local file system. Note that these do not work if the referenced content does not exist on the computer where the SHAREPOINT save set is restored. Also, you must restore any shortcuts to workspaces in My Network Places.

The SHAREPOINT save set does not contain:

- ◆ Scheduled tasks for processing subscriptions. SPS processes subscriptions based on default schedules at the time of the recovery operation.
- ◆ The gather log that SPS creates each time it updates an index. This file contains data about the URLs that SPS accesses while generating an index.

VSS SYSTEM save sets

The VSS SYSTEM save sets discussed in this section work with computers installed with Windows Server 2003, when VSS is enabled.

The VSS SYSTEM SAVE SET includes the VSS System Boot, VSS System Fileset, and VSS System Services.

Components of the VSS SYSTEM BOOT save set

The VSS SYSTEM BOOT save set includes all elements of the Windows system state. All other components are dynamically generated, and therefore may differ each time NetWorker software runs.

Components of the VSS SYSTEM FILESET save set

The VSS SYSTEM FILESET save set includes the VSS System Writer writer.

Most of the elements of the Windows system state are part of the VSS SYTEM BOOT save set. However, Windows System Writer can consist of thousand of files and gigabytes of data. To relieve the strain on resources, NetWorker software backs up the System Writer System Files component as part of the VSS SYSTEM FILESET save set.

Components of the VSS SYSTEM SERVICES save set

The VSS SYSTEM SERVICES save set includes all elements of the Windows System Services. It also includes the Disk Quota database legacy component, File Server Resource Manager (FSRM) Disk Quota, and the DFS Replication writer. In addition, for Windows Server clusters, it includes the Cluster writer.

The Disk Quota database legacy component of the save set is recovered by using the legacy method of recovery. All other components of this save set are dynamically generated and may differ each time NetWorker software runs.

Windows Server Cluster writers

Backup and recovery of Windows Server cluster is supported for Windows Server 2003. In Windows Server 2003, the NetWorker software uses the following specifications:

- ◆ Cluster writer name: Cluster Service Writer
- ◆ Backup and recovery save set: Under the VSS SYSTEM SERVICES system save set
- ◆ Mode: Regular recovery mode only, using the NetWorker user interface or command line.

VSS SYSTEM Recovery Considerations

To properly recover the entire system, especially in a disaster recovery situation, back up and recover VSS SYSTEM BOOT, VSS SYSTEM FILESET, VSS SYSTEM SERVICES, and all boot/system volumes. During recovery, if you choose not to mark all three save sets, a warning message will appear with the option to mark the other save sets. The *NetWorker Procedure Generator* provides more information.

NOTICE

When performing a disaster recovery in multiple Windows platforms and copying the registry, a failure may be reported during the recovery of **VSS SYSTEM BOOT**: due to the size of the **PendingRenameFileOperations** registry value, which is populated during the disaster recovery. The error message indicates a lack of system resources. If this error appears, it is recommended to set the variable NSR_RECOV_TEMP_CLEANUP to an appropriate value (for example, 1) in the system space, and then restart the disaster recovery. Setting this variable ensures that the error does not appear, and that the recovery and subsequent cleanup of the temporary recover files after restart occur without this interruption.

WINDOWS ROLES AND FEATURES save sets

The WINDOWS ROLES AND FEATURES save set discussed in this section work with these operating systems:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Windows 7
- Windows 8
- Windows 8.1

The WINDOWS ROLES AND FEATURES save set is introduced in NetWorker 8.1. It replaces the VSS SYSTEM BOOT, VSS SYSTEM FILESET and VSS SYSTEM SERVICES save sets.

The WINDOWS ROLES AND FEATURES save set supports online restores for the following:

- Active Directory
- DFSR
- Cluster

Considerations for the WINDOWS ROLES AND FEATURES save set

The WINDOWS ROLES AND FEATURES save set is not supported with BBB.

If you cancel a deduplication recovery, the state of the recovered data is not reliable and may contain corrupted data. To ensure the recovery is correct, you should restart the deduplication recovery process.

Missing VSS writer files are no longer tracked for backup. Only the files that are available at the time of backup are restored.

The Missing System Writer is still checked and if detected, will fail a DISASTER RECOVERY:\ or ALL backup.

Required volume is not available for the WINDOWS ROLES AND FEATURES save set.

The WINDOWS ROLES AND FEATURES save set and its children cannot be selected along with other file system backups.

Writer and component level restore of the WINDOWS ROLES AND FEATURES save set are not available from the NMC Recovery UI.

When cloning a WINDOWS ROLES AND FEATURES save set, you must clone the entire savegrp that created the save set. The actual data is backed up with the volumes included in that savegrp.

The ALL-DFSR save set applies to all supported platforms. Unlike the ALL save set, where the DFSR namespace is skipped because it is a junction point, every namespace, along with the associated replication folders are backed up.

The WINDOWS ROLES AND FEATURES save set supports component level granular restore. The system state and replication folders can be restore separately.

Do not restore the Windows Roles and Features system state multiple times in succession without rebooting the computer as required. Not rebooting the computer can result in a system with an unreliable operational state.

In NetWorker 8.1 SP1, the simultaneous recovery of WINDOWS ROLES AND FEATURES save sets from different browse times is not supported. You can only mark one WINDOWS ROLES AND FEATURES save set for each recovery.

APPENDIX B

Firewall Support

This appendix covers these topics:

- ◆ Overview..... 854
- ◆ Special considerations for firewall environments..... 855
- ◆ Configuring the port ranges..... 861
- ◆ Examples..... 867
- ◆ Troubleshooting..... 872

Overview

NetWorker uses a direct socket connection to communicate and move data across the network, to the required service with minimal overhead. While NetWorker opens some ports for TCP and UDP, NetWorker only requires TCP ports. UDP ports are optional.

The NetWorker software uses two types of ports during communication:

- ◆ “Service ports” on page 854
- ◆ “Connection ports” on page 854

Service ports

NOTICE

Service ports are also known as listener ports or destination ports.

The TCP server processes that run on each NetWorker host uses service ports to listen for inbound connections. NetWorker uses two types of service ports:

- ◆ Fixed ports—NetWorker uses two fixed ports: TCP/7937 and TCP/7938. You must include these ports in the service port range of each NetWorker host. NetWorker uses these ports to initiate connections.
- ◆ Variable ports—NetWorker dynamically opens ports. A NetWorker host can allocate any port in the defined service port range and the NetWorker daemons select the dynamic ports within that range randomly. The default range is 7937-9936.

You can narrow or expand this range, as discussed in this chapter. To increase security in the environment, reduce this range to specify only the minimum number of service ports that the NetWorker software requires. The minimum value depends on the installation type and number of hosted NetWorker devices, if any. NetWorker stores the service port range in the NSR system port ranges resource in the NSR Local Agent (NSRLA) database of each NetWorker host.

Connection ports

NOTICE

Connection ports are also known as communication ports, source ports, or outbound ports.

NetWorker processes use connection ports to connect to a service. The NetWorker software requires one connection port for any type of communication between the client, storage node, and server.

NetWorker uses a default range, 0-0 to indicate that the NetWorker software allows the operating system to select the port for TCP clients. The operating system reserves connection ports for short-term use and reuses them, as needed. The operating system might allow you to configure the dynamic port range, for example by using **netsh** on Windows. NetWorker does not require modifications to this range and EMC recommends that you use the default dynamic port range.

There are no security concerns in using the default port range and EMC recommends that you do not change the range for any NetWorker hosts in the datazone. NetWorker performance problems or random malfunctions can occur when the range is too narrow.

Special considerations for firewall environments

You can configure some firewall products to close an open connection that is inactive for a defined period of time. NetWorker uses persistent connections between daemons to transfer information as efficiently as possible. Connections open at the start of communication, and close when the communication finishes. For example, a running backup may have connections open to with: `nsrmmmd` to send the backup data, `nsrindexd` to send the client file index information, and `nsrjobjd` to send control and status information. NetWorker connections between hosts can remain idle for periods of time that exceed the idle timeout value on the firewall and as a result, the firewall ends the connection. For example, the status connection to `nsrjobjd` is frequently idle during a backup. When there are no error messages to report, the connection will not have traffic until the backup completes and NetWorker generates the success message.

To prevent the firewall from closing a NetWorker connection prematurely, configure the firewall to not close idle connections. If you cannot eliminate the firewall timeout, then configure the datazone to send a keep alive signal between the hosts at an interval that is shorter than the timeout period defined on the firewall.

Configure the keep alive signal in one of two ways:

- ◆ [“Configuring TCP keep alives at the operating system level” on page 855](#)
- ◆ [“Configuring TCP keep alives within the NetWorker software” on page 856](#)

NOTICE

When you configure TCP keep alives within NetWorker, NetWorker does not send a keep alive signal across some connections, for example between the `save` and `nsrmmmd` processes. EMC recommends that you configure TCP keep alive signals at the operating system level to ensure all connections do not close prematurely. EMC does not recommend reducing the `TIME_WAIT` and `CLOSE_WAIT` intervals on a host in an attempt to reduce the demand for connection or service ports. When the intervals are too low, the port for a process might close while NetWorker is resending data packets to the process. In some situations, a new instance of a process connects to the port and incorrectly receives the data packet. This can corrupt the new process.

Configuring TCP keep alives at the operating system level

You can change the TCP KeepAlive parameters temporarily on UNIX or permanently on UNIX and Windows operating systems. Restart all NetWorker services after you change the TCP KeepAlive parameters.

Firewall configurations commonly define a 1 hour idle timeout. EMC recommends that you set the Wait Time Before Probing and Interval Between Retry Probes parameters to 57 minutes. The exact value you use to define these parameters depend what unit of measure the operating system uses.

For example:

57 min = 3420 seconds = 6840 half seconds = 3420000 milliseconds

Note: If the firewall time out is shorter than the common one hour value, further decrease these values. The network overhead as a result of enabling TCP KeepAlive is minimal.

Table 135 on page 856 summarizes the Wait Time Before Probing and Interval Between Retry Probes parameters for each operating system.

Table 135 Setting TCP parameters for each operating system

Operating system	Temporary setting	Permanent setting
AIX	# no -o tcp_keepidle = 6840 # no -o tcp_keepintvl = 6840 where the TCP parameter value is defined in half-seconds.	/etc/rc.net
HP-UX	# ndd -set /dev/tcp tcp_time_wait_interval 3420000 # ndd -set /dev/tcp tcp_keepalive_interval 3420000 where the TCP parameter value is defined in milliseconds.	/etc/rc.config.d/nddconf
Linux	# sysctl -w net.ipv4.tcp_keepalive_time = 3420 # sysctl -w net.ipv4.tcp_keepalive_intvl = 3420 where the TCP parameter value is defined in seconds.	Add the net.ipv4.tcp_parameter=<i>tcp_value</i> commands to the /etc/sysctl.conf file, then issue the following command: <ul style="list-style-type: none"> • RHEL: chkconfig sysctl on • SLES: chkconfig boot.sysctl on
Solaris	# ndd -set /dev/tcp tcp_time_wait_interval 3420000 # ndd -set /dev/tcp tcp_keepalive_interval 3420000 where the TCP parameter value is defined in milliseconds.	Add the ndd commands to the /etc/rc2.d/S69inet file.
Windows	n/a	Modify the following registry keys: HKLM\System\CurrentControlSet\Services\Tcpip\Parameters\KeepAliveTime DWORD=3420000 HKLM\System\CurrentControlSet\Services\Tcpip\Parameters\KeepAliveInterval DWORD=3420000

Configuring TCP keep alives within the NetWorker software

Use the NSR_KEEPALIVE_WAIT variable to define how frequently (in seconds) NetWorker sends a keep alive signal between the nsrexecd process on one NetWorker host and the nsrexec process on the other NetWorker host.

NOTICE

If you do not set the NSR_KEEPALIVE_WAIT variable or you set it to an invalid value (0, a negative number, or a nonnumeric string), then NetWorker does not send a keep alive signal.

Determining service port requirements

This chapter describes how to determine the minimum number of service ports that the NetWorker software requires and how to view or update the service port range value. When the datazone uses an external firewall, you must open the service port range in the firewall, for TCP connections.

Some operating systems enable personal firewall software on a host, by default. For example, Windows 7 enables Windows Firewall and RedHat Linux 6 enables iptables. The NetWorker installation process on Windows adds firewall rules to the Windows firewall for NetWorker. The NetWorker installation process on UNIX does not add firewall rules to a personal firewall. When you use personal firewall software on a UNIX host, you must manually create the firewall rules for the NetWorker software.

When the NetWorker software interacts with other applications in the environment for example, a Data Domain appliance, you must define additional service ports on a firewall as described in this chapter.

Before you modify the service port range on the NetWorker host or on a firewall, determine the minimum number of required service ports for the NetWorker host. The number of ports that the NetWorker software daemons and processes require for communication depends on the NetWorker installation type.

This section describes how to calculate the service ports required for each NetWorker installation type (Client, Storage Node, Server or NetWorker Management Console Server).

NetWorker client

This section describes the port requirements for standard, NDMP, and Snapshot clients.

Standard client

A standard NetWorker client requires a minimum of 4 TCP service ports to communicate with the NetWorker server. [Table 136 on page 857](#) summarizes the TCP service port requirements and the RPC program number for each program on a NetWorker client.

Table 136 NetWorker client port requirements for NetWorker server

RPC program number	Port number	Daemon/program
TCP/390113	TCP/7937	nsrexecd/nsrexec
TCP/390113	TCP/7938	nsrexecd/portmap
TCP/390435	Dynamic TCP port from the service port range	nsrexecd/res_mirror
TCP/390436	Dynamic TCP port from the service port range	nsrexecd/gss_auth

A standard NetWorker client requires two TCP service ports to communicate with the NMC server. [Table 137 on page 858](#) summarizes the TCP service port requirements and the RPC program number for each program on a NetWorker client.

Table 137 NetWorker client port requirements for NMC server

RPC program number	Port number	Daemon/program
TCP/390113	TCP/7937	nsrexecd/nsrexec
TCP/390113	TCP/7938	nsrexecd/portmap

NDMP client

An NDMP client that backs up to an NDMP device requires access to TCP ports through the firewall only. The service port range in the NSRLA database on the NetWorker host does not require modifications. [Table 142 on page 864](#) provides more information.

Snapshot client

When you configure snapshot backups, each Snapshot client requires 2 TCP ports for the PowerSnap service, in addition to the 4 standard client ports.

Table 138 Snapshot port requirements

RPC program number	Port number	Daemon/program
TCP/390408 (Snapshot services)	Dynamic TCP port from the service port range	nsrpsd
TCP/390409 (Snapshot services)	Dynamic TCP port from the service port range	nsrpsd/nsrsnapckd

NetWorker storage node

When calculating the service port requirements only consider devices that the storage node manages. To accommodate growth in the environment for example, the addition of new devices, allocate extra service ports for the NetWorker storage node.

The minimum number of service ports that a storage node requires is 5. This number includes the four TCP service ports required for a NetWorker client and one service port for the storage management process, **nsrsnmd**. NetWorker requires additional ports that differ depending on the device type used.

Use these formulas to calculate storage node port requirements:

- ◆ For NDMP-DSA or SnapImage devices:

`5 + One service port for each backup stream`

- ◆ For tape devices:

`5+ #devices + #tape_libraries`

- ◆ AFTD or Data Domain Boost devices:

`5+ #nsrmmnds`

where:

- ◆ `#devices` is the number of devices connected to the storage node.

- ◆ *#tape_libraries* is the number of jukeboxes that the storage node accesses. The storage node has one nsrlcpd process for each jukebox.
- ◆ *#nsrmmds* is the sum of the **Max nsrmmmd count** attribute value of each device that the NetWorker storage node manages.

[Table 139 on page 859](#) summarizes the port requirements specific to the storage node programs.

Table 139 NetWorker Storage node specific port requirements

RPC program number	Port number	Daemon/program
TCP/390111	Dynamic TCP port from the service port range.	nsrsnmd
TCP/390429	Dynamic TCP port from the service port range.	nsrlcpd
TCP/390104	Dynamic TCP port from the service port range. Total port number depends on device type.	nsrmmmd

Note: In enterprise environments where unattended firewall ports need to be restricted for security reasons, use the storage node attributes **mmds for disabled devices** and **Dynamic nsrmmmds unselected** (static mode) to prevent a listener from starting on an inactive nsrmmmd port. [“Configuring storage nodes” on page 133](#) provides details.

NetWorker server

The minimum number of service ports that a NetWorker server requires is 15. This number includes:

- ◆ The four TCP service ports required for a NetWorker client.
- ◆ Two TCP ports for nsrjobd.
- ◆ One TCP port for each of the following processes: nsrd, nsrmmdbd, nsrindexd, nsrmmgd, nsrsnmd, nsrlogd, and nsrccd.
- ◆ One optional UDP port for the nsrd/nsrstat process.

Additional ports are required when the NetWorker server manages devices. Additional port requirements differ depending on the device type used.

Use the following calculation to determine the service port range:

- ◆ For NDMP-DSA or SnapImage devices:

$$14 + \text{One service port for each backup stream}$$
- ◆ For tape devices:

$$14 + \#devices + \#tape_libraries$$
- ◆ For AFTD or Data Domain Boost devices:

$$14 + \#nsrmmds$$

where:

- ◆ *#devices* is the number of devices connected to the NetWorker server.
- ◆ *#tape_libraries* is the number of jukeboxes that the server accesses. Each nsrlcpd process requires one TCP port.
- ◆ *#nsrmmds* is the sum of the **Max nsrmmdd count** attribute value of each device that the NetWorker server manages.

To accommodate growth in the environment for example, the addition of new devices, allocate extra service ports for the NetWorker server.

NOTICE

The Software Configuration Wizard requires one service port. The port is dynamic and closes when the wizard closes. If you use the Software Configuration Wizard, add one additional port to the service port range.

[Table 140 on page 860](#) summarizes the port requirements for each NetWorker server program.

Table 140 NetWorker server specific port requirements

RPC program number	Port number	Daemon/program
TCP/390103	Dynamic TCP port from the service port range	nsrd
TCP/390109	User defined UDP	nsrd/nsrstat Note: Optional, used for internal communications. For example, automatic discovery and initial ping (is alive) checks of the NetWorker server. Backup and recovery operations do not use this port. NetWorker does not require this port through an external firewall.
TCP/390105	Dynamic TCP port from the service port range	nsrindexd
TCP/390107	Dynamic TCP port from the service port range	nsrmmdbd
TCP/390437	Dynamic TCP port from the service port range	nsrncpd
TCP/390433	Dynamic TCP port from the service port range	nsrjobd/jobs
TCP/390439	Dynamic TCP port from the service port range	nsrjobd/rap
TCP/390438	Dynamic TCP port from the service port range	nsrlogd
TCP/390430	Dynamic TCP port from the service port range	nsrmmgd

Note: If you restrict unattended firewall for security reasons, then use the storage node attributes **mmds for disabled devices** and **Dynamic nsrmmms *unselected*** (static mode) to prevent a listener from starting on an inactive nsrmmmd port. [“Configuring storage nodes” on page 133](#) provides details.

NetWorker Management Console server

The minimum service port range for a NetWorker Management Console server is the same as a NetWorker client. [“NetWorker client” on page 857](#) provides more information.

Configuring the port ranges

After determining the service port requirements for a NetWorker host, you must confirm which port numbers are available between each host, then configure the port range on each NetWorker host and on the firewall.

- ◆ [“Determining the available port numbers” on page 861](#)
- ◆ [“Configuring the port ranges in NetWorker” on page 861](#)
- ◆ [“Configuring the port ranges on the firewall” on page 864](#)

Determining the available port numbers

Before you define ports in the service ports attribute for a NetWorker host, determine which ports are available. Use the **netstat -a** command to determine the current service port allocations for a host.

After you determine which ports are available, you can decide which ports to allocate for NetWorker host communications. Review this information when choosing the ports:

- ◆ The service port range for each NetWorker host must contain port 7937 and 7938. The nsrexecd daemon reserves these ports and they cannot be changed.
- ◆ EMC recommends specifying ports within the default range 7937-9936.
- ◆ To avoid conflicts with other daemons or services on the host, do not assign ports under 1024.

Configuring the port ranges in NetWorker

The TCP server processes running on each NetWorker host listens and connects only on the ports specified in the Service ports attribute in the NSRLA database of each NetWorker host.

Define the service port range on each NetWorker host in the datazone in one of the following ways:

- ◆ [“By using NMC” on page 862](#)
- ◆ [“By using the nsrports command” on page 862](#)

NOTICE

Anytime you change the service port range, you must restart all NetWorker processes on the host for those changes to take effect.

By using NMC

You can use NMC to view and modify the current port ranges for each NetWorker host.

1. Connect to the NetWorker server.
2. In the NMC **Configuration** window, select **Local Hosts**.
3. Right-click the NetWorker host and select **Configure Port Ranges...**
4. On the **General** tab, review the value in the **administrators** attribute:
 - If you see the message: **No privilege to view administrator list**, then the account used to log in to the Console server does not have permission to modify the port ranges. [“Enabling updates of the NSR system port ranges resource” on page 863](#) describes how to provide user accounts with the ability to modify the **service port** attribute.
 - If you see accounts in the **administrator** attribute, then update the **Service ports** attribute with the calculated service port range. For multiple ranges, type one range per line.

NOTICE

EMC recommends that you do not change the **Connection ports** attribute from the default value 0-0.

5. Click **Ok**.
6. Stop and start the NetWorker services or daemons on the NetWorker host.

By using the nsrports command

Use the **nsrports** command to view and modify the current port ranges for each NetWorker host, from a command prompt.

```
# nsrports -s target_hostname [-S|-C] range
```

Table 141 nsrports options

Option	Description
-s target_hostname	Optional, use this option when updating the port range for a remote NetWorker host. “Enabling updates of the NSR system port ranges resource” on page 863 describes how to enable remote access of the NSR system port ranges resource.
-S range	Sets the service ports range to the value specified by range. The default range is 7937-7941. If the range is not consecutive set of ports, use a space to separate the port values.
-C range	Sets the connection ports range to the value specified by range. EMC recommends that you do not change the connection ports attribute from the default value 0-0.

In this example, the NetWorker host, myclient.emc.com:

- ◆ Uses the default service port range 7937-7940.
- ◆ Requires 4 service ports.

To modify the service port attribute in the NSR system port ranges resource on myclient.emc.com:

1. View the current port range.

For example:

```
#nsrports -s myclient.emc.com

Service ports: 7937-7940
Connection ports: 0-0
```

2. Update the service port range. Separate multiple port ranges with a space.

For example:

```
nsrports -s myclient.emc.com -S 7937-7938 7978-7979
```

NOTICE

If you do not have permission to update the **NSR system port ranges** attribute an error message similar to the following appears: **nsrexecd: User 'username' on machine 'hostname' is not on 'administrator' list.** [“Enabling updates of the NSR system port ranges resource” on page 863](#) describes how to enable user access to update the NSR system port ranges resource.

3. Confirm the service port attribute updated successfully.

For example:

```
#nsrports -s myclient.emc.com

Service ports: 7937-7938 7978-7979
Connection ports: 0-0
```

4. Stop and start the NetWorker services or daemons on myclient.emc.com.

Enabling updates of the NSR system port ranges resource

The NSRLA database on each NetWorker host has its own administrators list. By default, only users that login to the NetWorker host locally can update the NSR system port ranges resource.

To add users to the administrator list of the NSR system port ranges resource and enable remote updates of the attribute:

1. Connect to the target NetWorker host.
2. To access the NSRLA database, from a command prompt, type:

```
nsradmin -p nsrexec
```

3. To print the current administrator list, type:

```
p NSR system port ranges
```

In this example, only local users can update the attributes in the NSR system port ranges resource:

```
NetWorker administration program.
Use the "help" command for help, "visual" for full-screen mode.
nsradmin> p NSR system port ranges
type: NSR system port ranges;
```

```

service ports: 7937-9936;
connection ports: 0-0;
administrator: *@localhost;

```

- To update the **administrator** attribute with the remote account, type

```
update administrator: *@localhost, username@system
```

For example, if you connect to the Console server with the console user administrator from the Console client mnd.mydomain.com, type:

```
update administrator: *@localhost, administrator@mnd.mydomain.com
```

- When prompted, type **y**.
- To exit the **nsradmin** program, type **quit**.

Configuring the port ranges on the firewall

After changing the service port range you must make corresponding modifications to the firewall rules. The firewall must allow bi-directional TCP connections for the service port range defined on each NetWorker host. NetWorker opens the ports in this range for UDP services, but the ports are not required. The NetWorker software selects random port numbers in the service port range for the daemons.

The NetWorker software may communicate with other applications on ports outside of the service port range for example, to communicate with a Data Domain or Avamar Utility node. To enable communication between the NetWorker host and other applications, configure additional firewall rules. [Table 142 on page 864](#) summarizes the firewall requirements for each NetWorker installation type and third-party application.

Table 142 Firewall port requirements for NetWorker

Source host	Destination host	Protocol	Ports to open on the firewall to reach the destination host (1 of 4)
NetWorker client	NetWorker server	TCP	Port range determined in “NetWorker client” on page 857
NetWorker client	NetWorker storage node	TCP	Port range determined in “NetWorker client” on page 857
NetWorker client	NMC server	TCP	Port range determined in “NetWorker client” on page 857
NetWorker client	Data Domain	TCP TCP/UDP	2049,2052 111 (Portmapper)
NetWorker client	Avamar All Nodes	TCP TCP	27000 29000 (For SSL only)
NetWorker client	Avamar Utility Node	TCP	28001
NetWorker storage node	NetWorker client	TCP	Port range determined in “NetWorker client” on page 857
NetWorker storage node	NetWorker server	TCP	Port range determined in “NetWorker storage node” on page 858
NetWorker storage node	Data Domain	TCP TCP/UDP	2049,2052 111 (Portmapper)

Table 142 Firewall port requirements for NetWorker

Source host	Destination host	Protocol	Ports to open on the firewall to reach the destination host (2 of 4)
NetWorker storage node	ESX Cluster	TCP	902
NetWorker storage node	vCenter server	TCP	443
NetWorker storage node (NDMP-DSA or SnapImage)	NetWorker server	TCP	10000 Port range determined in “NetWorker storage node” on page 858
NetWorker server	ATMOS server	TCP	80, 443
NetWorker server	AlphaStor	TCP	44475
NetWorker server	NDMP filer	TCP TCP	10000 One user defined port in the range of 0-1024
NetWorker server	NetWorker storage node (NDMP-DSA or SnapImage)	TCP	10000 Note: When a Windows NetWorker server uses Windows Firewall, manually create an inbound rule in for the nsrdsa_save program to allow communications over TCP port 10000. Port range determined in “NetWorker storage node” on page 858
NetWorker server	NetWorker client	TCP	Port range determined in “Standard client” on page 857 .
NetWorker server	NetWorker storage node	TCP UDP	Port range determined in “NetWorker server” on page 859 Note: Open the 2 required UDP service ports on the firewall for TCP connections but there is no need to allow UDP connections through the firewall.
NetWorker server	Data Domain	TCP TCP/UDP	2049,2052 111 (portmapper) 161 (Port used by SNMPd to query the Data Domain system)
NetWorker server	Avamar Utility Node	TCP	7937,7938 2 ports in range 7939-9936
NetWorker server	DPA	TCP	3916,4001
NetWorker server	vCenter server	TCP TCP	443 Port range determined in “NetWorker client” on page 857
NetWorker server	VMware Backup Appliance (EBR/VBA)	TCP	8543 Port range determined in “Standard client” on page 857

Table 142 Firewall port requirements for NetWorker

Source host	Destination host	Protocol	Ports to open on the firewall to reach the destination host (3 of 4)
NetWorker server	Console server	TCP	Port range determined in “NetWorker Management Console server” on page 861
NetWorker server	NetWorker Module for Microsoft Applications	TCP	6278 (Control port) 6279 (Data port) Port requirements determined in “Snapshot client” on page 858 .
NetWorker server	AlphaStor server	TCP	44475
NetWorker Module for Microsoft Applications	NetWorker server	TCP	6278 (Control port) 6279 (Data port) Port requirements determined in “Snapshot client” on page 858 .
Avamar Utility Node	NetWorker client	TCP	28002
Console server	NetWorker server	TCP	Port range determined in “NetWorker Management Console server” on page 861 .
Console server	NetWorker client	TCP	Port range determined in “NetWorker Management Console server” on page 861 .
Console server	Data Domain	TCP TCP	161 (Port used by SNMPd to query the Data Domain system) 162 (Port used by SNMPtrapd to capture Data Domain SNMP traps)
Console client	Console server	TCP UDP TCP	9000 (Port used by HTTPd to download the Console user interface) 9001 (Port used to perform RPC for calls from the Console Java client to the Console server) 2638 (Port used by Tabular Data Stream (TDS) for database queries) You can modify default ports values. “How to confirm NetWorker Management Console server service ports” on page 867 provides more information.
DPA	NetWorker server	TCP	3741
DPA	Data Domain	TCP TCP/UDP	22 161 (Port used by SNMPd to query the Data Domain system)
DPA	Avamar Utility Node	TCP	55555

Table 142 Firewall port requirements for NetWorker

Source host	Destination host	Protocol	Ports to open on the firewall to reach the destination host (4 of 4)
Data Domain	Console server	TCP/UDP	162 (Port used by SNMPtrapd to capture Data Domain SNMP traps)
Data Domain	DPA	TCP/UDP	162 (Port used by SNMPtrapd to capture Data Domain SNMP traps)
VMware Backup Appliance (VBA/EBR)	NetWorker server		8080 Port range determined in “Standard client” on page 857

How to confirm NetWorker Management Console server service ports

You define the service ports that the Console server uses during the installation process. To confirm the defined port numbers, review the `gstd.conf` file and look for the following lines:

- ◆ `http_svc_port = http_service_port`
- ◆ `clnt_svc_port = client_service_port`
- ◆ `int db_svc_port = client_db_port`

where `http_service_port`, `client_service_port`, and `client_db_port` are port numbers. By default, the HTTP service port is 9000 and the client service port used to make RPC calls is 9001.

If you change the port values in the `gstd.conf` file, you must restart the `gstd` daemon.

Note: The `gstd.conf` file is located in the `Console_install_dir\GST\etc` on UNIX and `Console_install_dir\GST\etc` on Windows.

Examples

This section provides three examples to determine firewall port requirements. In each example, the NetWorker server resides in the secure network. Each example uses the following IP addresses and host names:

```
192.167.10.101 client_A
192.167.10.102 client_B
192.167.10.103 client_C
192.167.10.104 client_D
192.167.10.105 client_E
192.167.10.106 client_F
196.167.10.124 storage_node_X
192.167.10.125 storage_node_Y
192.167.10.127 storage_node_Z
192.167.10.126 NW_server
```

Example 68 Service port ranges on a bi-directional firewall

In this example:

- ◆ The Service port attribute on each client specifies a minimum of 4 service ports, for example: 7937–7940.

NOTICE

To simplify the configuration, configure each client to use the same 4 service port numbers.

- ◆ The firewall must allow outbound traffic, to the IP address of each NetWorker client, on each of the service ports defined in the **Service port** attribute on the NetWorker client. Because each client can specify the same port numbers, the firewall only needs to allow 4 ports for each client IP address. These port numbers can be a subset of the port numbers used by the NetWorker server, as in this example.
- ◆ In pseudo syntax, the firewall rule for the service ports would look like this:

```
TCP, Service, src 192.167.10.*, dest 192.167.10.101, ports
7937-7940, action accept
TCP, Service, src 192.167.10.*, dest 192.167.10.102, ports
7937-7940, action accept
TCP, Service, src 192.167.10.*, dest 192.167.10.103, ports
7937-7940, action accept
...
```

In the previous pseudo syntax, the firewall configuration allows:

- ◆ Incoming service connections to the IP address of the NetWorker server on ports 7937–7958, from the IP addresses of each storage node, client, and any other host on the subnet.
- ◆ Connections to the IP addresses for each storage node on ports 7937–7948, and to each client IP address on ports 7937–7940. Ensure that you configure each NetWorker host with the appropriate port range, then restart the NetWorker services each host.

This is the most stringent configuration possible, but difficult to maintain.

To simplify the configuration and administration of the datazone, assign a range of 22 ports, 7937–7958 to each host, and then configure the firewall to allow traffic to these ports on any host, from any host.

In pseudo syntax, the firewall rule for the service ports would look like this:

```
TCP, Service, src 192.167.10.*, dest 192.167.10.*, ports 7937-7958,
action accept
```

Example 69 Calculating service ports in a storage node environment

This example describes how to apply the basic rules of service port calculations to a sample network. In this example there is one NetWorker storage node on either side of the firewall. Clients D, E, and F in the secure network back up data to the storage node in the secure network. Clients A, B, and C in the insecure network back up data to the storage node in the insecure network. The firewall protects each host in the secure network. The firewall does not protect hosts in the insecure network. The firewall blocks network traffic from insecure to secure.




Figure 62 Service port ranges in a storage node environment

This example requires you to only open service ports for the NetWorker server on the firewall to allow inbound traffic. Calculate the service port requirements for the NetWorker server with this formula:

$$14 + (\text{num devices}) + (\text{num libraries}) + 1 (\text{client push}) = 14 + 6 + 1 + 1 = 22$$

In this example:

- ◆ The service ports attribute of the NetWorker server contains the range: 7937-7958.
- ◆ The firewall must allow inbound traffic, to the IP address of the NetWorker server, on each service port with the exception of the UDP port. In this example, 22 ports in the range of 7937 to 7958 must allow inbound traffic to the NetWorker server.
- ◆ In pseudo syntax, the firewall rule for the service ports would look like this:

```
TCP, Service, src 192.167.10.*, dest 192.167.10.126, ports
7937-7958, action accept
```

Example 70 Calculating service ports in a Data Domain environment

This example shows how to apply the basic rules to a sample network with clients A, B and C, one storage node X, and a Data Domain appliance in an insecure network. The NetWorker server and Console server are in a secure network. A single firewall separates

the secure network from the insecure network. The NetWorker server has a tape library and six drives. The client sends backup data to the Data Domain appliance and each client acts as a Console client.




Figure 63 Service port ranges in a Data Domain environment

NetWorker server system port requirements

Calculate the service port requirements for the NetWorker server with this formula:

$$14 + (\text{num devices}) + (\text{num libraries}) = 14 + 6 + 1 = 21 \text{ service ports.}$$

In this example:

- ◆ Configure the Service port attribute on the NetWorker server to use a minimum of 21 service ports, for example: 7937–7957.
- ◆ Configure the firewall to allow inbound traffic, to the IP address of the NetWorker server:
 - On the 21 service ports specified in Service port attribute of the NetWorker server. The UDP port is not required.
 - On TCP ports 2049 and 2052 for Data Domain connectivity.
 - On TCP ports 111 and 161 for Data Domain connectivity.

In pseudo syntax, the firewall rules for the service ports would look like this:

```
TCP, Service, src 192.167.10.*, dest 192.167.10.126, ports
7937-7957, action accept
TCP, Service, src 192.167.10.*, dest 192.167.10.126, ports 2049,
action accept
TCP, Service, src 192.167.10.*, dest 192.167.10.126, ports 2052,
action accept
TCP, Service, src 192.167.10.*, dest 192.167.10.126, ports 111,
action accept
UDP, Service, src 192.167.10.*, dest 192.167.10.126, ports 111,
action accept
TCP, Service, src 192.167.10.*, dest 192.167.10.126, ports 161,
action accept
UDP, Service, src 192.167.10.*, dest 192.167.10.126, ports 161,
action accept
```

NetWorker storage node system port requirements

The storage node is in the insecure network and uses a Data Domain appliance. There are two data domain devices and each device uses a Max nsrmmmd count value of 4. The Dynamic nsrmmmds attribute is enabled on the storage node.

Calculate the service port requirements for the NetWorker storage node with this formula:

5 + 8 = 13 service ports.

In this example:

- ◆ The system port attribute on the NetWorker storage node must specify a minimum of 13 service ports, for example: 7937–7949.
- ◆ The firewall must allow outbound traffic, from the NetWorker server to IP address of the NetWorker storage node:
 - On the 13 service ports specified in Service port attribute of the NetWorker storage node.
 - On TCP ports 2049 and 2052 for Data Domain connectivity.
 - On TCP/UDP port 111 for Data Domain connectivity.

In pseudo syntax, the firewall rules for the service ports would look like this:

```
TCP, Service, src 192.167.10.126, dest 192.167.12.125, ports
7937-7949, action accept
TCP, Service, src 192.167.126.*, dest 192.167.10.125, ports 2049,
action accept
TCP, Service, src 192.167.126.*, dest 192.167.10.125, ports 2052,
action accept
TCP, Service, src 192.167.126.*, dest 192.167.10.125, ports 111,
action accept
UDP, Service, src 192.167.126.*, dest 192.167.10.125, ports 111,
action accept
```

NetWorker client service port requirements

There are NetWorker clients in the insecure network. Each client requires four service ports. Two ports must be 7937 and 7938.

In this example:

- ◆ The Service port attribute on each client specifies a minimum of 4 service ports, for example: 7937–7940.

NOTICE

To simplify the configuration, configure each client to use the same 4 service port numbers.

- ◆ The firewall must allow outbound traffic, to the IP address of each NetWorker client, on the 4 service ports defined in the Service port attribute of the NetWorker client. Because each client can specify the same port numbers, the firewall only needs to allow 4 ports for each client IP address. These port numbers can be a subset of the port numbers that the NetWorker server uses, as in this example.
- ◆ In pseudo syntax, the firewall rules for the service ports would look like this:

```
TCP, Service, src 192.167.10.*, dest 192.167.10.101, ports
7937-7940, action accept
TCP, Service, src 192.167.10.*, dest 192.167.10.102, ports
7937-7940, action accept
TCP, Service, src 192.167.10.*, dest 192.167.10.103, ports
7937-7940, action accept
```

Troubleshooting

This section contains solutions to some common problems.

- ◆ [“Backups appear to stop responding or slow down dramatically” on page 872](#)
- ◆ [“Fallback to RPC portmapper service on port 111” on page 873](#)
- ◆ [“Cannot bind socket to connection port range on system hostname” on page 874](#)
- ◆ [“Failed to bind socket for service_name service: Can't assign requested address” on page 874](#)
- ◆ [“Service is using port port_number which is outside of configured ranges: range” on page 875](#)
- ◆ [“Connection refused” on page 875](#)
- ◆ [“Connection reset by peer” on page 875](#)
- ◆ [“Unable to obtain a client connection to nsrmmgd \(version #\) on host hostname” on page 875](#)
- ◆ [“nsrmdmp_save: data connect:failed to establish connection” on page 875](#)
- ◆ [“Unable to execute savefs job on host hostname: Remote system error - No route to host” on page 876](#)

Backups appear to stop responding or slow down dramatically

When you configure a firewall to drop packets outside an allowed range, and the firewall configuration does not allow for proper NetWorker connectivity:

- ◆ NetWorker will not get proper notification that a connection is not possible.
- ◆ The socket connections may not close correctly and remain in a TCP FIN_WAIT state. As a result, NetWorker will require more ports for client connectivity.

To avoid these issues, configure the firewall to reject packets outside the allowed range. When the firewall rejects packets, NetWorker receives an immediate notification of any connection failures and the remaining operations continue.

If you cannot configure the firewall to reject packets, reduce the TCP timeout values on the NetWorker server's operating system to reduce the impact of the problem. The *Performance Optimization and Planning Guide* describes how to change TCP timeout values.

Fallback to RPC portmapper service on port 111

NetWorker requires a fully functional RPC portmapper service to discover available program services and their current connection points. NetWorker can use either the default operating system SunRPC portmapper on port 111 (if present) or the internal NsrRPC portmapper available inside the nsrexecd process (by default on port 7938). The `/etc/services` file on a UNIX host or the `%SYSTEMROOT%\System32\Drivers\etc\services` file on a Windows host determines which portmapper service the operating system uses.

On a NetWorker host:

- ◆ The operating system uses the NsrRPC portmapper service for all RPC connections, by default.
- ◆ When the NetWorker software cannot reach the NsrRPC portmapper on the expected port, NetWorker attempts to use the SunRPC portmapper on port 111. In this default configuration, the services file contains entries for the sunrpc portmapper only:

For example:

```
sunrpc  111/tcp  rpcbind portmap #Sun RPC
sunrpc  111/udp  rpcbind portmap #Sun RPC
```

NOTICE

If a firewall rule blocks SunRPC, then delays in connectivity between a client and the server can occur. The NetWorker software will wait for the SunRPC connection attempts made by the operating system, to time out.

- ◆ If required, you can modify the operating systems 'services' file to change the default portmapper selection behavior.

For example, when you add entries for nsrrpc in addition to sunrpc, NetWorker will not try to use the SunRPC portmapper when NsrRPC is not available:

```
sunrpc  111/tcp  rpcbind portmap #Sun RPC
sunrpc  111/udp  rpcbind portmap #Sun RPC
nsrrpc  7938/tcp  lgtomapper      #EMC NetWorker RPC
nsrrpc  7938/udp  lgtomapper      #EMC NetWorker RPC
```

[Table 143 on page 874](#) summarizes the portmapper selection behavior for each service file configuration.

NOTICE

For a dedicated NetWorker server, EMC recommends that you disable the system SunRPC portmapper service.

Table 143 Services file entries

Services file entry	Action taken
sunrpc and nsrpc	NetWorker uses NsrRPC as the portmapper on port 7938. NetWorker never attempts to use the SunRPC portmapper.
nsrpc and no sunrpc	NetWorker uses NsrRPC as the portmapper on the defined port. NetWorker never attempts to use the SunRPC portmapper.
sunrpc and no nsrpc	NetWorker uses NsrRPC as the portmapper on port 7938. NetWorker uses SunRPC as a fallback when nsrpc is unavailable. This is the default Operating System configuration.
no sunrpc and no nsrpc	NetWorker portmapper used on port 7938. NetWorker never attempts to use the SunRPC portmapper.

When RPC portmapper program restarts, all applications registered to use RPC services must also restart because the portmapper database removes all previous registration information. For example, when the NetWorker software uses SunRPC service and the SunRPC service restarts, the NetWorker services must restart as well.

Cannot bind socket to connection port range on system *hostname*

This message appears in the savegroup messages or in stdout during manual operations when there are insufficient connection ports available and NetWorker cannot establish a connection.

To resolve this issue, ensure the Connection port attribute in the NSR System Port ranges resource is 0-0 on the host specified by *hostname*.

Failed to bind socket for *service_name* service: Can't assign requested address

This messages appears when a NetWorker daemon cannot register to a port within the service port range because all ports are in use by other daemons and process.

To resolve this issue, increase port range in the 'service ports' attribute in the NSR System port ranges resource on the NetWorker host and make a corresponding change in the firewall rules.

Service is using port *port_number* which is outside of configured ranges: *range*

This message appears in the Logs window when a NetWorker daemon attempts to register to a port that is not within the service port range. This can occur because the port requirements of the NetWorker host exceed the number of service ports defined in the range.

To resolve this issue, increase port range in the 'service ports' attribute in the NSR System port ranges resource on the NetWorker host and make a corresponding change in the firewall rules.

Note: Communications between NetWorker processes on the same host do not follow defined rules. For example, the NetWorker server daemons communicate internally outside of the defined port range. Do not configure a firewall to limit the range for TCP traffic inside a single system.

Connection refused

This message appears when the NetWorker host cannot establish a portmapper connection on port 7938 or 111.

To resolve this issue, ensure that the NetWorker software can register an RPC portmapper connection on port 7938 or 111. [“Fallback to RPC portmapper service on port 111” on page 873](#) provides more information.

Connection reset by peer

This message appears when the connection between two NetWorker hosts closes prematurely.

To resolve this issue, configure the datazone to send a keep alive signal between the hosts at an interval that is shorter than the time out period defined on the firewall. [“Special considerations for firewall environments” on page 855](#) describes how to configure the keep alive signal.

Unable to obtain a client connection to nsrmmgd (version #) on host *hostname*

This message appears on a Windows host when the Windows firewall **Allow** list on the NetWorker server does not contain the nsrmmgd process.

When this error message appears:

- ◆ A library configured on the NetWorker storage node will not enter “ready” state.
- ◆ Multiple **nsrlcpd** processes are started on the storage node.

To resolve this issue, ensure that the firewall is turned on, then add the **nsrmmgd** process to the **Allow** list of the Windows firewall on the NetWorker server host.

nsrndmp_save: data connect:failed to establish connection

This message appears during an NDMP-DSA backup when a Windows NetWorker server uses Windows firewall but an inbound rule for port 10000 does not exist.

To resolve this issue:

1. As Administrator, log in to the NetWorker server.
2. Go to the **Advanced** properties of the **Windows Firewall** application and select **Inbound Rules > New Rule...**
3. Select **Program** and click **Next**.
4. Select **This Program Path**.
5. Click **Browse...** and select the binary **nsrdsa_save.exe**, then click **Next**.
6. Select **Allow the connection**, click **Next**.
7. Leave the default **Profiles** selections enabled, click **Next**.
8. Provide a name for the rule and click **Finish**.
9. Edit the new rule.
10. On the **Protocols and Ports** tab:
 - a. In **Protocol type**, select **TCP**
 - b. In **Local Port**, select **Specific Ports** then specify port number **10000**.
 - c. Click **Ok**.

Unable to execute savefs job on host *hostname*: Remote system error - No route to host

This message appears during a scheduled backup when the NetWorker server can reach the client but cannot contact the nsrexecd process to start the savefs process.

To resolve this issue, ensure that you configure:

- ◆ Any external firewalls between the two hosts to allow communication on the required service ports.
- ◆ A personal firewall on the client, for example iptables on Linux, to allow communication between the two hosts on the required service ports.

APPENDIX C

Backing Up and Restoring a Microsoft DFS

This appendix covers these topics:

- ◆ Overview of a Microsoft DFS 878
- ◆ DFS topology information 879
- ◆ DFS topology information 879

Overview of a Microsoft DFS

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.

Benefits of DFS include:

- ◆ Easy browsing of servers
- ◆ Simplified searches for files and data
- ◆ Server load balancing

Domain-based DFS

Domain-based DFS has the DFS topology information stored in Active Directory (AD). Because this information is replicated on multiple domain controllers, domain-based DFS is fault tolerant. The DFS host server can be any Windows domain controller or member server.

Registry-based DFS

Registry-based DFS, also called stand-alone DFS, has the DFS topology information stored in the Windows registry on the DFS host server.

DFSR

DFS Replication (DFSR) provides basic file replication between servers. DFSR identifies modified or new files, and copies only those parts of files that have changed or been added.

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.

Save Set ALL-DFSR

NetWorker has a save set type called All-DFSR which includes all DFS related save sets for a specified backup.

Unlike other all-inclusive save set types, ALL-DFSR is not related to any particular file system. ALL-DFSR backs up all components defined by DFS\FRS writers.

The syntax for this save set is ALL-DFSR. It is not case sensitive.

ALL-DFSR 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.

If ALL-DFSR is specified for a system where DFS or FRS is not installed, backup will fail.

For Windows Server 2008 and later

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.

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. [Appendix A, "SYSTEM, VSS SYSTEM, and WINDOWS ROLES AND FEATURES Save Sets,"](#) provides more information about the SYSTEM and VSS SYSTEM save sets.

Configuring a scheduled DFS backup

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 on a computer running Windows 2003:

1. In the **Administration** window, 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. [Chapter 2, "Backing Up Data,"](#) provides more information about configuring a scheduled backup.

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.

- Using registry-based DFS only, include the following SYSTEM save sets:

SYSTEM STATE

SYSTEM FILES

If VSS is licensed and enabled, include the VSS SYSTEM save sets. (what was here in Steam Whistle?)

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:

4. For domain-based DFS only, include the following SYSTEM save sets in the domain controller's Save Set attribute:

SYSTEM STATE:

SYSTEM FILES:

If VSS is licensed and enabled, include the following (what was here in Steam whistle)

To configure a scheduled backup for a DFS on a computer running Windows 2008 or later:

1. In the **Administration** window, 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. [Chapter 2, "Backing Up Data,"](#) provides more information about configuring a scheduled backup.

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:

Restoring a DFS

For Windows 2003, the VSS SYSTEM save set and the AD can be restored as a single unit only.

For Windows 2008 and later, restore DFSR through Windows Roles & Features.

[Appendix A, “SYSTEM, VSS SYSTEM, and WINDOWS ROLES AND FEATURES Save Sets,”](#) provides more information.

To restore a DFS:

Windows 2003

The SYSTEM STATE save set, or the VSS SYSTEM BOOT save set and the AD can be restored as a single unit only. [Appendix A, “SYSTEM, VSS SYSTEM, and WINDOWS ROLES AND FEATURES Save Sets,”](#) provides more information.

To restore a DFS:

1. Restore the DFS topology information:
 - To restore a domain-based system, restore the SYSTEM STATE and SYSTEM FILES (or VSS SYSTEM BOOT and VSS SYSTEM FILESET) save sets on the domain controller.
 - To restore a registry-based system, restore the SYSTEM STATE and SYSTEM FILES (or VSS SYSTEM BOOT and VSS SYSTEM FILESET) save sets on the DFS host server.

[Chapter 14, “Recovering Filesystem Data,”](#) provides recovery procedures.

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 necessary, restore any local DFS destination directories.

3. If necessary, restore the remote DFS destination directories.

Window 2008 and later

1. Restore the DFS topology information:
 - To restore a domain-based system, restore the Windows Roles and Features save sets on the domain controller.

[Chapter 14, “Recovering Datafile system,”](#) provides recovery procedures.

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 necessary, restore any local DFS destination directories.

3. If necessary, restore the remote DFS destination directories.

Authoritative restores of DFS Replication writers

Windows 2003

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 on Windows 2003 systems, use the **-U** option with the **recover** command.

To restore the DFS Replication writer on Windows 2003 systems, type the following command:

```
recover -s server -U -N "VSS USER DATA:\DFS Replication service writer"
```

NOTICE

You cannot select individual components within the writer for recovery.

For Windows 2008 and later

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 on Windows 2008 systems and later, 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 restores of DFS Replication writers

Windows Distributed File System Replication (DFSR) granular recovery is supported on Windows Server 2008, Windows Server 2008 R2, Windows Server 2012 and Windows Server 2012 R2.

DFSR Shared Directories supports granular DFSR folder and file recoveries on computers that run Windows Server 2008 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.

DFS backups and restores for Windows 2003

A complete NetWorker backup of a DFS requires backups of the following:

- ◆ DFS topology information
- ◆ DFS junctions
- ◆ DFS destination directories (shared directories connected to DFS links)

The NetWorker software does not traverse DFS links and therefore does *not* back up DFS destination directories as a part of the DFS junctions backup. To properly protect data, back up the DFS destination directories.

APPENDIX D

Additional Features of the Microsoft Windows Server

This appendix covers these topics:

◆ NetWorker Module for Microsoft	886
◆ Active Directory	886
◆ Encrypting file system	886
◆ Event logs	887
◆ Internet Information Server	888
◆ Windows registry	888
◆ Sparse files	888
◆ Windows Change Journal.....	888
◆ Advanced Configuration and Power Interface.....	891
◆ Windows print queues	892
◆ Windows Optimized Deduplication	892

NetWorker Module for Microsoft

NetWorker Module for Microsoft (NMM) provides VSS-based backup and recovery of Windows, as well as Microsoft server applications such as Microsoft Exchange Server, Microsoft SQL Server, Microsoft Data Protection Manager (DPM), and Microsoft SharePoint Services.

If NMM is installed on the client computer, refer to the *NetWorker Module for Microsoft Administration Guide* for documentation about that product.

There is also a module available for Microsoft SQL Server called the **EMC NetWorker Module for Microsoft SQL Server** which supports back up and recovery of Microsoft SQL Servers.

Active Directory

Active Directory (AD) is the Windows directory service and the foundation for the Windows Distributed File System. AD is a component of the Windows system state on Windows Server 2008 and 2003 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.

Backing up Active Directory

The NetWorker software automatically backs up AD as a component of the SYSTEM STATE or VSS SYSTEM BOOT save set. An AD backup or restore includes the AD log files, database, patch files, and expiry token.

Recovering Active Directory

The NetWorker *Procedure Generator* contains information about the online recovery of Active Directory. [Chapter 25, “Windows Bare Metal Recovery,”](#) provides information about the bare metal recovery of an Active Directory host.

NOTICE

If you must recover an AD backup whose expiry token is older than its tombstone lifetime value, refer to Microsoft Knowledge base (KB) article 216993, which can be found at:

<http://support.microsoft.com/kb/216993>

Encrypting file system

Windows Encrypting 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.

Consider these when backing up or recovering files or folders that are encrypted with EFS:

- ◆ NetWorker software will not encrypt or compress a file already encrypted by Windows.
- ◆ Do not use AES encryption when backing up files that are encrypted using EFS.
- ◆ Files can become unusable if the encryption keys change on the domain controller. Reasons include:

- The domain controller functionality is moved from one computer to another
- The domain controller crashes.
- ◆ NetWorker software does not back up encryption keys. If the EFS is reinstalled after a disaster, the new security keys will not match the recovered keys and the recovery will fail. Keep a copy of the keys to ensure a successful recovery.
- ◆ You cannot perform a directed recovery of encrypted files.
- ◆ 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. Microsoft Windows documentation provides instructions on encrypting the re-created folder.
- ◆ Windows EFS encrypted data is backed up and recovered in its encrypted state.

Event logs

Event logs can be used for troubleshooting hardware problems as well as monitoring security conditions, and system and application software problems.

If VSS is disabled, the NetWorker software backs up event logs for services that are running at the time of the backup. You can restore event logs to the same location or to a new location on the computer and then view them with the Microsoft Event Viewer.

The size of a restored event log might be smaller than the size of the corresponding backed-up log. This is characteristic of Windows event logs and does not cause any loss or modification of data. The recovered, smaller log can still be viewed in the Microsoft Event Viewer.

If more than one active event log is marked for backup (for example, SecEvent.Evt and SysEvent.Evt), all event logs are backed up.

Event logs can be recovered to a location different than the location from which they were backed up. However, event logs cannot be recovered to a FAT16 or FAT32 partition if they were backed up from an NTFS partition.

If VSS is enabled, event logs are backed as a component of the VSS SYSTEM SERVICES save set.

NOTICE

Windows Server 2008 and Windows Vista do not have an event log writer. The event logs will not be backed up as part of the VSS system save sets. The event logs are backed up as part of the file system. To back up the event logs, perform a regular (non-VSS) backup of the `system32\winevt\logs` folder.

Internet Information Server

Internet Information Server (IIS) is a web server that enables the publication of information on the Internet or a corporate intranet by using HTTP.

The NetWorker software backs up IIS using its active metabase and restores backup versions to the metabase location, which can be at the default location (%SystemRoot%\system32\inetsrv\MetaBase.bin) or in a location specified in the registry. The Microsoft documentation provides information about creating a registry key to specify an alternate metabase location.

Relocation of the IIS metabase is not supported in IIS version 6.0.

NOTICE

If you perform a NetWorker recovery of the SYSTEM STATE or VSS SYSTEM BOOT save set (which includes the active metabase), and reboot, the Network News Transfer Protocol (NNTP) virtual server might not start. In that case, rebuild the NNTP index and hash table files. The Microsoft NNTP documentation provides more information.

Windows registry

For supported Windows Server versions, the registry is a component of the system state. The registry can be backed up and restored only as part of the SYSTEM or VSS SYSTEM save set to which it belongs.

In the NetWorker User program, the registry is a component of the SYSTEM STATE or VSS SYSTEM BOOT save set. The NetWorker software automatically backs up or recovers the registry as well as this save set.

The registry is always saved and restored at level full.

Sparse files

The NTFS sparse files feature enables a program to create huge files without actually committing disk space for every byte. The NetWorker software provides complete backup and recovery support for sparse files.

Windows Change Journal

Microsoft Windows Change Journal is a Windows Server 2003, and Windows XP Professional file system feature that logs a record of each change as it occurs to the files and directories on a local NTFS 5.0 volume.

The Change Journal enables NetWorker software to detect more types of changes and save more changed files than is possible when not using the Change Journal. Additionally, the Change Journal improves NetWorker performance.

The Change Journal can be enabled or disabled independently for each NTFS 5.0 volume. When enabled, the Change Journal stores records of the volume's file and directory changes in System Volume Information\tracking.log.

Note: When VSS is used, the Microsoft Change Journal is not used. Microsoft Windows Vista and Windows Server 2008 are VSS only, so Windows Vista and Windows Server 2008 do not use the Windows Change Journal.

NetWorker support for Change Journal

Windows does not have an administrative interface for enabling or disabling the Change Journal. That functionality is provided with the NetWorker Change Journal Manager.

The NetWorker Change Journal Manager is installed during the NetWorker software setup and can be run from **Start>Programs>NetWorker**.

The NetWorker Change Journal Manager allows you to:

- ◆ Enable or disable the Change Journal for each NTFS 5.0 volume.
- ◆ Enable or disable the NetWorker software's use of each volume's Change Journal.
- ◆ Set parameters that control the size of the Change Journal log file.

How NetWorker software uses the Change Journal

When configured to use the Change Journal, the NetWorker software bases its save decisions for level and incremental backups on the Change Journal log rather than on the traditional save criteria of modification time and Archive attribute.

[“Backup levels” on page 267](#) provides information about level and incremental backups.

The NetWorker software does not use the Change Journal for the following types of backups:

- ◆ Full backups
- ◆ Client-initiated backups
- ◆ Backups with an undefined level
- ◆ DFS backups
- ◆ VSS file system backups
- ◆ Backups of pseudo-volumes, such as the SYSTEM save sets

The Change Journal is only used when the NetWorker **save** command's path argument specifies an entire volume (for example, C:\ but not C:\MyDir). The time a change occurred continues to control the file save decision.

For a file to be selected for backup, it must have changed more recently than the *changed after time* (also called the *as of time*), as specified in the NetWorker **save** command's **-t** argument. The *EMC NetWorker Command Reference Guide* or the UNIX man page provides information about the NetWorker **save** command.

NetWorker save criteria when using the Change Journal

A NetWorker save that uses the Change Journal occurs when *all* of the following conditions are met:

- ◆ The target volume is NTFS 5.0.
- ◆ The Change Journal is enabled on the target volume.
- ◆ The NetWorker software is configured to use the Change Journal on the target volume.
- ◆ The save path is the entire volume (for example, C:\).
- ◆ The save set is a level or incremental backup.
- ◆ One or more of the save triggers occurs.

[Table 144 on page 890](#) shows the NetWorker save triggers when using the Change Journal.

Table 144 NetWorker save triggers

Change	Trigger
A file was deleted (or moved to another directory).	Save of the deleted file's directory.
A directory was deleted.	Save of the parent directory.
A directory was renamed.	Save of all files and subdirectories within the directory.
A file or directory was compressed or decompressed.	Save of the file or directory.
A file or directory was encrypted or decrypted.	Save of the file or directory.
The NTFS-extended attributes of a file or directory were changed.	Save of the file or directory.
The access rights of a file or directory were changed.	Save of the file or directory.
An NTFS hard link was added to or removed from a file or directory.	Save of the file or directory.

Configuring NetWorker software to use the Change Journal

The NetWorker Change Journal Manager enables one to view or edit the Change Journal configuration of each volume in the NetWorker server or client host computer.

How to configure NetWorker software to use the Change Journal

To view or edit the Change Journal configuration:

1. On the NetWorker server or client host computer, select **Start>Programs>NetWorker>NetWorker Change Journal Manager**.
2. In the **NetWorker Change Journal Manager** dialog box, select one of the following:
 - All NTFS Volumes — to view or edit the Change Journal configuration of all local NTFS 5.0 volumes at once.
 - A drive letter — to view or edit the Change Journal configuration of an individual volume.

Note: The status of the selected volume or volumes appears in the scrolling text box. If a selected volume does not support the Change Journal (FAT volumes, for example) the configuration options are dimmed.

3. To enable the Change Journal for use with NetWorker software, ensure that the **NetWorker Uses Change Journal** checkbox is selected. To disable the Change Journal, ensure that this checkbox is clear.

If you select **All NTFS Volumes**, and some (but not all) of the local NTFS 5.0 volumes already have this option enabled, the checkbox appears shaded. To change the setting of this option for all NTFS 5.0 volumes, click the checkbox until it is selected or clear, but not shaded.

4. To enable the Change Journal for use with Microsoft, select the **Enable Change Journal For Selected Volumes** checkbox.
 - To disable the Change Journal, ensure that this checkbox is clear.
 - If you select **All NTFS Volumes**, and some (but not all) of the local NTFS 5.0 volumes already have this option enabled, the checkbox appears shaded.
 - To change the setting of this option for all NTFS 5.0 volumes, click this box until it is selected or clear, but not dimmed.

Note: When you enable the Change Journal, there may be a delay of several minutes before logging begins.

5. Edit the following values to control the size of the Change Journal log file. If the Change Journal is already enabled for the selected volume or volumes, disable it before changing either of these values.
 - **% Of Volume For Log File** — The maximum amount of the volume's storage space that can be used for the Change Journal log file. The allowable range is 0.01% to 2.0% of the volume's capacity.
 - **% Of Log For Allocation Delta** — The amount by which the Change Journal log file can expand if additional space is needed. This is also the amount that will be purged from the beginning of the log when the file has reached its maximum size. The allowable range is 12% to 25% of the % Of Volume For Log File.
6. Click **Apply** to save any configuration changes.
7. Click **OK** to exit **NetWorker Change Journal Manager**.

Advanced Configuration and Power Interface

NetWorker software supports the following:

- ◆ Windows Server 2003
- ◆ Windows XP Professional
- ◆ Power Interface (ACPI), which is also called OnNow

NetWorker support for ACPI

Support for ACPI is provided by the *NetWorker Power Monitor* service. The executable for this service (<NetWorker_install_path>\bin\nsrpm.exe) is installed and configured for automatic startup during NetWorker setup.

For scheduled backups of Windows Server 2003, and Windows XP Professional clients, each client must be able to respond when a NetWorker server contacts it for backup. Therefore, by default, the NetWorker Power Monitor service does not allow a NetWorker server or client host operating system and its network interface to enter the ACPI standby mode.

However, the NetWorker Power Monitor service does not prevent a user from specifying that a computer can power down to standby state. Also, if line power is lost and the uninterrupted battery power reaches a critically low state, the NetWorker software does not prevent the host's power-management policies from forcing the system to power down. The NetWorker software shuts down any storage management operation that is in progress when standby is forced by a user action or a critical power event.

Considerations for ACPI usage

Before using ACPI, review the following conditions:

- ◆ Do not place a NetWorker server host or client host in standby mode during a time period when either computer is to participate in a scheduled backup.
- ◆ Do not put a NetWorker client or server host in standby or hibernation mode *while* a NetWorker backup or restore operation (involving that host) is in progress. Doing so will yield unpredictable results.
- ◆ If a NetWorker server is powered down while NetWorker operations are in progress, the server's peripheral devices might be powered down as well.

If this occurs, when the server's power is restored, the tape devices may rewind and NetWorker's tape processes will have incorrect positioning information.

Windows print queues

The following considerations apply to Windows print queues:

- ◆ Print queues are backed up and recovered as part of the file system and not as part of any VSS writer.
- ◆ During a recover operation, you may have to reboot depending on the status of the print queue.

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.

The data deduplication feature is supported on Windows Server 2012, Windows Server 2012 R2, Windows Storage Server 2012 and Windows Storage Server 2012 R2. It is not supported 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, the feature is supported only on volumes using the NTFS file system. These volumes can be part of a fail over cluster, including CSV volumes .

Backup of an optimized deduplication volume is an optimized deduplication backup by default, unless the backup path is a subdirectory of the volume or the non-optimized deduplication save option is not present. If the deduplication save option flag is present, the backup will not be deduplicated. In the case the path is a subdirectory of a volume, the backup created is not optimized.

A NetWorker 8.1 or later client is required to back up and restore Windows Server deduplication volumes or files. Additionally, deduplication backups can only be restored to computers that run on supported versions of Windows Server and 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*.

For more information on *mminfo*, refer to the *EMC NetWorker Command Reference Guide* or the *mminfo* man pages.

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:

1. From the NetWorker Administration console select **Properties** menu.
2. On the **Client Properties** text box, select the **Apps & Modules** tab.
3. 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.

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

This restore 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.

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

The following bullet points are recommended as best practices when backing up 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.

APPENDIX E

UNIX and Linux Platform-Specific Notes

This appendix covers these topics:

◆ Solaris	898
◆ Linux.....	899
◆ HP-UX	900
◆ AIX.....	904

Solaris

This section provides information specific to NetWorker software that runs on the Solaris platform.

Support for Solaris zones

The NetWorker software provides support for 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 machine running in a local zone. [“Dedicated storage nodes” on page 139](#) provides more information about storage node support in a local zone.

NOTICE

The NetWorker Console (NMC) and the NetWorker License Manager can only be installed in a global zone.

NetWorker executables not found for Solaris client

On Solaris, NetWorker executables are installed by default in `/usr/sbin`. If you start a group backup on a NetWorker server that does not have `/usr/sbin` in the search path for root, the backup fails on a client that has its NetWorker executables in `/usr/sbin`. This is because the `savefs` command is not in the search path.

To solve this issue, set the **Executable Path** attribute for the client.

How to set the Executable Path attribute

To set the Executable Path attribute:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Clients**.
3. In the right pane, select the client name.
4. From the **File** menu, select **Properties**.
5. For the **Executable Path** attribute on the **Globals 2 of 2** tab, enter the path of the executables, `/usr/sbin`.
6. Click **OK**.

Alternate solution: Modify the search path for root on the NetWorker server to include `/usr/sbin` even if it does not exist locally.

How to obtain support for devices not supported by Solaris

For devices that are not directly supported by Sun Microsystems for use with your operating system, obtain a `st.conf` file from the device manufacturer.

Extended file attribute data included in Save Set File Size attribute

The save set file size shown in NetWorker appears to be slightly larger than one might expect. This is because the extended file attribute data is included in the calculation of the save set file size.

The inquire command and Solaris 10

On Solaris 10, the **inquire** command does not show library information after the library has been configured for NetWorker.

Linux

This section provides information about NetWorker software that runs on the Linux platform.

Backup considerations for Linux raw disk partitions

The following considerations apply to backing up Linux raw disk partitions:

- ◆ The Linux raw device must be unbound before it can be saved.
- ◆ The save set must be `/dev/sd` or `/dev/hd`.
- ◆ The backup will fail if the `/dev/raw` device is used.

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 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 will only configure 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 be seen 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 will only display up to the number of `st` devices (`/dev/nst`) defined by the **ST_MAX_TAPES** value.

To resolve this issue, the `st` module of the Linux kernel must be modified and recompiled to increase the maximum number of allowable `st` devices created by the OS to exceed the default value. Refer to Linux documentation for details about 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, device files may need to be created so that the **inquire** command can detect all devices.

For example, on a NetWorker server that is running Red Hat Linux, if devices sg0 through sg15 already 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.

Linux Journaled file system support

Backup and recovery operations are supported on the following Linux journaled file systems:

- ◆ ext3
- ◆ reiserfs
- ◆ jfs
- ◆ xfs

NOTICE

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 this file system is excluded from a backup. [Chapter 9, “Directives”](#) provides information about directives.

HP-UX

This section provides information specific to NetWorker storage nodesoftware that runs on the HP-UX platform.

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.

To select unused SCSI addresses for an autochanger:

1. Log in as root on the NetWorker server or storage node and enter the **ioscan -f** command.
2. 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.

The following sections provide examples of the command and output to use with different combinations of hardware and operating systems.

Installing the SCSI pass-through driver

The following procedure describes how to install a GSC, HSC, or PCI pass-through driver.

How to install a GSC, HSC, or PCI pass-through driver

The following procedure assumes you are using the SAM terminal mode.

To install a GSC, HSC, or PCI pass-through driver:

1. Run **SAM**.
2. Select **Kernel Config** and press **Enter**.
3. Select **Drivers** and press **Enter**.
4. Select SCTL from the list. The SCSI_ctl driver is represented by the name SCTL.
 - If the current state is in, proceed to [“How to verify a device file” on page 901](#).
 - Select any unreserved name for the device. For example, do not select a name such as /dev/null.
5. From the **Actions** menu, select **Add Drivers to Kernel** and press **Enter**.
6. From the **Actions** menu, select **Create a New Kernel** and press **Enter**.
7. When prompted with “Are you sure?” indicate Yes, and press **Enter**.
8. The **Creating Kernel** message appears, followed by the Move Kernel Message. Select **OK** and press **Enter**. The system reboots.
9. Proceed to [“How to verify a device file” on page 901](#).

How to verify a device file

To verify a device file:

1. Verify that the *spt* was successfully installed with the following command:


```
ioscan -kfn
```
2. 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 has been completed properly.
3. If the device entry was defined by the operating system, use the OS-defined entry and proceed to verify the installation.

Major number

To determine the value for *majorum*, type the following commands:

```
lsdev -d sctl
```

The output should resemble the following. The assigned number may differ from those displayed in this example:

Character	Block	Driver	Class
HP-PB	75 -1	spt	spt
HSC or PCI	203 -1	sctl	ctl

The value for *majorum* is the number in the Character column.

Minor number

To determine the value for *minorum*, use the **ioscan** command. 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:

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:

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					

How to test 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:

inquire

NOTICE

Use the inquire command with caution. Running inquire sends the SCSI inquiry command to all devices detected on the SCSI bus. Using inquire 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 0.14Autochanger (Jukebox), /dev/rac/c0t1d0
scsidev@0.2.0:Quantum DLT4000 CC37Tape, /dev/rmt/c0t2d0BESTnb
scsidev@0.3.0:Quantum DLT4000 CC37Tape, /dev/rmt/c0t3d0BESTnb
scsidev@0.4.0:Quantum DLT4000 CC37Tape, /dev/rmt/c0t4d0BESTnb
scsidev@0.5.0:Quantum DLT4000 CC37Tape, /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 you attach a Tape drive to the HP-UX 11i V2 64-bit host and run the inquire command, the tape drive is not detected, even if the device is configured, labeled and mounted and a save was successful.

Workaround

Identify the drive path in the /dev/rmt folder, and using this path configure the device, as usual.

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 will rebuild 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 types of error messages may appear in the specified situations:

- ◆ When the **nsrmm** or **nsrjb** command is used to label the tape:

```
nsrmm: error, label write, No more processes (5)
```

- ◆ When the **scanner -i** command 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 attached to an HP-UX server or storage node, all of the tape drives inside the autochanger are unloaded to their respective slots. To unload a single drive to its corresponding slot, use the **nsrjb -u -f devicename** command instead.

SCSI pass-through driver required for HP-UX autochangers

If an autochanger with a NetWorker HP-UX server is used, refer to the *NetWorker Installation Guide*. Read the required procedures to follow before the **jbconfig** program is run. Even if the SCSI pass-through driver is installed, follow the procedures to rebuild the kernel. Then run the **jbconfig** program to configure the autochanger.

Symbolic link entries in the fstab file

For HP-UX operating systems, do not use symbolic entries in the `/etc/fstab` file. If symbolic links are used in the `fstab` file, the NetWorker server will *not* back up the file system that the symbolic link points to.

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 savegroup. Use the korn shell instead (`/bin/ksh`).

AIX

This section provides information specific to NetWorker software that runs on the AIX platform.

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

The operation of the LUS driver on AIX has been changed with NetWorker release 7.6 SP2. When a library comes online, NetWorker now obtains an exclusive lock on the library. This lock is maintained as long as the library is enabled. As a result, diagnostic tools such as **inquire** and the **sjj** utilities cannot be used to access the library during this time. To access the library using these tools, it will be necessary to take the library offline.

Recovering set-group-id or setuid binaries and files

On AIX, non-root users who are performing a recovery, will not be able to restore group ownership (the `set-group-id-on-execution` or `setuid` permission bit) on binaries or files. This is expected behavior.

APPENDIX F

MAC OS X Support

This appendix covers these topics:

- ◆ Support for Mac OS X 906
- ◆ Mac OS X backup considerations 906
- ◆ Recovering files and directories on Mac OS X using the command prompt 908
- ◆ Recovering files and directories on Mac OS X using NetWorker Recover 909

Support for Mac OS X

This section describes NetWorker client support for the Mac OS X platform. Mac hosts can be set up as NetWorker clients by using the NetWorker client for Mac OS X. Mac hosts that are set up as NetWorker clients can be backed up and restored by using any supported NetWorker server on UNIX, Linux, or Windows. Currently the NetWorker server and the Console server are not supported on Mac OS X.

Mac OS X metadata support

The NetWorker client on Mac OS X supports backup and recovery of all file system metadata including:

- ◆ Finder information
- ◆ Resource forks
- ◆ Extended attributes
- ◆ Access Control Lists

Supported file systems

The NetWorker client for Mac OS X software supports these file systems:

- ◆ HFS+ (including journaled)
- ◆ HFS
- ◆ UFS

Mac OS X backup considerations

Use this section to help plan successful backups for NetWorker clients on the Mac OS X platform.

Scheduling a NetWorker client backup on Mac OS X

This section provides information on configuring backups for a NetWorker client on Mac OS X.

MAC OS X required directives

To ensure a consistent state after recovery, certain files and directories must not be backed up on Mac OS X systems.

To ensure that appropriate files and directories are not backed up:

1. Create or edit the Mac OS Client resource. [“Task 6: Create a backup Client resource” on page 64](#) provides information about creating a Client resource.
2. Select one of these directives from the Directive attribute list:
 - Mac OS Standard Directives
 - Mac OS with Compression Directives

[“Preconfigured global directive resources” on page 294](#) provides more information about Mac OS directives.

3. Click **OK**.

[“Scheduled backups” on page 58](#) provides more information about scheduling a backup.

Backing up Mac OS X server’s open directory for disaster recovery

This section describes how to back up the Mac OS X Server’s Open Directory. Open Directory contains system configuration information that is essential for disaster recovery. The *NetWorker Procedure Generator* provides more information about disaster recovery.

NOTICE

The NetWorker Mac OS directives do not back up Open Directory database files.

To ensure complete protection of a Mac OS X Server system in the event of a catastrophic failure:

1. Use the **savenpc** script to automatically export and backup the required Open Directory database files.
2. Open Directory database files remain available during the backup.

To automatically back up Open Directory files:

1. Enter **savenpc** in the Backup Command attribute when configuring the Mac OS X client as a NetWorker Client resource.

[“Using the savenpc command with a customized backup program” on page 124](#) provides more information about enabling the Client resource to use the **savenpc** command.

2. Create a custom **savenpc** script in the `/nsr/res` directory with the name `<group_name>.res`.

where `<group_name>` is the Group that was selected for the Client resource.

3. Include entries in the **savenpc** script to perform these functions:

- Back up Open Directory's LDAP directory domain:


```
# slapcat -l /var/backups/networker.ldif
```
- If your LDAP server uses SSL, back up Open Directory's Password Server database:


```
# mkdir -p /var/backups/networker.odpdb
# mkpassdb -backupdb /var/backups/networker.odpdb
```
- Back up the local NetInfo directory domain:


```
# nidump -r / . > /var/backups/networker.nidump
```

Example 71 Custom savenpc script for Mac OS X

A Mac OS X NetWorker client that belongs to the Default group will have a `/nsr/res/Default.res` script with this content:

```
type: savenpc;
precmd: "/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"
```

In this script, the **savenpc** command backs up Open Directory's LDAP directory, Password Server, and NetInfo databases before each scheduled save.

Performing a manual backup on Mac OS X

Manual backups for Mac OS X clients must be performed from the command prompt. To perform a manual backup, use the **save** command, in a Terminal session, as follows:

```
$ save "file_or_directory_to_back_up"
```

By default, the **save** command contacts the NetWorker server defined in the `/nsr/res/servers` file that comes first in alphabetical order.

To specify an alternative NetWorker server, use the **save** command with the **-s** `NetWorker_server` option.

Recovering files and directories on Mac OS X using the command prompt

These sections provide information on recovering individual files and directories from a NetWorker client on Mac OS X using the command prompt:

- ◆ [“Task 1: Browse backed-up Mac OS X data” on page 909](#)
- ◆ [“Task 2: Recover individual files or directories” on page 909](#)

Each task in this section uses the NetWorker **recover** command. The *NetWorker Command Reference Guide* or the UNIX man pages provide more information about this command.

Task 1: Browse backed-up Mac OS X data

To browse backed-up Mac OS X data:

1. From the Mac OS X Terminal application, start a recover session with a NetWorker server by using this command:

```
$ recover
```

By default, the **recover** command contacts the NetWorker server defined in the `/nsr/res/servers` file that comes first in alphabetical order. To specify an alternative NetWorker server, use the **recover** command with the **-s** `NetWorker_server` option.

2. At the recover prompt, browse backed-up Mac OS X data by using common UNIX shell commands such as **cd** and **ls**.

Task 2: Recover individual files or directories

To recover individual files or directories from the client's **recover** prompt:

1. At the **recover** prompt, add all the directories and files to be recovered, by using the **add** command, for example:

```
recover> add directory_name
```

2. (Optional) To automatically overwrite existing files, enter the **force** option at the recover prompt.
3. Start the recovery by typing this command:

```
recover> recover
```

NOTICE

Do not recover any Mac OS X operating system boot files. For example do not recover the Mac OS X operating system kernel, `/mach_kernel`.

Recovering files and directories on Mac OS X using NetWorker Recover

NetWorker Recover is the client interface program that can be used to recover files from a NetWorker server. As a NetWorker client, your computer can connect to a server to recover files and save sets.

Starting NetWorker Recover for the first time

When starting NetWorker Recover for the first time, the Connect to Server dialog appears:




Figure 64 Connect to Server dialog

In the dialog, you have the option to manually enter the server address, or select a server from the list of Available servers. This list is generated from the hostnames in the `/nsr/res/servers` file.

If no servers are found in the `/nsr/res/servers` file, you can build a list of servers on the network by entering the server address and clicking Update in the Connect to Server dialog.

Once you have selected the server, click **Connect** to initiate a browse session for the local host.

After a successful connection, the NetWorker Recover window appears.

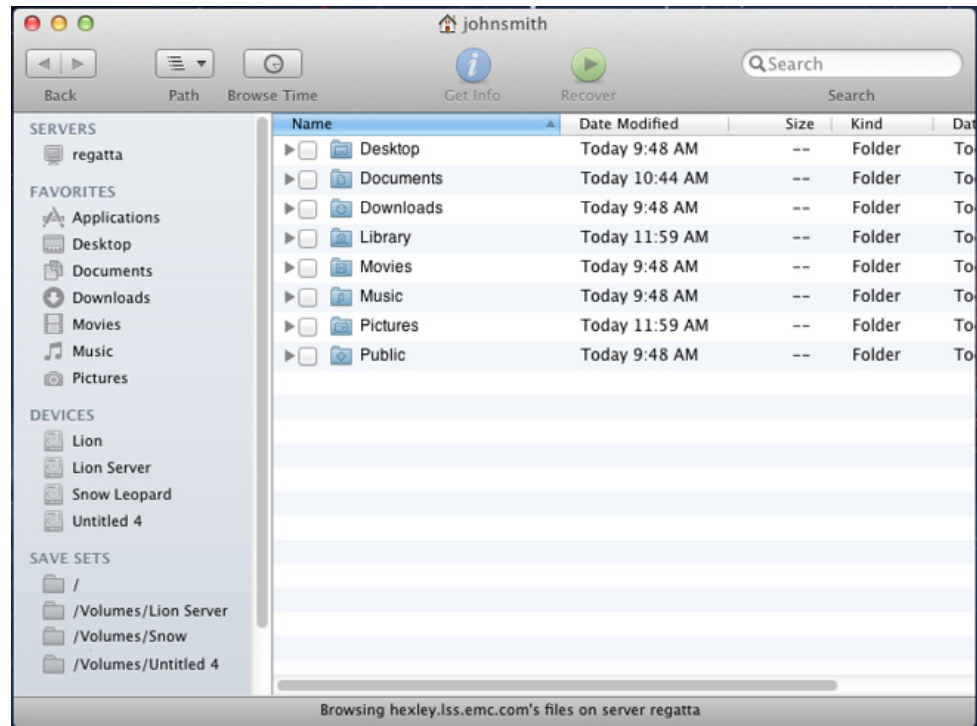


Figure 65 NetWorker Recover window

Navigating the NetWorker Recover window

The NetWorker Recover window consists of a browse session view, along with the following menu commands and toolbars.

Menu Commands

The following menu commands are available in the NetWorker Recover window.

- ◆ **File > Get Info** - displays an information dialog of the selected view object
- ◆ **File > Recover** - initiates a recover command
- ◆ **File > Required Volumes** - when file indices or save sets are marked for recovery, selecting this command launches a dialog that displays which volumes are required to recover those files
- ◆ **File > Find...** - when selected, this command will switch the browser view to the index search view
- ◆ **Edit > Mark/Unmark File for Recover** - toggles the mark of a selected file or save set
- ◆ **View > Show/Hide Hidden Files** - toggles the visibility of hidden files
- ◆ **View > Show/Hide File Versions** - toggles the visibility of the index versions sidebar
- ◆ **View > Monitor Server** - displays the NetWorker Server Monitor dialog
- ◆ **Go > Browse Time** - sets the browse time to a user supplied date
- ◆ **Go > Browse Client** - initiates a browse session with a different client

- ◆ **Go > Connect to Server** - opens the Connect to Server dialog where you can enter the address for the server you want to connect to
- ◆ **Window > Recover Log** - select this command to view the previous recover's log

Toolbar

The following toolbar buttons and options are available for frequently used commands:

- ◆ **Back** - navigate to previously viewed folders
- ◆ **Path** - lists the folders hierarchy of the currently browsed directory
- ◆ **Browse Time** - displays the Browse Time view, along with a field that, when selected, allows you to change the browse time.
- ◆ **Get Info** - opens a dialog displaying information about the currently selected object
- ◆ **Recover** - starts a recover session
- ◆ **Search** - searches file indexes for the given string. When you enter text in the Search field, the Browse session view displays the results. You can then refine the search using the Search filter bar (for example, you can search based on the last backup time, or combine search criteria).

Each control is enabled or disabled based on the browse session's current state. For example, the Recover button is only enabled when either file indices or save sets have been marked for recovery.

Browse session view

The NetWorker Recover window browse session view contains two panes, similar to the Mac OS X Finder. The bottom of the window displays a status message that identifies the current state of the browse session (for example, the current connection status, or the number of files marked for recover).

Any object displayed in the sidebar or browser views can have its basic information queried by highlighting the icon and then right-clicking and selecting Get Info from the drop down, by clicking Info on the toolbar, or by clicking File > Get Info from the application menu.

Sidebar

The sidebar contains the following:

- ◆ **SERVERS** - displays each previously connected or currently connected server. When a specific server is highlighted, a connection to that server is attempted, if it is not already connected.
- ◆ **DEVICES** - displays each file system backed up by NetWorker
- ◆ **FAVORITES** - displays any entry that has been backed up
- ◆ **SAVE SETS** - displays each unique save set. You can select a save set to reveal each instance of that save set, including cloned save sets, in the browser view.
- ◆ **RECOVER SETS** - All marked file indices and save sets are added to a recover set. There are always two recover sets to separate index and save set recovers.

Folder icons, when expanded, display the folder contents. Folder icons that are double-clicked will refresh the browser session's right pane with its contents.

Selecting a specific server from the **SERVERS** category displays basic information about the NetWorker Server and lists each client configured on that server in the browser view. If the connection to the selected server is active, the **Monitor...** button appears below the computer icon in the Details pane. If the connection to this server is not currently established, the **Connect...** button displays:




Figure 66 Browse session

Selecting entries from the **DEVICES** or **FAVORITES** categories display their contents in the browser view.

Selecting a specific save set from the **SAVE SETS** category displays the instances for that save set in the browser view.

Selecting **RECOVER SET** displays all entries marked for recovery in the browser view, as well as summary information for the recover set.

Right pane (browser view)

In the right pane, you can browse and navigate the file index and save sets for categories selected in the sidebar. Files and save sets are presented in the Finder's list view format. The browser lists the following information for file indices and save sets:

- ◆ File Index includes Name, Date Modified, Size, Kind, and Backup Date
- ◆ Save Sets includes Save Set ID, Save Set Level, Save Set Status, Size, Files, and Backup Date

Configuring NetWorker Recover

The following sections provide the steps required to configure NetWorker Recover to perform a recovery operation using the NetWorker Recover window.

Note: To perform a default recovery, where files from the most recent backup are recovered to their original location using their original filenames, skip to Step 6.

- ◆ [“Changing the NetWorker Server” on page 914](#)
- ◆ [“Changing the NetWorker Client” on page 914](#)
- ◆ [“Step1: Selecting files and save sets for recovery” on page 915](#)
- ◆ [“Step 2: Search for files” on page 915](#)
- ◆ [“Step 3: Check the status of required volumes” on page 916](#)
- ◆ [“Step 4: Viewing file versions” on page 916](#)
- ◆ [“Step 5: Changing the browse time” on page 916](#)
- ◆ [“Step 6: Starting the Recovery” on page 916](#)
- ◆ [“Monitor the Server” on page 917](#)
- ◆ [“Recover logging” on page 917](#)

Changing the NetWorker Server

If the NetWorker server indicated in the Status Bar is not the server you want to recover from, change it to a different host by selecting **Go > Connect to Server** from the menu. The Connect to Server dialog displays, where you can select the server from the list of Available Servers and click **Connect**, or enter the address of the server you want to connect to, and click Update. You can then select this server from the list of Available Servers, and click **Connect**.

Upon exiting NetWorker Recover, the last server connection made is cached in User Preferences.

Changing the NetWorker Client

After selecting the server, to change NetWorker clients, select **Go > Browse Client** from the menu, and select a client from the dropdown list. This lists the hostname of each client computer known to the NetWorker server. To specify a client as the source for the data you want to recover, select the name of the desired client in the list. A browse session is then initiated for the client.

Clients are displayed in the lower half of the browser view. The Clients filter bar is located above the client names so you can limit the list of clients to a particular view. In the following graphic, the Client filter bar displays All, Mac OS X Clients, Windows Clients, and UNIX Clients.




Figure 67 Clients

Step1: Selecting files and save sets for recovery

1. Select **RECOVER SETS** in the sidebar of the NetWorker Recover window and then select **Files** to display files in the browser view. Mark files for recovery by selecting the checkbox next to each file you want to recover.
2. Select **SAVE SETS** in the sidebar of the NetWorker Recover window and then select a save set name to display instances of that save set in the browser view. Mark instances for recovery by selecting the checkbox next to each instance you want to recover. Only one instance (or clone) of a save set can be marked for recovery at a time. Save set instances or cloned save sets can be filtered using the save set filter bar, which works similar to the Client filter bar.

Each marked item will be added to RECOVER SETS and separated depending on whether the item is a file index or save set. A number next to each RECOVER SET in the sidebar displays the number of items selected for recovery. Selecting Files or Save Sets from the RECOVER SETS category displays the files or save sets that have been marked for recovery, in addition to summary information related to the RECOVER SET.

Step 2: Search for files

If the files you want to recover are not visible in the browser view, you can search for files using the Search field in the top right of the NetWorker Recover window. Entering text in the search field will automatically search the backed up file indices for the text specified in the field. Successful matches are displayed in the browser view.

Search results can be marked for recovery and navigated similar to other browser entries. Once a search result is selected, the path displayed at the bottom of the browser view lists the file's folder hierarchy.

The scope of the search can be controlled by clicking the **Search Scope** bar. All of the client's indices and the folder currently being browsed are listed as scope options.

You can also specify a folder you want to search for by selecting **Go > Go to Folder...** from the menu and entering the path and name of the folder. The folder then displays in the browser view.

Step 3: Check the status of required volumes

To view the status of the backup volumes required for the recovery operation:

1. Highlight Files under **RECOVER SETS** in the sidebar. The summary information contains the names of any required volumes.
2. Click the **Volume Status** button located to the right of the summary information. The Required Volumes dialog appears.

Ensure that the status of the required volume(s) indicates on-line, then close the dialog.

Step 4: Viewing file versions

To view a list of all versions of a selected file, go to View > Show File Versions from the application's main menu. The versions sidebar appears.

Selecting a file displays all backed up versions of the file.

To recover a specific version of a file, you can:

- ◆ Drag and drop the file to the browser view
- ◆ Drag the file to a folder for recovery
- ◆ Right-click the file to select for recovery.

Step 5: Changing the browse time

If you want to browse or recover files from an earlier backup, the recovery browse time can be changed using several methods:

- ◆ Browser view context menus - in each browser view, right click on a file index icon to display the Set Browse Time To... menu item.
- ◆ The Browse Time toolbar button, when clicked, displays the Browse Time view, displaying the current browse time and the option for later or earlier browse times.
- ◆ Go > Browse Time displays the Change Browse Time dialog.

Step 6: Starting the Recovery

To start the recovery:

1. From the menu, select **File > Recover**, or click the **Recover** button in the toolbar. The Recover Options dialog appears.
2. Specify where to relocate the recovered files to.
3. Select an option for conflict resolution in the event a conflict occurs between a local file and a file being recovered. If you select Prompt me for an action, a prompt will appear each time a conflict is encountered.
4. Select **Exit** when an error is encountered and you want to stop the recovery.
5. Select a host from the drop-down to direct the recovery to.
6. Click **OK**.

The recover status dialog appears. At any time during the recovery, you can click the Stop button to cancel the operation.

Monitor the Server

From the recover progress dialog, you can click the **Monitor Server...** button to launch the NetWorker Monitor dialog.

The NetWorker Monitor dialog displays the following tabs:

Info -- general server information including name, IP, OS type, NetWorker version, Save totals and Recover totals)

Messages -- server messages logged during the recovery (including 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

Recover logging

When the recovery completes, the success or failure of the operation is reported by NetWorker Recover. You can view the recover log by clicking the **Recover Log** button. Upon clicking this button, the Console application is launched, displaying the log that was written to the user's ~/Library/Logs/recover.log file.

APPENDIX G

Direct SCSI Backup and Recover

This appendix covers these topics:

- ◆ Introduction to direct SCSI backup and recover 920
- ◆ System requirements 920
- ◆ Performing direct SCSI backup 921
- ◆ Performing direct SCSI recover 924
- ◆ Licensing 927

Introduction to direct SCSI backup and recover

Direct SCSI backup and recover enables direct backup and recover of Small Computer System Interface (SCSI) devices without the requirement of mounting it on the backup host if an access path is available to these devices over a Storage Area Network (SAN). You can also use this feature to migrate EDM servers to the NetWorker software to perform backup and recover of Business continuance volume (BCV) devices on a Symmetrix server (as well as backup and recover of raw devices) over a SCSI bus.

Backup technologies often protect information as files and directories or as file systems. But backups also allow information contained on raw disks to be protected as raw devices. However, this type of backup (known as raw backup) usually does not provide granular recover capabilities.

The direct SCSI backup and recover feature enables raw backups for the NetWorker software directly by using a SCSI target, which is usually accessible from a SAN proxy host. Typically, in an EMC Symmetrix[®] storage environment, these devices can be viewed from a primary application host and from a proxy backup host. The direct SCSI backup and recover feature allows you to protect BCV devices from a proxy backup host as a raw backup.

Be aware that since backup and recover is performed on the proxy client (which is also a storage node), and a Symmetrix device is accessible, data is accessed from the proxy client, which may not be the client that originally created the data

System requirements

Before the configuration of Direct SCSI backup and restores, review the following list of requirements:

- ◆ The direct SCSI feature is only supported on Solaris SPARC storage nodes
- ◆ EMC Solution Enabler version 5.5 or later must be installed.
- ◆ The following hardware devices are supported:
 - Raw device path of a SCSI device
 - Sym

Unsupported features

If performing backup and recover over a SCSI bus, the following features are not supported with NetWorker:

- ◆ Archiving
- ◆ Save set consolidation
- ◆ Index browsing
- ◆ Conventional recovery by using the command line recover utility

- ◆ File-by-file recovery
- ◆ Conventional save set recovery by using the command line recover utility

The EMC Symmetrix device is supported for use with SCSI backup and recover, but no other vendor device is supported with this feature as of NetWorker release 7.4.

Performing direct SCSI backup

When performing Direct SCSI backup, you can perform either of the following:

- ◆ Single device backup
- ◆ Backup of a device set that consists of a list of devices specified in a resource file.

If the device is vendor-specific, the program loads the vendor-specific plug-in shared library.

The NetWorker software uses the `nsrscsi_save` program to start the backup thread for each device to be backed up. Each backup thread performs the following operations:

- ◆ Finds the host accessible raw device path for the given vendor device.
- ◆ Starts a save session with the NetWorker server.
- ◆ Runs **scsi asm** on the raw device path to move data from the SCSI device to a storage node (**nsrmmmd**) by using SCSI commands.

NOTICE

Before starting a backup, set the backup device to offline or read-only mode, and the file systems that reside on the device to read-only. If the device is a Symmetrix BCV, keep the BCV detached from the standard during the backup.

Backing up data on a Symmetrix BCV device

To back up data on a Symmetrix BCV device:

1. Create the `.res` file (for example, `/nsr/res/deviceset.res`) if performing backup on a device set. Within the file, do the following:
 - a. Specify a list of devices to be backed up as part of a device set. Associate the device set with each entry in the file. You can identify the devices by their Symmetrix volume IDs (SYMMIDs) and category names, as in the following example:


```
000182504581/011 ## These two will be
000182504581/012 ## grouped as OracleDisks
# This is a comment line
000182504581/07D ## These two will be
000182504581/07E ## grouped as ExchDisks
```
 - b. Store this file in the `nsr/res` directory on the storage node.

2. For the AppHost or the host that controls the data, perform either of the following steps:
 - Create a new NetWorker Client resource. [“Task 6: Create a backup Client resource” on page 64](#) provides more information.
 - Edit an existing client by right-clicking the client in the in the **Configuration** screen of the **Administration** window and selecting **Properties**.
3. For the **Save Set** attribute, do one of the following:
 - If performing multiple device backup, type the following:


```
<<emc_symm>>/{deviceset.res}
```

where <<emc_symm>> is the name of the device and {*deviceset.res*} is the name of the .res file.
 - If performing single device backup without the .res file, type the following:


```
<<emc_symm>>
```

where <<emc_symm>> is the name of the device.
4. On the **Apps and Modules** tab, select **SCSI** for the **Proxy backup type** attribute.
5. In the **Proxy Backup host** attribute, enter the name of the storage node on which the nsrscsi_save command will be run.
6. Click **OK**.

Once you create the Client resource, the **nsrscsi_save** program reads the .res file for all the device IDs listed or reads the ID for the single device, and starts the backup thread. The backup thread performs the following:

- ◆ validates the SYMMIDs
- ◆ finds the host-accessible raw device path for the SYMMIDs
- ◆ starts the save session with the NetWorker server.

The NetWorker software creates a save set for each category name, but does not create an index for the content within the BCV.

Backing up data on a raw device

The process for backing up raw device data is similar to backing up data on a BCV device. To perform a SCSI backup of a raw device:

1. For the AppHost or the host that controls the data, perform either of the following steps:
 - Create a new NetWorker Client resource. [“Task 6: Create a backup Client resource” on page 64](#) provides more information.
 - Edit an existing client by right-clicking the client in the in the **Configuration** screen of the **Administration** window and selecting **Properties**.
2. For the **Save Set** attribute, type the path for the raw device with the device name. For example, if the device name is c1t2d0s2 and the path is */dev/rdisk*, type the following:


```
/dev/rdisk/c1t2d0s2
```

3. On the **Apps and Modules** tab, select **SCSI** for the **Proxy backup type** attribute.
4. In the **Proxy Backup host** attribute, enter the name of the storage node on which the `nsrscsi_save` command will be run.
5. Click **OK**.

Once you create the Client resource, associate it to scheduled group. At the scheduled time, the group starts the `nsrscsi_save` program on the storage node.

The `nsrscsi_save` program starts the backup thread. For a single device backup, a single backup thread is created by `nsrscsi_save`. The program starts the save session with the NetWorker server.

Backing up data from the command line

To perform backup from the command line, run the following command:

```
nsrscsi_save [ -c clientname ] [ -g group ] [ -N save-set-name ]
[ -I input filename ] [ -s server ] [ -b pool ] [-e expiration ]
[ -y retention time ] Path
```

where:

- ◆ `-c clientname` is the client name for starting the save session. The client name is the local host by default. If using the local host, you do not need to specify the client name.

Note: The client-name is not necessarily the host accessible device. For example, a Symmetrix BCV device may be accessible on a different host for backup than the client where the standard device is attached (the client that you want to register the backup against).

- ◆ `-N save-set-name` is the name of the save set. The save set name is the pathname by default. If the pathname is the name of the device set, `-N` is ignored.
- ◆ `-I input-filename` is the filepathname for the file that contains the list of devices to be backed up (for example, `/tmp/testdisks.res`). If the input-filename is not specified, then the default input-filename is taken from device-set-name. For example, if the device-set-name is `oracledisks`, the input-filename would be `/nsr/res/oracledisks.res`. Also, if `device-set-name` is used in the path but `-I` is not specified, the default location is set to `/nsr/res/device-set-name.res`.

Note: The input file should contain only entries for the devices that need to be backed up. Multiple device entries should be separated by a new line. When `-I` is specified, device-set-name should also be specified.

- ◆ Path can be any of the following formats:
 - For a raw device path: /dev/rdisk/c1t2d0s2
 - For a device-set-name: {OracleDevices}

Note: Braces are required to distinguish device-set-name from a single device path.

- ◆ `-g group` is used by **savegrp** and **savefs** to specify the group of the save. It is also used by the NetWorker server to select the specific media pool.
- ◆ `-b pool` specifies a particular destination pool for the save. All the save sessions go to the same pool.

Performing direct SCSI recover

Direct SCSI recover is performed by using the **nsrscsi_recover** program. This program starts the recover thread for each save set to be recovered. From the command line, you can recover single save sets, or multiple save sets specified in a resource file, with a unique destination for each save set. If the device being recovered from is vendor-specific, the program loads the vendor-specific plug-in DLL.

Each recover thread performs the following operations:

- ◆ Finds the host accessible raw device path for the given target vendor device.
- ◆ Starts a recover session with the NetWorker server.
- ◆ Runs **scsi asm** on the raw device path to move data from the storage node (nsrmmmd) to a raw device via SCSI CDB commands.

NOTICE

Before starting the recovery, set the backup device to offline or read-only mode from the application host, and the file systems that reside on the device to read-only. If the device is a Symmetrix BCV, keep the BCV detached from the standard during the backup. Also, note that the data on the target device being used for recovery will be rewritten and the original data will be lost when using the target device ID and raw device path for recovery.

Recovering data to a Symmetrix BCV device

Recovering data on a Symmetrix BCV device must be performed from the command line. You cannot use NMC to perform this function.

To recover a single save set:

1. Perform an **mminfo** query on the device set name, for example:

```
mminfo -avVot -q "name=xxx"
```

The query returns a list of Save Set IDs (SSIDs) that are part of the device set.

2. Find the SSID (for example, 3697521281), and then run the following command:

```
mminfo -aS -q "ssid=3697521281"
```

The list of connected save sets displays. The connected save sets are the save sets that are backed up as part of this device set.

3. Select a save set from this list.
4. If applicable, prepare the target device and retrieve the target vendor device ID. The target vendor device ID can be original.
5. Run the following command:

```
nsrscsi_recover -S ssid -T target device
```

where *ssid* is the ID for the save set being recovered and *target device* is the SYMMID and the device ID (for example, 0034567/0366, where 0034567 is the SYMMID and 0366 is the device ID).

The NetWorker software recovers the contents of the save set to the destination location. The destination location is the original location by default.

To recover multiple save sets:

1. Perform an mminfo query on the device set name. For example:

```
mminfo -avVot -q "name=xxx".
```

The query returns a list of Save Set IDs (SSIDs) that are part of the device set.

2. Find the appropriate SSID (for example, 3697521281), and then run the following command:

```
mminfo -aS -q "ssid=3697521281",
```

The list of connected save sets displays. The connected save sets are the save sets that are backed up as part of this device set.

3. Choose the SSIDs from the list of connected save sets (in the following example, the SSIDs are 3697521281 and 3680744065):

```
mminfo -avVot -r "volume,name,savetime(25),ssid" |grep oraclediskset
scip2b081.networker.com.001 oraclediskset:000187910217/0365
11/21/06 06:14:25 PM 3697521281
scip2b081.networker.com.001 oraclediskset:000187910217/0366
11/21/06 06:14:26 PM 3680744065
```

4. If applicable, prepare all target devices and retrieve the target vendor device IDs. The target vendor device ID can be the same as the original backup device ID.

5. Create a .res file under /nsr/res (for example, /nsr/res/restorelist.res), and specify an entry for each save set to be recovered.

Each entry in the .res file must have an SSID for every save set being recovered, mapped to a target device (SYMMID/Device ID), as in the following example:

```
3697521281=>000187910217/0366
3680744065=>000187910217/0365
```

6. Run the **recover** command:

```
nsrscsi_recover -I input filename
```

where *input filename* is the name and location of the .res file (for example, nsr/dev/restorelist.res).

Once you create the Client resource, the **nsrscsi_recover** program reads the .res file for all the device IDs listed, and starts the recover thread. The recover thread validates the device IDs, finds the host-accessible raw device path for the IDs, and starts the recover session with the NetWorker server.

Recovering data to a raw device

To recover data on a raw device:

1. Perform an **mminfo** query and select an SSID (for example, 3697521281), as in the following example:

```
mminfo -aVvot
volume      client size level name
ssid save  time      date      time      browse  clretent
first      last file  rec volid      total fl

scip2b081.networker.com.001 scip2b081.networker.com 8839 MB full
/dev/rdisk/c1t1d0s2
3697521281 1164161665 11/21/06 06:14:25 PM 12/21/06 11/21/07
0 9051405795 0 0 3731075692 9051405796 cr
```

2. If applicable, prepare the target device and retrieve the target raw device path. The target raw device path can be the same as the original raw device path.
3. Run the following command:

```
nsrscsi_recover -S ssid -T target device
```

where *ssid* is the ID for the save set being recovered and *target device* is the raw device path (for example, /dev/rdisk/c1t1d0s2).

The NetWorker software recovers the contents of the save set to the destination location. The destination location is the original location by default.

NOTICE

The destination path for recover must be specified and it must be a raw device or a vendor device.

Licensing

EMC Solution Enabler version 5.5 or later is required to use the Direct SCSI feature. There are no other licensing requirements for this feature.

APPENDIX H

Security Configuration Settings

This appendix covers these topics:

- ◆ Access control settings..... 930
- ◆ Log settings 932
- ◆ NetWorker Accountability 934
- ◆ Communication security settings 943
- ◆ Encrypting backup data..... 944
- ◆ Federal Information Processing Standard Compliance 944

Access control settings

Access control settings protect resources against unauthorized access.

User authentication

User authentication settings control the process of verifying an identity claimed by a user when accessing the product.

Default accounts

[Table 145 on page 930](#) describes the default login accounts.

Table 145 Login accounts

User account	Description
root@localhost — for NetWorker server on UNIX platforms	User ‘root’ on the NetWorker server host is automatically added to the Administrator list.
system@localhost — Windows platforms	User ‘root’ on the NetWorker server host is automatically added to the Administrator list.
@	All users at all hosts are added to the ‘users’ attribute of an instance of the ‘NSR Usergroup’ resource. All privileges associated with that ‘NSR Usergroup’ instance go to ‘all users @ all hosts’ with no explicit denial of access.
administrator	Default NMC user. First login procedure forces password change.

Authentication configuration

The NMC Console has two modes of authentication: native NMC authentication mode and LDAP mode. By default, NMC user authentication is set to native mode.

You can set up user authentication in LDAP mode by using the Configure Login Authentication wizard. You can also revert back to native NetWorker user authentication by using the wizard.

[“NMC server authentication” on page 504](#) provides more information.

User authorization

User authorization settings control the privileges that are granted to a user when accessing a resource managed by the product. [“NetWorker User Groups” on page 559](#) describe how to set up users and user groups on the NetWorker server.

[“Managing server access” on page 558](#) describes how to set up users and to control user permissions for the Console server.

Component access control

Component access control settings define the control over access to the product by external and internal systems or components.

Component authentication

NetWorker hosts and daemons are authenticated by using the nsrauth mechanism, which is available for hosts that run NetWorker release 7.3 or later. The nsrauth authentication mechanism is a strong authentication and is based on the Secure Sockets Layer (SSL) protocol provided by the OpenSSL library or RSA BSAFE SSL, depending on the platform.

Each NetWorker host has a nsrexecd service, which provides authentication services. Each nsrexecd service has its own private key and self-signed certificate for authentication. The private key is generated by nsrexecd when it starts. The private key can also be loaded from a file. The corresponding self-signed certificate is generated by the private key. The private key is RSA and is 1024 bits in length. The encryption method that is used after an SSL session is set up is AES-128.

The session information sent over the SSL connection includes:

- ◆ Session keys
- ◆ Session ID
- ◆ User's information
- ◆ User's NetWorker permissions

[“NetWorker authentication” on page 612](#) provides more information about configuring nsrauth authentication.

Component authorization

NetWorker uses the contents of the `/nsr/res/servers` file for UNIX or the `NetWorker_installation_path\res\servers` file for Windows, on each NetWorker client to control the client-tasking rights.

The client-tasking rights are the rights to request the execution of a program on another client and might be any of the following:

- ◆ Server that performs an archive request
- ◆ Scheduled backup
- ◆ Another client that requests a directed recover

If the server file is empty, then any NetWorker host can have tasking rights.

- ◆ Add the names of the additional NetWorker servers to the server file so that the client with the tasking rights can back up to other NetWorker servers.
- ◆ Add the client names to the servers file so that other clients can perform directed recovers to the client with the tasking rights.
 - You can add the names of NetWorker servers to the server file during the software installation.
 - To add additional hosts later, use a text editor and add the hostnames to the server file.
 - After adding the additional hosts in the server file, restart the nsrexecd on that client to enable permissions for the additional server hosts.

Log settings

A log is a chronological record that helps to examine the sequence of activities surrounding or leading up to an operation, procedure, or event in a security-related transaction from beginning to end.

Log files and their descriptions

[Table 146 on page 932](#) shows the log files location.

Table 146 Log files (1 of 2)

Component	Default Location
NetWorker server and client daemons	UNIX: /nsr/logs/daemon.raw Windows: <NetWorker_install_path>\nsr\logs\daemon.raw
NetWorker server generated syslog messages and daemon.notice	UNIX: OS log file defined by system log configuration file Windows: <NetWorker_install_path>\nsr\logs\messages
NetWorker server generated syslog messages local0.notice and local0.alert	UNIX: OS log file defined by system log configuration file. Unlike previous versions of the NetWorker software, NetWorker 8.0 and later does not modify the syslog.conf file to configure local0.notice and local0.alert. Refer to Vendor specific documentation to configure local0.notice and local0.alert.
NetWorker server disaster recovery command line wizard, nsrdr program	UNIX: /nsr/logs/nsrdr.log Windows: <NetWorker_install_path>\nsr\logs\nsrdr.log
Windows BMR logs on the recover host (Windows only)	Windows: X:\Program Files\Legato\nsr\logs\ <ul style="list-style-type: none"> • Ossr_director.raw • recover.log • WinPE_Wizard.log • winpe_nw_support.raw
NMC gstd logs	AIX & Linux: /opt/lgtomc/management/logs/gstd.raw Solaris: /opt/LGTONmc/management/logs/gstd.raw Windows: C:\Program Files\EMC NetWorker\Management\logs
NMC database conversion log	Solaris: /opt/LGTONmc/logs/gstdbupgrade.log AIX and Linux: /opt/lgtomc/logs/gstdbupgrade.log Windows: C:\Program Files\EMC NetWorker\Management\logs\gstdbupgrade.log
NMC Web Server logs	AIX & Linux: /opt/lgtomc/management/logs/web_output Solaris: /opt/LGTONmc/management/logs/web_output Windows: C:\Program Files\EMC NetWorker\Management\logs\web_output
NMC DB logs	AIX & Linux: /opt/lgtomc/management/logs/db_output Solaris: /opt/LGTONmc/management/logs/web_output Windows: C:\Program Files\EMC NetWorker\Management\logs\web_output
Client push log	UNIX: /nsr/logs/nsrccd.raw Windows: C:\Program Files\EMC NetWorker\logs\nsrccd.raw

Table 146 Log files (2 of 2)

Component	Default Location
Savegroup Job Logs	UNIX: /nsr/logs/sg/ <i>groupname</i> Window: C:\Program Files\EMC NetWorker\logs\sg\ <i>groupname</i>
rap log	This log file records configuration changes that are made to the NetWorker server resource database. UNIX: /nsr/logs/rap.log C:\Program Files\EMC NetWorker\logs\rap.log
security audit log	UNIX: /nsr/logs/ <i>NetWorker_server_sec_audit.raw</i> Window: C:\Program Files\EMC NetWorker\logs\ <i>Networker_server_sec_audit.raw</i>
User log (Windows only)	UNIX: /nsr/logs/networkr.raw Window: C:\Program Files\EMC NetWorker\logs\networkr.raw

Log management and retrieval

This section explains how to view and manage logs.

Viewing log files

The following log files are viewed by using the non-interactive command line program, **nsr_render_log** (for UNIX and Linux) or **nsr_render_log.exe** (for Microsoft Windows):

- ◆ Daemon log file: daemon.raw
- ◆ gstd log file: gstd.raw
- ◆ User log file: networkr.raw (for Microsoft Windows only)

[“Viewing log files” on page 803](#) provides information about using the **nsr_render_log** program.

Managing log files

The Console server gstd.raw log file is managed with the various Console server environment variables. [“Setting environment variables” on page 533](#) provides more information.

The NetWorker server daemon.raw log file is managed with various NetWorker server environment variables. [“Log file size management” on page 600](#) provides more information.

NetWorker Accountability

The NetWorker 8.0 and later software provides a centralized logging mechanism to log security related events that occur in a NetWorker datazone. This mechanism is called security audit logging.

When the NetWorker 8.0 software is installed in a datazone, each client is automatically configured to use security audit logging. Any audit logging configuration changes that are set in the NetWorker server are automatically communicated to all NetWorker 8.0 and later clients in the datazone:

- ◆ When the NetWorker server software is updated, existing NetWorker client resources are automatically configured to send security audit messages to the nsrlogd daemon.
- ◆ When new client resources are created, each client is automatically configured to send security audit messages to the nsrlogd daemon.

Examples of security audit events that generate security audit messages include:

- ◆ Console server authentication attempts.
- ◆ Account management events including password and privilege changes.
- ◆ Authorization changes including the creation or deletion of peer certificates

The following sections provide more information about security audit logging:

- ◆ [“Security audit logging overview” on page 935](#)
- ◆ [“Security audit logging configurations” on page 935](#)
- ◆ [“Security audit logging interoperability” on page 939](#)
- ◆ [“Modifying the security audit log resource” on page 939](#)
- ◆ [“Audit message format” on page 942](#)

Additional information about configuration changes to NetWorker server resources and their attributes can be found in the rap.log file. [“Monitoring Changes to NetWorker Server Resources” on page 599](#) provides more information.

Security audit logging overview

NetWorker 8.0 and later enables security audit logging by default. With security audit logging, the following applies:

- ◆ The NetWorker 8.0 and later server in each datazone contains a new resource, NSR auditlog. Use this resource to configure security audit logging.
- ◆ The NSR auditlog resource on the NetWorker server is mirrored to all NetWorker 8.0 and later clients in the datazone. The client side security audit log resource is stored in the nsrexec database. This resource provides each client with the hostname of the machine that hosts the nsrlogd daemon and the types of security audit messages to send to the nsrlogd daemon.
- ◆ The security audit messages are assigned a severity. Security audit messages that are at least as severe as the level defined in the NSR security audit log resource are recorded in the log. [“Modifying the security audit log resource” on page 939](#) describes how to change the severity level defined in the NSR auditlog resource.
- ◆ The NetWorker client processes send audit messages to the nsrlogd daemon.
- ◆ The nsrlogd daemon records the security audit messages to the security audit log file.

Security audit logging configurations

While any NetWorker 8.0 or later client in the datazone can be configured to run the nsrlogd daemon, there are certain performance and reliability advantages to using the NetWorker server for this task.

The following sections provide examples of security audit logging configurations and the advantages and disadvantages of each configuration.

- ◆ [“Single datazone — The NetWorker server hosts the nsrlogd daemon” on page 936](#)
- ◆ [“Multiple datazones — The Console server hosts the nsrlogd daemon” on page 937](#)
- ◆ [“Multiple datazones — Each NetWorker server hosts the nsrlogd daemon” on page 938](#)

Single datazone — The NetWorker server hosts the nsrlogd daemon

By default, the nsrlogd daemon runs on the NetWorker 8.0 or later server.

In this configuration, the nsrlogd daemon receives security audit messages from:

- ◆ The gstd and nsrexecd processes on the Console server.
- ◆ The nsrexecd process on each NetWorker client in the datazone.
- ◆ The daemons running on the NetWorker server.

Advantages:

- ◆ The NetWorker server daemons generate the majority of the security audit messages. In this configuration, the audit log messages are not sent over the network and will not increase network traffic.
- ◆ Security audit messages from each NetWorker client are sent to the NetWorker server. Additional network ports and routes to other networks are not required to send security audit messages.

Figure 68 on page 936 provides an example of this configuration.




Figure 68 The NetWorker server hosts the nsrlogd daemon

Multiple datazones — The Console server hosts the nsrlogd daemon

In this configuration, the nsrlogd daemon runs on the Console server and the Console server manages multiple NetWorker datazones. The Console server must be configured as a client, on each NetWorker server.

Advantages:

- ◆ Centralized logging of the security audit messages. The security audit log for each NetWorker server is stored on the Console server.

Disadvantages:

- ◆ If the nsrlogd daemon is not accessible, either because it fails or due to some message routing difficulty, security related events are not recorded.
- ◆ The NetWorker server daemons generate the majority of the security audit messages. In this scenario, the security audit log messages are sent over the network and will increase network traffic.
- ◆ Each NetWorker host in each datazone must have a route to the Console server.

Figure 69 on page 937 provides an example of this configuration.




Figure 69 The Console server hosts the nsrlogd daemon for all datazones

Multiple datazones — Each NetWorker server hosts the nsrlogd daemon

In this configuration, each NetWorker server acts runs the nsrlogd daemon and records the messages for a single datazone.

Each NetWorker client in the datazone sends security audit messages to the NetWorker server.

The Console server is:

- ◆ Configured to send security audit messages to one nsrlogd daemon.
- ◆ A client of the NetWorker server in Datazone 1.

Advantages:

- ◆ The NetWorker server daemons generate the majority of the security audit messages. In this configuration, the audit log messages are not sent over the network and will not increase network traffic.
- ◆ Security audit messages from each NetWorker client are sent to the NetWorker server. Additional routes in other networks are not required to send security audit messages.

Disadvantages:

- ◆ If the NetWorker server is compromised, the security audit log might not be accessible to review.
- ◆ Multiple security audit logs must be managed.

Figure 70 on page 938 provides an example of this configuration.

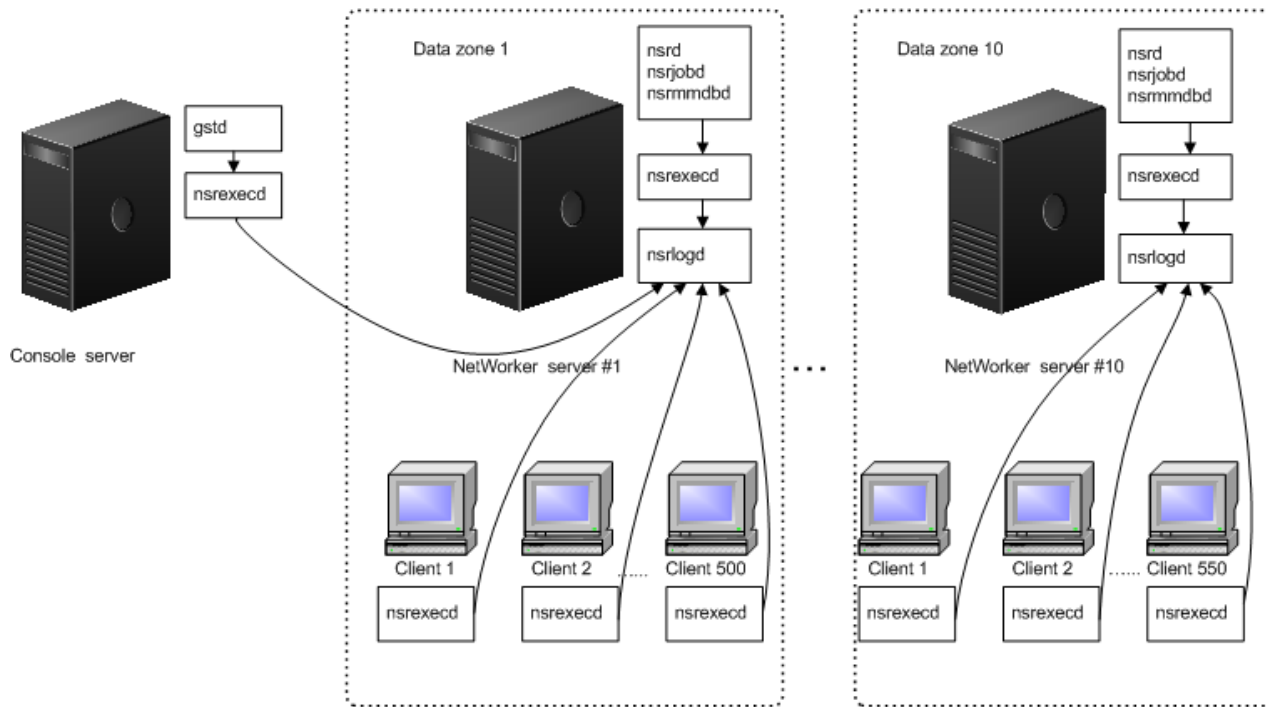


Figure 70 Each NetWorker server hosts the nsrlogd daemon

Security audit logging interoperability

The security audit log is a feature new in NetWorker 8.0. Previous versions of NetWorker do not support logging security events or hosting the nsrlogd daemon. [Table 147 on page 939](#) provides a summary.

Table 147 Interoperability matrix for security audit logging

NetWorker server version	NetWorker client version	security audit logging behavior
8.0 and later	8.0 and later	<ul style="list-style-type: none"> • Audit messages generated by the NetWorker server are logged to the nsrlogd daemon. • Audit messages generated by the NetWorker client are logged to the nsrlogd daemon.
8.0 and later	7.6.x	<ul style="list-style-type: none"> • Audit messages generated from the NetWorker server are logged to the nsrlogd daemon. • Audit message are not generated by the NetWorker client. • A NetWorker client cannot run the nsrlogd daemon.
7.6.x	8.0 and later	<ul style="list-style-type: none"> • Audit messages are not generated by the NetWorker server. • Audit messages are generated by the client but without a NetWorker 8.0 server or later, the client cannot be configured to run the nsrlogd daemon.

Modifying the security audit log resource

To modify the audit log server resource:

1. Log in to the Console server as a **Console Security Administrator**.
2. Connect to the NetWorker server.
3. In the **Configuration** Window, select **Security Audit log** in the left pane.
4. Right-mouse click the **Security Audit Log** resource and select **Properties**.
5. Optionally, specify a hostname in the auditlog hostname attribute to specify the NetWorker client that will run the nsrlogd daemon. Only clients that are defined for the NetWorker server can run the nsrlogd daemon.

NOTICE

All security audit logging is disabled if the specified hostname does not support audit logging.

6. Optionally, specify a valid path on the audit log server in the **Auditlog filepath** attribute.

This changes the location of the security audit log file.

The default location is `/nsr/logs` on a UNIX Audit Log server and `<NetWorker_install_path>\nsr\logs` on a Windows Audit Log server.

7. Optionally, change the maximum size of the security audit log in the **auditlog maximum file size (MB)** attribute.

When this maximum is reached, the security audit log file is renamed for archival purposes and a new security audit log file is created using the default name.

The default value is 2 MB.

8. Optionally, change the maximum number of the audit log file versions that are maintained in the **auditlog maximum file version** attribute.

When the maximum number of versions is reached, the oldest archived version of the security audit log file is removed before the new log file is created.

The default value 0, means that all versions are maintained.

9. Optionally, change the audit message severity to increase or decrease the volume of messages saved in the security audit log in the **auditlog severity** attribute.

Changes to the attribute apply to each client that generates security related events. For example, if the security audit log severity attribute is Information, all clients will send messages with the Information severity level.

The following severity levels are available:

- Information
- Notice
- Warning
- Error — selected by default
- Severe
- Critical

The Information and Notice level audit messages are very common. If the security audit log records too much or too little detail, adjust the severity level accordingly.

10. Optionally, use a third party logging service to send security audit log messages to by using the **auditlog rendered service** attribute. [Table 148 on page 940](#) provides a description of the available options.

Table 148 Available auditlog rendered locale options (1 of 2)

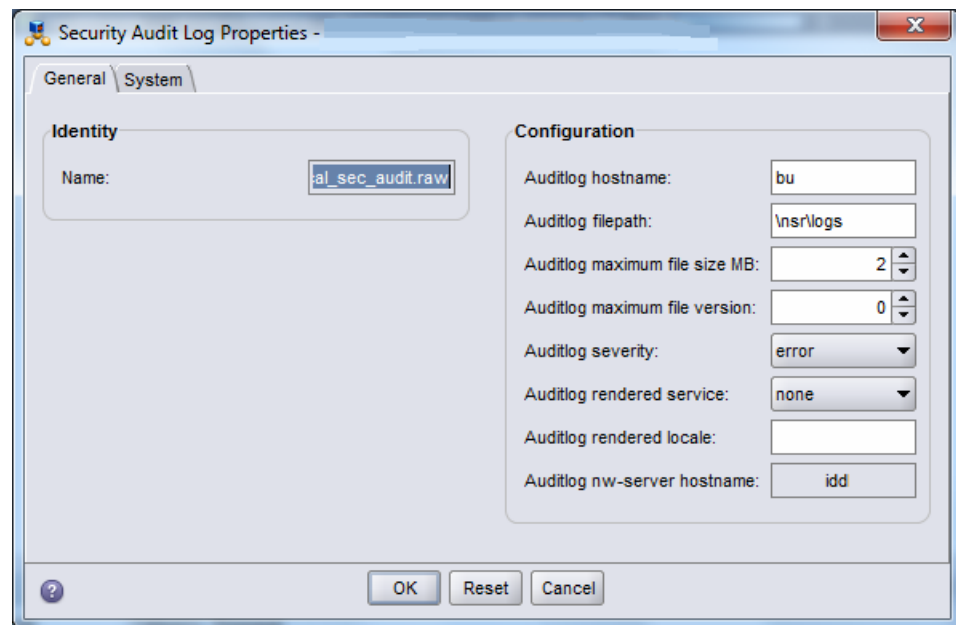
Option	Description
none	<ul style="list-style-type: none"> • The default value. • Writes unrendered security audit log messages to the <i>NetWorker_server_sec_audit.raw</i> file only. • Use the nsr_render_log program to render the log file into a readable format. “Audit message format” on page 942 provides more information.

Table 148 Available auditlog rendered locale options (2 of 2)

Option	Description
Local	<ul style="list-style-type: none"> Writes rendered security audit log messages to the <i>NetWorker_server_sec_audit.log</i> file. Writes unrendered security audit log messages to the <i>NetWorker_server_sec_audit.raw</i> file.
syslog	<ul style="list-style-type: none"> Writes rendered security audit log messages to the UNIX syslog. Writes unrendered security audit log messages to the <i>NetWorker_server_sec_audit.raw</i> file.
eventlog	<ul style="list-style-type: none"> Write rendered security audit log messages to the Windows Event Log. Writes unrendered security audit log messages to the <i>NetWorker_server_sec_audit.raw</i> file.

11. Optionally, specify the locale for the rendered audit log file in the **auditlog rendered locale** attribute. If this attribute is empty, the default locale **en_US** is used. The *Multi-locale datazone considerations* section in the *NetWorker Installation Guide* describes how to install and configure the NetWorker software on a machine that uses a non-English locale.

[Figure 71 on page 941](#) provides an example of the Security Audit Log Properties resource.

**Figure 71** The Security Audit Log Properties resource

12. Click **Ok**.

13. Review **Monitoring > Log** window to ensure that the configuration change is successful.

For example:

- If the host specified in the **auditlog hostname** attribute supports security audit logging and the nsrlogd daemon is successfully started, a message similar to the following appears:

```
The process nsrlogd was successfully configured on host
'security_audit_log_hostname' for server 'NetWorker_server'.
```

- If the host specified in the **auditlog hostname** attribute does not support security audit logging or the nsrlogd daemon does not start successfully, a message similar to the following appears:

```
The security audit log daemon nsrlogd is probably not running.
'Unable to connect to the nsrexecd process on host 'client_name'.
'355:Program not registered.'. Ensure that the host 'client_name'
can be reached. If required, restart the host.
```

- If a service port is not available on the host specified in the **auditlog hostname** field, the nsrlogd daemon fails to start and a message similar to the following appears:

```
Process nsrlogd was spawned on 'security_audit_log_hostname', but
nsrlogd could not open an RPC channel. 'Unable to connect to the
nsrlogd process on host 'security_audit_log_hostname'.
'352:Remote system error'.
```

- If the path specified in the **auditlog filepath** does not exist, a message similar to the following appears:

```
Unable to open the output file
'/proc/NetWorker_server_sec_audit.raw'
for the security audit log. No such file or directory.
```

NOTICE

Users that belong to the Security Administrators User Group but not the Application Administrators User Group cannot see messages in the Logs window.

Audit message format

The security audit log file contains the following information about the each security audit message:

- ◆ TimeStamp
- ◆ Category
- ◆ ProgramName
- ◆ RenderedMessage

Use the **nsr_render_log** program on a UNIX server or the **nsr_render_log.exe** program on a Windows server to render the audit log file into a readable format.

For example, the following message is displayed when the security audit log file is rendering by using the command `nsr_render_log Security_Audit_Log_filename:`

```
03/03/12 14:28:39 0 nsrd Failed to modify Resource type: 'NSR
  usergroup', Resource name: 'Users' for Attribute: 'users' by user:
  'administrator' on host: 'nwserver.emc.com'
```

- ◆ The TimeStamp is: 03/03/12 14:28:39.
- ◆ The Category is 0.
- ◆ The ProgramName is nsrd.
- ◆ The RenderedMessage is: Failed to modify Resource type: 'NSR usergroup', Resource name: 'Users' for Attribute: 'users' by user: 'administrator' on host: 'nwserver.emc.com'.

Communication security settings

Communication security settings enable the establishment of secure communication channels between:

- ◆ Product components
- ◆ Product components and external systems or components

Port usage

[Table 149 on page 943](#) lists all components, protocols, ports, and services.

Table 149 Port usage (1 of 2)

Component	Service	Protocol	Port	Description
Service Port Range (SPR)		TCP	7937 - 9936	This is the default range of ports that all NetWorker daemons should use when they need to start a service. This range can be configured. The daemons might start in any order, therefore, there is no guarantee that any one daemon will always use any singular port from this range.
nsrd	RPC	TCP		1 from SPR
nsrindexd	RPC	TCP		1 from SPR
nsrmmdbd	RPC	TCP		1 from SPR
nsrmmgd	RPC	TCP		1 from SPR
nsrjobd	RPC	TCP		1 from SPR
nsrlogd	RPC	TCP		1 from SPR
nsrexecd	RPC	TCP	7937, 7938	Plus 2 from SPR. Regardless of SPR, nsrexecd always listens to these two ports. 7938 must be allowed through a firewall, either by NetWorker or another portmapping service, or NetWorker will not work.
nsrlcpd	RPC	TCP		Per instance running.
nsrmmmd	RPC	TCP		1 from SPR, per instance running.
nsrsnmd	RPC	TCP		1 from SPR

Table 149 Port usage (2 of 2)

Component	Service	Protocol	Port	Description
NMC Web Server	HTTP	TCP	9000	Jumpstart to launch NMC
NMC	GSTD	TCP	9001	Communicates between Java client and main daemon.
NMC SQLAnywhere DB	DB		2638	Database listening port

Encrypting backup data

Backup and archive data on UNIX and Windows hosts can be encrypted with the AES Application Specific Module (ASM). The AES ASM provides 256-bit data encryption. Backup data is encrypted based on a user-defined pass phrase.

Do not use AES encryption when backing up files that are encrypted using the Microsoft Windows Encrypting File System (EFS).

Encryption for cloud backup data

A cloud backup device can also be set up to encrypt data sent to the cloud. If encryption is already enabled for the NetWorker host and you enable encryption on the cloud backup device, backups will be slower because encryption functions will occur twice. [“Cloud devices” on page 185](#) provides more information about cloud backups.

Federal Information Processing Standard Compliance

NetWorker utilizes encryption technologies from RSA BSAFE that are compliant with the Federal Information Processing Standard (FIPS 140-2). RSA BSAFE is deemed compliant under certificate 1092.

NetWorker 8.0 SP1 is the minimum version of the NetWorker software that contains the RSA BSAFE FIPS compliant encryption technologies. A NetWorker server 8.0 SP1 or later requires NetWorker client 7.6 SP4 or later.

[Table 150 on page 945](#) lists the NetWorker 8.1 SP1 and higher supported platforms that contain RSA BSAFE FIPS compliant encryption technologies.

The *EMC NetWorker Software Compatibility Guide* provides detailed information about the latest NetWorker support compatibility matrix.

Table 150 NetWorker 8.1 SP1 supported platforms that contain RSA BSAFE FIPS compliant encryption technologies (1 of 3)

Supported Platform	Supported Server and Storage Node OS	Supported Versions or Service Packs	Supported Client OS	Supported Versions or Service Packs
Windows x86	Windows Server 2008 (all editions) Storage Node only	SP1, SP2	Windows Server 2008 core	SP1, SP2
	Windows Server 2008 R2 (all editions) Storage Node only	SP1	Windows Server 2008 (all editions)	SP2
	Windows Server 2008 without Hyper-V [Standard, Enterprise and Datacenter Edition] Storage Node only	SP1, SP2	Windows Server 2008 R2 (all editions)	SP1
			Windows Server 2008 without Hyper-V [Standard, Enterprise and Dacenter Edition]	
			Windows Server 2003 (all editions)	SP1, SP2
			Windows Server 2003 R2 (all editions)	SP1, SP2
			Windows 7	SP1
			Windows VISTA (Business, Ultimate Edition)	SP1, SP2
			Windows XP (all editions)	SP2, SP3
			Windows Storage Server (WSS) 2003 R2	SP1, SP2
		Windows Unified Data Storage Server (WUDSS) 2003	SP1, SP2	
		Virtual Server 2005 R2 <ul style="list-style-type: none"> • Windows Server 2003 Standard Edition • Windows Server 2003 Enterprise Edition • Windows Server 2003 Web Edition • Windows Server 2003 R2 Standard Edition • Windows Server 2003 R2 Enterprise Edition 	SP1	

Table 150 NetWorker 8.1 SP1 supported platforms that contain RSA BSAFE FIPS compliant encryption technologies (2 of 3)

Supported Platform	Supported Server and Storage Node OS	Supported Versions or Service Packs	Supported Client OS	Supported Versions or Service Packs
Windows x64	Windows Server 2008 (all editions)	SP1, SP2	Windows Server 2008 core	SP1, SP2
	Windows Server 2008 R2 (all editions)	SP1	Windows Server 2008 (all editions)	SP1, SP2
			Windows Server 2008 R2 (all editions)	SP1
			Windows Server 2003 (all editions)	SP1, SP2
			Windows Server 2003 R2 (all editions)	SP1, SP2
			Windows 7	SP1
			Windows VISTA (Business, Ultimate edition) Windows XP (all editions)	SP1, SP2 SP2, SP3
			Windows Storage Server (WSS) 2003 R2	SP1, SP2
			Windows Unified Data Storage Server (WUDSS) 2003	SP1, SP2
			Virtual Server 2005 R2 Windows Server 2003 Standard Edition Windows Server 2003 Enterprise Edition Windows Server 2003 Web Edition Windows Server 2003 R2 Standard Edition Windows Server 2003 R2 Enterprise Edition	SP1
Linux x86	Red Hat Enterprise Linux AS, ES, WS	5, 6	Red Hat Enterprise Linux AS, ES, WS	5, 6
	SuSE Linux Enterprise Server (SLES)	10, 11	SuSE Linux Enterprise Server (SLES)	10, 11
	Oracle Linux	5	Oracle Linux	5
	Novell Open Enterprise Server (OES)	OES, OES SP2, OES 2, OES SP3	Novell Open Enterprise Server (OES)	OES, OES SP2, OES 2, OES SP3
	Redflag Asianux Server	3	Redflag Asianux Server	3
	CentOS Linux	5	CentOS Linux	5

Table 150 NetWorker 8.1 SP1 supported platforms that contain RSA BSAFE FIPS compliant encryption technologies (3 of 3)

Supported Platform	Supported Server and Storage Node OS	Supported Versions or Service Packs	Supported Client OS	Supported Versions or Service Packs
Linux x64	Red Hat Enterprise Linux AS, ES, WS	5, 6	Red Hat Enterprise Linux AS, ES, WS	5, 6
	SuSE Linux Enterprise Server (SLES)	10, 11	SuSE Linux Enterprise Server (SLES)	10, 11
	Oracle Linux	OES, OES SP2, OES 2, OES SP3	Oracle Linux	OES, OES SP2, OES 2, OES SP3
Linux Itanium			Red Hat Enterprise Linux AS, ES, WS	5
			SuSE Linux Enterprise Server (SLES)	10, 11
Oracle Sparc (64-bit)	Oracle Solaris	10	Oracle Solaris	10
	Oracle Solaris Non-global zones	10	Oracle Solaris Non-global zones	10
Oracle x64 (AMD64 and Intel EM64T)	Oracle Solaris	10	Oracle Solaris	10
HP Itanium	HP-UX	11i v2 (Storage node only) 11i v3 (Server only)	HP-UX	11i v2 11i v3
IBM Power AIX (32-bit) IBM Power AIX (64-bit)	IBM AIX	6.1 7.1	IBM AIX	6.1 7.1

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.
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	<ol style="list-style-type: none">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 also "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.
Atmos	EMC cloud storage product.
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 also "license key."
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.

autochanger [See "library."](#)

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. [See "retention policy."](#)

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.

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.

client-initiated backup [See "manual backup."](#)

clone

1. Duplicate copy of backed-up data that is indexed and tracked by the backup 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.
cloud	Configuration of backup disks that uses EMC Atmos.
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 also "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.
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. <i>See also</i> "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.
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.
datazone	Group of clients, storage devices, and storage nodes that are administered by a NetWorker server.
deduplication backup	Type of backup that removes redundant blocks of data to decrease storage space usage. When the deduplication 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	<ol style="list-style-type: none"> 1. Storage unit or folder that can contain a backup volume. A device can be a tape, 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	<ol style="list-style-type: none"> 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. <p>See also "distributed File System (DFS)."</p>
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	Instructions to take special actions on a given set of files for a specified client during a backup.
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.
domain controller	Server that stores directory data and manages user access to a network.
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 also "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 also "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 also "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 also "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 the application services continue despite a hardware or software failure.
- 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 also "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)** Capability of the software to display and output data in the same language fonts and numeric formats that are passed to it by localized operating systems or applications.

J

- 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 archive (JAR)** File that contains compressed components needed for a Java applet or application.
- 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: <ul style="list-style-type: none"> ◆ 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 automates the loading and mounting of movable storage media during backup and recovery processes. 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	Combined enabler code and authorization code for a specific product release to permanently enable its use. Also called an activation key or license enabler.
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 or AlphaStor. 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

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.
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.
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 also "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 tape or disk volume 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	Method of backing up or restoring multiple parallel streams of data simultaneously between a database and multiple storage devices. Also known as multistripe.
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: <ul style="list-style-type: none"> ◆ 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 either 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.
network file system (NFS)	Communications protocol that enables users to access shared files on different types of computers over a network.
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	EMC 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.
NFS server	Host that contains exported file systems that NFS clients can access. See also "network file system (NFS)."
node	See "cluster."
noncritical volume	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	Method that backs up or recovers data for multiple clients, or multiple save sets for one client, at the same time.
pathname	Set of instructions to the operating system for accessing a file: <ul style="list-style-type: none"> ◆ 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 also "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	<ol style="list-style-type: none"> 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. See also "conventional storage."

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 also "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 filesystem 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	<ol style="list-style-type: none"> 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.
requestor	A VSS-aware application that creates and destroys a shadow copy. NetWorker software is a requestor. See also "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. See also "recover."
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. See also "browse policy."
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 also "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 live backup or nonpersistent backup.
root	<ol style="list-style-type: none"> 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 also "hash."
save	NetWorker command that backs up client files to backup media volumes and makes data entries in the online index.
save set	<ol style="list-style-type: none"> 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	Attribute that indicates whether a save set was successfully backed up and whether it is currently browsable, recoverable, or recyclable.

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	See "conventional storage."
security event	Operation related to authorization, authentication, or configuration.
server index	See "client file index."
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 also "Volume Shadow Copy Service (VSS)."
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.
snapup	Term previously used by Avamar software to refer to backup.
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 also “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	<ol style="list-style-type: none"> 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 also “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 also "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 Identifiable unit of physical storage medium, such as magnetic tape or disk file system used to store data.

volume ID (volid) 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. [See also "label."](#)

Volume Shadow Copy Service (VSS) 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 also "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 also "Volume Shadow Copy Service \(VSS\)."](#)

INDEX

A

- aborting a recover 815
- ACPI
 - NetWorker support 892
 - OnNow 891
 - recover considerations 892
 - scheduled backup considerations 892
- actions, notifications 486
- Active Directory
 - backup 886
 - domain controller 886
 - explained 886
 - recover 886
 - prerequisites 395
 - SYSTEM STATE save set 846, 886
- adding
 - annotation 459
 - enabler code 538
 - folder 546
 - host 544
 - license 538
 - multiple hosts 548
- addresses
 - server, changing 832
- Administration window 37
- administration window
 - opening 39
- administrator
 - privileges 550
- administrators
 - group 45
- Advanced Configuration and Power Interface. See ACPI
- Advanced File Type Device
 - concurrent access 183
 - simultaneous recovery of save sets 183
- aes encrypted data
 - implementing 108
 - recovering 389
- AFTD load balancing 181
- alert priority 458, 469, 473
- alerts
 - lists priority, category, time, and message 464
- Alias attribute 830
- aliases, host sharing restriction 813
- annotation
 - adding 459
 - attribute 458
 - icon 459
 - viewing 459
- application data
 - recovering on NetWorker client
 - UNIX 909
- application specific module (ASM) 299
- archive
 - changing time of 333
 - request, starting automatically later 475
 - stopping request in progress 474
- Archive attribute 272
- Archive button 45
- archiving
 - archive pools
 - errors 834
 - archive requests
 - creating 332
 - defined 331
 - deleting 334
 - disabling 337
 - editing 333
 - archive services
 - enabling 330
 - archive volume pool 326
 - compared to backups 326
 - copying an archive request 333
 - disable scheduled archive 475
 - enabling 328
 - grooming files 326
 - indexed 327
 - licensing 326
 - managing 337
 - manual 330
 - naming archive requests, errors 834
 - NetWorker User program
 - toolbar function 45
 - nonindexed 328
 - nsrchive program considerations 834
 - permissions 328, 329
 - pools
 - configuration 310
 - creating 318
 - types of pools 329
 - remote requests, failure 834
 - requesting an archive 474
 - requesting an archive's status 464
 - requirements 326
 - retrieving from client machine 334
 - save sets 326
 - multiple 834
 - retrieving 334, 335
 - scheduled 331
 - scheduling 337
 - starting 337
 - stopping 337
 - time change 333
 - tracking entries 327
 - troubleshooting 834
 - viewing details 337
 - volumes, cloning 353
- ASM (application specific module) 299
- ASR save set 710

- ASR. See Automated System Recovery
 - attribute
 - Administrator 566
 - annotation 458
 - auth code 538
 - category 458
 - enabler code 538
 - Grooming
 - limitation 332
 - message 458
 - note 458
 - priority of managed event 457
 - See also specific attribute name
 - server name 457
 - time 458
 - auth code attribute 538
 - authentication
 - considerations for 613
 - nsrauth 612
 - oldauth 612
 - authorization code
 - error message 836
 - Auto Media Management
 - recyclable volumes 279
 - auto media verification
 - media position errors 810
 - Autochanger parallelism 556
 - autochangers
 - AIX considerations 904
 - attributes
 - new 820
 - autodetection 821
 - control ports access 823
 - destination component 821
 - HP-UX considerations 904
 - installation, HP-UX considerations 900
 - maintenance commands 820
 - STK-9840 904
 - Automated System Recovery
 - 2008, 2008 R2, Windows 7 719
 - ASR save set 710
 - cluster database, special handling 717
 - COM+ database, special handling 717
 - disk quota database, special handling 717
 - documentation 710
 - FAT16 partitions 711
 - OEM recovery CD 711
 - overview 710
 - recovering a client 715
 - scheduled backup of ASR save set 712, 713, 749
 - verifying client recovery 718
 - WMI database, special handling 717
 - Autorestart attribute 251
 - Autostart attribute 251
- B**
- back ups, manual 111
 - backing up Console 110
 - backing up renamed directories 65, 129
 - Backup button 45
 - backup configuration wizard 59
 - backup groups. See groups
 - backup levels
 - 1-9 268
 - described 267
 - options 268
 - overriding 267
 - planning 268, 270
 - types
 - consolidated 263, 264, 268, 271
 - full 268, 270, 271, 813
 - incremental 268
 - level 271
 - skip 268
 - usage 270, 271
 - backup schedules. See schedules
 - backups
 - adhoc 50
 - Backup and Recover Server service 585
 - balancing resources 249, 263
 - bootstrap 259
 - client-initiated 309
 - commands 119
 - example 120
 - savepnpc program 124
 - completed 258
 - consolidated 263, 264
 - customization scripts 119
 - cycle, using levels 262
 - directives 290
 - failed 117
 - filesystems 265
 - force incremental 255
 - groups. See groups
 - hard links 116
 - incremental
 - pool for 309
 - large client filesystems 622
 - large filesystems 265
 - levels 270
 - log file 117
 - managing 257
 - manual 58, 71
 - policies 286
 - pool for 309
 - NetWorker User program
 - browse windows 45
 - nonscheduled 50, 58, 71
 - operations
 - stopping 810
 - operators group 45, 585
 - permissions, backup operators group 585
 - pools 304
 - previewing 257
 - recoveries 807
 - RPC errors 815
 - save sets. See save sets
 - server, forced 817
 - time intervals, setting 255
 - troubleshooting 807

- types of backups. See backup levels
 - Windows NT registry 73
 - bar chart 432
 - basic reports 424
 - best practices for cloud backups 186
 - BMR support
 - recovery process 389
 - bootstrap
 - emailing 259
 - pools 306
 - printing 259
 - failure 817
 - boot-time file 833
 - browse policies
 - about 275, 276
 - client file index growth 586
 - clones, storage nodes 355
 - defined 276
 - save set recoveries 399, 403
 - usage 276
 - browse windows
 - described 45
 - toolbar 45
 - browser
 - unresponsive 836
- C**
- capture of managed event 457
 - case sensitivity 452, 534
 - category attribute of managed event 458
 - centralized
 - license management 36
 - Certificate Server
 - recovery prerequisites 395
 - SYSTEM STATE save set 846
 - certified protocols 833
 - changing
 - name of folder 547
 - changing servers 46
 - characters
 - illegal 546, 547
 - characters, not permitted 505, 546, 547, 572
 - chart formats 429
 - checkpoint restart backups 95
 - CHKDSK, running 71
 - client 606
 - client alias, changing 813
 - client backup configuration wizard 59
 - Client Direct 34, 64, 95, 162, 165, 176
 - client file index
 - backup level, pool for 309
 - browse policies 276, 591
 - checking 814
 - clones and storage nodes 355
 - cross-checking 589
 - defined 585
 - entries
 - adding 586
 - removing 587, 591, 593
 - growth 586
 - location, designating 589
 - managing size 591
 - operations
 - checking 587
 - moving 590
 - recovery 372
 - policies 591
 - pools 306
 - save sets
 - cycles, removing 592
 - entries 279
 - removing 591
 - size 586
 - management 591
 - notification 815
 - Client parallelism 554
 - Client Retries attribute 251
 - client-initiated backups
 - pool for 309
 - clients
 - aliases, problems 813
 - backup commands 119
 - client ID
 - creating new client 404
 - client/server communication errors 832
 - cloning
 - examples 348
 - configuration 606
 - defined 606
 - DHCP 585
 - DNS name resolution 585
 - editing 606
 - groups 248
 - large filesystems 265
 - manual backups 58
 - multiple 622
 - NetWorker User program 44
 - operations
 - archive retrieve 334
 - backing up renamed clients 814
 - backups 622
 - creating 606
 - editing 606, 607, 614, 615, 618
 - indexes, moving 590
 - installation 606
 - manual backups 58
 - recovery, failure 814
 - policies, multiple 282
 - priority 625
 - save sets 622, 623
 - Solaris binary location 898
 - clone pools, configuring 310
 - Clone reports 423, 439
 - Clone Volumes dialog box 350
 - cloning
 - archives 353
 - defined 340
 - destination volume, defined 340
 - examples 348
 - manual 346

- online indexes, storage nodes 355
 - recovery
 - save sets 351
 - volumes 351
 - save sets 341, 346
 - manually 348
 - performing 346
 - source volume, defined 340
 - storage nodes
 - online indexes 355
 - volumes 350
 - creating 350
 - details, viewing 350
 - process 350
 - cloning to cloud 189
 - cloud
 - best practices 186
 - cloning 189
 - compared to other device types 185
 - data consumption information 189
 - prerequisites 185
 - reports 448
 - setting up a cloud device 186
 - staging 189
 - support for 185
 - Cluster Server
 - recovery prerequisites 395
 - SYSTEM STATE save set 846
 - COM+ Database
 - recovery prerequisites 395
 - SYSTEM STATE save set 846
 - command
 - export 534
 - gst 54
 - savepsm 110
 - command line
 - reporting 425, 454, 455
 - reporting program 454
 - Common Device Interface
 - SCSI command 199
 - Completion Data Retention 421
 - completion data retention 421
 - completion message retention 421
 - compression, data 109
 - computer damage, recovery from 528
 - configuring
 - Host reports 448
 - reports 426
 - connection
 - problem 836
 - refused 836
 - Console
 - environment variables 531
 - Console client
 - starting, after the first time 39
 - Console Configuration Wizard 535
 - console security administrator
 - resetting administrator password 528
 - Console software
 - HTTP service port 38
 - improving performance 531
 - License Manager 539
 - logging on 38
 - managing window 35
 - URL 38
 - Console window
 - opening 38
 - console window 36
 - consolidated backups 271
 - contacting server, problem with 836
 - Content Index Server
 - backing up 114
 - backup 114
 - defined 114
 - recover 114, 398
 - SYSTEM DB save set 114, 398, 848
 - control zone 542
 - copying host 545
 - corrupted database 836
 - creating
 - folder 546
 - host 544
 - label templates 322
 - pools 311
 - staging policies 363
 - criteria for organizing hosts 542
 - critical priority 458, 469, 473
 - cross-checking online indexes 589
 - cross-platform
 - name resolution 585
 - CSV 453
 - cutting and pasting host 545
- ## D
- daemon log file 600
 - daemon.log 601
 - daemons
 - nsrexecd 625
 - damaged computer 528
 - data
 - compression 109
 - encryption with aes ASM 108
 - life cycle
 - cloned data 349
 - managing 281
 - relocating, errors 815
 - sorting
 - into pools 305, 307, 308
 - to storage devices 310
 - verify 76
 - data compression 822
 - Data Domain deduplication 33
 - Data Retention dialog box 422
 - database
 - backing up 110
 - corrupted 836
 - corruption 836
 - delete failed 837
 - failed to store 837
 - fetch operation 836

- database, fetch operation failure 836
 - db-output 838
 - dbsrvr9 840
 - dbstop output 838
 - DD Boost 34, 162, 165, 345
 - debug level 534
 - debug messages
 - logging stopped 838
 - deduplication 32, 34, 162
 - Data Domain 33
 - deduplication backups 107
 - default backup schedules 261
 - default pool 304, 308, 309
 - delete operation failed 837
 - deleting
 - folder 546
 - host 545
 - label templates 323
 - license 538
 - managed event note 459
 - multiple hosts 548
 - note 459
 - problem 837
 - staging policies 365
 - device
 - covert to read-only 183
 - delete 184
 - disable 184
 - erase data 184
 - lists of 464
 - related messages 464
 - Device Access Information 171, 174, 178
 - Device configuration wizard 168
 - devices
 - device drivers
 - maintenance commands 820
 - device ordering 199
 - correcting device order problems 201
 - detecting device order problems 201
 - disk label errors 816
 - filesystem
 - staging 362
 - labeling errors 816
 - load balancing 181
 - nonrewinding 825
 - pools 310
 - DFS (Distributed File System)
 - junctions 878
 - not in SYSTEM save sets 848
 - recover 881
 - DHCP (Dynamic Host Configuration Protocol)
 - database 115, 397
 - static IP address for NetWorker server 832
 - DHCP (dynamic host configuration protocol)
 - clients 585
 - dialog box
 - Data Retention 422
 - differences 164
 - directed recovery
 - access 369
 - advantages 368
 - defined 369
 - use of 368
 - directives
 - ASM (application specific modules) 299
 - copying 291
 - creating 290
 - defined 290
 - deleting 291
 - editing 291
 - naming restrictions 812
 - preconfigured 294
 - disabling
 - managed event capture 457
 - disallowed characters 505, 546, 547, 572
 - disappearing managed event 460, 838
 - disaster recovery 405
 - Windows 2003, XP 709
 - Windows 2008, 2008 R2, Windows 7 719
 - disk quota database
 - recovery prerequisites 396
 - SYSTEM DB save set 848
 - disk space
 - insufficient 836
 - disk space, gstd log file 839
 - disk space, gstd.log 534
 - display problem 837
 - displaying
 - annotation 459
 - reports 428
 - distributed segment processing (DSP) 162
 - DNS
 - hostname alias, troubleshooting 813
 - DNS (Domain Name System)
 - host name determination 585
 - document view 434
 - domain controller
 - Active Directory, configured by 886
 - defined 886
 - encryption keys
 - not supported 887
 - drag-and-drop, column 39
 - drill-down reports, Managed Event 446
 - DSA
 - DSA and NDMP Tape Server distinctions 645
 - duplicating host 545
 - dynamic addressing 585
- ## E
- ECB (Event Control Block) 811
 - editing
 - clients 606
 - folder 547
 - label templates 323
 - pools 317
 - staging policies 364
 - emailing the bootstrap report 259
 - EMC online support website 25
 - emergency priority 458, 469, 473
 - enabler code 538

- enabler code, entering 538
 - enabler code, problem 837
 - enabling
 - debug information 534
 - JavaScript 840
 - software 538
 - encrypting data
 - aes 108
 - Encrypting File System
 - backup 886
 - directed recover 886
 - encryption keys 887
 - explained 886
 - not in SYSTEM save sets 848
 - recover 886
 - entering
 - enabler code 538
 - license 538
 - enterprise
 - button 36
 - enterprise hierarchy
 - adding folder 546
 - adding host 544
 - copying folder 547
 - copying host 545
 - deleting folder 546
 - deleting host 545
 - folder 542
 - host 542
 - managing host 544
 - managing multiple hosts 548
 - moving folder 547
 - moving host 545
 - renaming folder 547
 - viewing 543
 - Enterprise Summary report 447
 - environment variable 534
 - GST_DEBUG 839
 - GST_MAXLOGSIZE 839
 - GST_MAXLOGVERS 839
 - setting 534
 - environment variables
 - NSR_DEV_BLOCK_SIZE_MEDIA_TYPE 194, 197
 - NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE 198
 - NSR_DEV_LOAD_TIME_MEDIA_TYPE 198
 - NSR_DEV_LOAD_TRY_TIMEOUT_MEDIA_TYPE 198
 - NSR_DEV_TAPE_FILE_SIZE_MEDIA_TYPE 197
 - NSR_DEV-DEFAULT_CAPACITY_MEDIA_TYPE 198
 - error
 - database store 837
 - relocating data 815
 - RPC (remote procedure call) 815
 - error messages 836
 - copy violation 817
 - destination component 821
 - disk label 816
 - environmental variables 835
 - illegal record size 811
 - media verification 810
 - print server 488
 - RPC errors 832
 - save sets 813
 - server, unavailable 832
 - Event Control Block (ECB) 811
 - event logs
 - backup 887
 - explained 887
 - not in SYSTEM save sets 848
 - recover 887
 - Event Time 446
 - Event Viewer 490
 - example
 - organizing hosts 542
 - sorting managed event 39
 - export command 534
 - export formats 453
 - exporting
 - non-ASCII characters 454
 - reports 453
- ## F
- failed
 - operation 837
 - failed operation 836
 - FAT16 partitions, ASR limitation 711
 - fetch operation 836
 - Figure 204
 - file conversion, sparse to fully-allocated 817
 - file handling
 - indicators 46
 - file index missing, message 814
 - File Manager 813
 - File Replication Service
 - backup 888
 - SYSTEM STATE save set 846
 - filemarks 822
 - filename, support for short names 116
 - files
 - compressing 109
 - encrypting 108
 - HOSTS 585
 - log 838
 - open files, backing up 259
 - servers 830
 - verify 76
 - filesystem devices 362
 - filesystems
 - backups, large 622
 - firmware, verification 826
 - folder 542
 - adding 546
 - deleting 546
 - editing 547
 - Force Incremental attribute 266
 - force incremental attribute 255
 - formats
 - export 453
 - full backups 270

G

- grooming 332
 - limitation 332
- grooming files 326
- group
 - backup, stopping 466
 - restarting backup 466
 - start immediately 466
 - viewing control details 466, 471
 - viewing status 464
- Group resource
- groups 249
 - attributes 251
 - backup operators 585
 - backups
 - management 257
 - previewing 257
 - bootstrap 259
 - client policies, multiple 282
 - completed 258
 - containing bootstrap 811
 - defined 248
 - filesystems, large 265
 - force incremental, setting 255
 - listing of per server 464
 - naming restrictions 812
 - NetWorker User program privileges 45
 - operations
 - copying 254
 - deleting 254
 - editing 253
 - previewing 257
 - time intervals, setting 255
 - types 250
- gst command 54
- GST server process 840
- GST_DEBUG 839
- GST_MAXLOGSIZE 839
- GST_MAXLOGVERS 839
- gstd 840
 - log 534, 839
 - process 54
 - size of log 534, 839
- gstd log file 839
- gstd service 53
- gstmodconf command 548, 549

H

- hard links, backup and recovery 116
- hardware
 - upgrade 528
- host
 - adding 544
 - copying 545
 - deleting 545
 - managing 544
 - moving 545
 - organization of 542
 - transfer affidavit 530

- Host List 447
- host name determination
 - DHCP clients 585
 - TCP/IP 585
- Host Reports 423
 - types and their configuration 448
- hostname alias 813
- hostname file 548
- HOSTS file 585
- How 426
- HP-UX
 - creating device files 901
 - installing autochangers 900
 - pass-through drivers 901
- HTML 453
- hung browser 836

I

- icon
 - annotation 459
 - priority 458
- IIS (Internet Information Server)
 - defined 888
 - recovery prerequisites 395
 - SYSTEM STATE save set 846
- Immediate cloning 345
- Importing save set information to a different server 404
- Inactivity Timeout attribute 251
- incremental backups
 - pool for 309
- Index Save Sets dialog box 587
- indexed archives 327
- individual user authentication 532
- info priority 458, 469, 473
- information, sorting table 39
- inquire program 900
- install log 839
- Interactive view 428
- interface
 - overview 35
- Internet Information Server See IIS
- Interval attribute 251
- Invalid Object id, error message 837
- ioscan program 900

J

- Java Runtime Environment 455
- JavaScript 840
- jbconfig program
 - autochangers
 - HP-UX 904
 - hanging 821
- jbexercise program
 - NDMP, not supported 654
- jobsd service 53
- JRE 38, 454, 455
- Jukebox parallelism 556

- L**
- label templates
 - attributes 321
 - components 321
 - creating 322
 - deleting 323
 - editing 323
 - naming restrictions 812
 - naming strategies 321
 - number sequences 321
 - preconfigured 319
 - labeling
 - tips 322
 - labels
 - silos, volumes 318
 - launch button 38
 - level of debug 534
 - Library parallelism 556
 - license
 - adding 538
 - deleting 538
 - License allocation failed, error 837
 - License Manager 539
 - licensing
 - archiving 326
 - copy violation 817
 - License Manager 539
 - List Report, User 447
 - LMHOSTS 585
 - load balancing 181
 - localized environments 427
 - location of
 - gst command 54
 - lockbox for pass phrases 611
 - log file 838, 839
 - log files
 - backup/recovery attempts 117
 - cloning information 340
 - logging on
 - to Console software 38
 - logging, events 490
 - login
 - name 505, 572
 - password 505, 572
 - lsdev program 902
- M**
- Mac OS X 906
 - managed application 542
 - managed event
 - define 456
 - deleting note 459
 - disabling capture of 457
 - disappearing 460, 838
 - priority 458
 - sorting example 39
 - Managed Event configuration
 - parameters 446
 - Managed Event drill-down reports 446
 - Managed Event Reports 423
 - Managed Events button 37
 - managed node 542
 - adding 544
 - copying 545
 - deleting 545
 - moving 545
 - managing
 - host 544
 - license 539
 - manual backups 111
 - manual backups. See backups, manual 58
 - Max active devices 555
 - Max parallelism, media libraries 556
 - Max parallelism, pools 557
 - Max Sessions 170, 172, 175
 - Max sessions 557
 - max sessions 176
 - media database
 - cloned data 349
 - clones and storage nodes 355
 - compression 592, 594
 - cross-checking 589
 - entries, removing 592, 593
 - managing size 585, 591
 - restoration 403
 - retention policies 276
 - save sets, entries 279
 - Media Library parallelism 556
 - media pools. See pools
 - media position errors 810
 - memory requirements for AFTDs 167
 - message attribute of managed event 458
 - message logs
 - failed backup/recovery attempts 117
 - message retention 421
 - messages file
 - Event Viewer 490
 - Microsoft Automated System Recovery. See Automated System Recovery
 - Microsoft Windows
 - backup operators group 585
 - databases
 - not in SYSTEM save sets 848
 - SYSTEM DB save set 848
 - Event Viewer 490
 - Windows Management Instrumentation
 - SYSTEM DB save set 848
 - mminfo program
 - reports 287
 - moving host 545
 - multiple hosts
 - adding or deleting 548
 - Multiplexing 556
 - multiplexing
 - performance issues 33
- N**
- name
 - folder, editing of 547

- resolution cross-platform 585
- naming restrictions 812
- NDMP (Network Data Management Protocol)
 - configuration
 - options 645
 - nsr resource attributes 654
- Nested Mountpoints 396
- network
 - TCP/IP certified 833
- NetWorker 480
 - startup commands 808
 - User program. See NetWorker User program
- NetWorker Management Console See Console
- NetWorker User program
 - adhoc backups 50
 - backing up registry, Windows NT 73
 - browse windows 45
 - changing servers 46
 - compression 109
 - connecting to a server 46
 - encryption 109
 - manual backups 71
 - overview 44
 - password protection 109
 - privileges 45
 - server connection 46
 - starting 44
 - toolbar 45
- NMC See Console
- node
 - adding 544
 - copying 545
 - deleting 545
 - managed 542
 - moving 545
- nonindexed
 - archives 328
- note
 - attribute 458
 - deleting from managed event 459
- notifications 490
 - defined 479
 - deleting 491
 - operations
 - customizing 484
 - preconfigured 480
 - printing 487
 - priorities 489
 - programs 486
 - SNMP
 - configuring 694
 - creating 696
 - modifying 695
 - nsrtrap 695
- nsr_getdate program 286, 287
- nsr_shutdown program 811
- nsradmin program
 - editing the nsrla.res database 608
 - starting 47
- nsrarchive program 327

- nsrauth authentication 612
- nsrd service 52
- nsrexecd 55
 - nsrexecd daemon 625
 - nsrexecd service 52
- nsrindexd service 52
- nsrjb program
 - troubleshooting, HP-UX 903
- nsrla.res database 608
- nsrldr program 487
- nsrmm program 286, 591, 811
- nsrmmmd daemon 198
 - NDMP
 - unsupported options 654
- nsrmmdbd service 52
- nsrmmgd service 52
- nsrtrap 694, 695
 - command line options 695
 - verbose mode 695
- numeric order 39
- nwrecover program
 - browse policy 276

O

- object, reference 837
- OEM recovery CD, ASR limitation 711
- offline disaster recovery, Windows 719
- oldauth authentication 612
- online indexes
 - cross-checking 589
 - entries
 - checking 587, 814
 - removing 587, 592
 - information
 - refreshing 589
 - viewing 587
 - management
 - manual 585
 - size 591
 - moving 590
 - recovery
 - location 812
 - restoration 399
 - save sets, viewing 587
 - size considerations 586
 - volumes
 - removing 593
- OnNow
 - Advanced Configuration and Power Interface (ACPI) 891
 - recover considerations 892
 - scheduled backup considerations 892
- open files, backing up with VSS 260
- operation failed 836, 837
- optimizing Console 531
- organizational criteria 542
- organizational structure, labeling for 322
- out of memory 836

- P**
- page cannot be displayed 837
 - Parallel save streams 93
 - Parallelism 553
 - parallelism
 - performance 33
 - pathname restrictions 813
 - PDF report format 453
 - performance
 - features 33
 - Performance Counters
 - SYSTEM STATE save set 846
 - performance of Console software 531
 - permissions
 - Archive feature 328, 329
 - backup operators group 585
 - Persistent binding 200
 - Persistent naming 200
 - pie chart 433
 - pie report format 432
 - plot chart 432
 - policies 284
 - backups, manual 286
 - browse
 - about 275, 276
 - data life cycle 281
 - defined 276
 - modifying 286
 - usage 276
 - clients 282
 - data life cycle 281
 - multiple 282
 - naming restrictions 812
 - overriding 286
 - planning 264
 - retention 278
 - about 275, 276
 - data life cycle 281
 - defined 276
 - modifying 286
 - usage 276
 - volume relabeling 276
 - setting expiration 422
 - policy
 - deleting 284
 - pool
 - configuration 749
 - consolidated backup 307
 - copying resource 317
 - Pool parallelism 557
 - Pool Type attribute, with archive pools 329
 - pools
 - archive 310
 - errors 834
 - archives, creating 318
 - archiving 318
 - bootstrap 306
 - client file index 306
 - clones 310
 - configuration
 - archive 310
 - clone 310
 - creating 311
 - criteria 305, 307
 - data sorting 305, 308
 - default 304, 308, 309
 - clone pool 310
 - defined 304
 - devices 310
 - editing 317
 - expression matching 306
 - incremental backups 309
 - manual backups 309
 - precedence 307
 - restrictions 305
 - save set consolidation 88
 - sorting data 310
 - volume labels 318
 - POSIX hard links, problems recovering 116
 - PostScript report format 453
 - Power Monitor service 892
 - precedence for pools 307
 - preconfigured notifications 490
 - priorities, notifications 489
 - priority
 - of managed event 458
 - symbol 458, 469, 473
 - Priority attribute 457
 - privileges
 - administrator 550
 - probe based backups 94
 - problem
 - contacting server 836
 - process, stopping and restarting 54
 - Program not registered 837
 - program not registered 837
 - Properties dialog box
 - NetWorker User program 46
 - protocols, certified 833
 - provider 773
 - ps command 840
 - PSS 93
- R**
- rearranging enterprise hierarchy 545
 - rearranging information in table 39
 - Recover button 45
 - recover program
 - media database 279
 - retention policy 278
 - recoveries
 - archives 334
 - Backup and Recover Server services 585
 - backup operators group 585
 - clients, renamed 814
 - clone volumes 351
 - directed
 - access 369
 - advantages 368
 - defined 369

- usage 368
 - disaster-related 405
 - failed 117
 - files, finding for recovery 685
 - hard links created by POSIX 116
 - index-based
 - advantages 372
 - log file 117
 - NetWorker User program
 - browse windows 45
 - planning 264
 - save sets
 - client file index 399
 - media database 403
 - recovery
 - aborting 815
 - ASR 715
 - recycling. See volumes
 - reference object 837
 - refused connection 836
 - registering program 837
 - registry
 - backups 888
 - explained 888
 - SYSTEM STATE save set 846, 888
 - relocating data 815
 - remote access
 - recoveries 831
 - Remote Access list 831
 - remote archives, failure 834
 - Removable Storage Manager
 - recovery prerequisites 396
 - SYSTEM DB save set 848
 - removing
 - folder 546
 - host 545
 - license 538
 - renamed directories, backing up 65, 129
 - renaming folder 547
 - report
 - background processing 435
 - chart types 431
 - document mode 430
 - export formats 453
 - interactive mode 428
 - restricted views 434
 - User List 447
 - reports
 - basic and drill-down 424
 - command line reporting program 454
 - customized 425
 - daemon log file 600
 - daemon.log 601
 - date and time formats 426
 - Managed Event Drill-Down 446
 - save set policies 287
 - saved 451
 - viewing 428
 - Reports button 36, 37
 - requirements
 - NetWorker User groups 45
 - resetting administrator password 528
 - resolved events 460
 - resource
 - archive
 - changing archive time 333
 - request
 - copying 333
 - status 464
 - label template
 - copying 323
 - notification
 - copying 491
 - deleting 491
 - policy
 - deleting 284
 - pool
 - copying 317
 - directing data from consolidated backup 307
 - staging
 - copying a policy 364
 - user group
 - copying 568
 - creating 567
 - customizing privileges 563, 579
 - deleting 568
 - editing 564, 565
 - preconfigured 563
 - Restart Window attribute 251
 - restrictions on pathnames 813
 - retention policies 278
 - about 275, 276, 278
 - clones, storage nodes 355
 - defined 276
 - usage 276
 - volumes, relabeling 276
 - retention, completion data 421
 - retention, completion message 421
 - retention, save set 421
 - retrieval
 - annotations, empty 835
 - troubleshooting 834
 - retrieving archives 334
 - retrieving save sets 334, 335
 - RPC (remote procedure call)
 - errors 815
- ## S
- save program 119, 286, 287
 - save set
 - All save set 68
 - Save Set Details 438
 - Save Set Name 438
 - save set retention 421
 - save sets
 - archives, retrieving 334, 335
 - backup commands 119
 - backups 622
 - client combination 623
 - client file index, entries 279

- client priority 625
- cloning
 - automatic 345
 - described 341
 - manually 348
 - status 346
- consolidation
 - pools 88
 - usage 76, 78
- defining 623
- indexes, viewing 587
- information, viewing 587
- load balancing 622
- media database, entries 279
- multiplexing 556
- policies
 - modification 286
 - reports 287
- predefined 66
- recoveries
 - media database 403
 - online indexes 399
- staging. See staging
- status
 - clone 346
 - retention policy 278
- suspected 351
- SYSTEM
 - manual backups 72
 - point-in-time recovery
 - command prompt 394
 - recover
 - command prompt 392
- SYSTEM DB
 - backup levels 273
 - basic components 848
 - databases not included 848
 - recovery prerequisites 396
- SYSTEM FILES
 - backup levels 273
 - basic components 847
- SYSTEM STATE
 - backup levels 273
 - basic components 846
 - optional components 846
 - recovery prerequisites 395
- VSS ASR DISK
 - backup levels 274
- VSS OTHER
 - backup levels 274
- VSS SYSTEM
 - manual backups 72
 - point-in-time recovery 394
- VSS SYSTEM BOOT
 - backup levels 274
 - components 849
- VSS SYSTEM FILESET
 - backup levels 274
 - components 850
- VSS SYSTEM SERVICES
 - backup levels 274
 - components 850
 - VSS USER DATA
 - backup levels 274
- savegroup completion report 482
- Savegroup parallelism 555
- savegrp program
 - backup limitation 817
- savepnpc program
 - message logging by 127
 - using with customized backup program 124
- savepsm 110
- savestream 304
- scalability 35
- SCANDISK, running 61
- scanner program
 - record size 811
 - recovering clone volumes 351
 - recovering save sets from volumes 282
 - retention policy 278, 279
 - volumes, read-only 811
- Schedule attribute 251
- Schedule resource 248
- schedules 260
 - attributes 265
 - backup cycle 262
 - balancing 263
 - copying 267
 - default 261, 263
 - deleting 267
 - editing 83, 266
 - large filesystems 265
 - load balancing 622
 - naming restrictions 812
 - overriding 267
 - planning 264
 - preconfigured 261
 - staggered 263
 - usage 260
- Schedules window 265
- scheduling backups 110
- SCSI address selection for HP-UX 900
- SCSI ID 205
- security
 - application authentication 612
 - lockbox for pass phrases 611
 - overview of settings 929
- server
 - problem contacting 836
 - setting up 550
 - web address 38
- server name
 - attribute of managed event 457
- Server parallelism 555
- servers
 - address, changing 832
 - backup
 - operators Group 585
 - client/server communication errors 832
 - DCHP 585

- DNS name resolution 585
- dynamic addressing 585
- errors, binding to 832
- file 830
- index
 - backup, failure 817
 - management 586
- notifications
 - priorities 489
- service
 - gstd 53
 - jobsd 53
 - nsrd 52
 - nsrindexd 52
 - nsrmmdbd 52
 - nsrmmgd 52
- services
 - Backup and Recover Server 585
 - described 52
 - Power Monitor service 892
- session management 34
- sessions
 - lists of backup, recover, or browse sessions 464
- Setting 422
- setting
 - data retention policies 422
 - environment variable 534
 - expiration policies 422
- severity of managed event 457
- shadow copy 772
- short filename support 116
- Simple Network Management Protocol. See SNMP
- size of
 - gstd log 839
- sleep state
 - defined 891
- snapshot 772
- snapshot policy
 - creating 285
- SNMP (Simple Network Management Protocol) 694
 - configuring 694, 696
 - defined 694
 - notifications
 - configuring 694
 - creating 696
 - modifying 695
 - nsrtrap 694, 695
 - traps 694
- Software Administration button 36
- Solaris
 - troubleshooting 898
 - unsupported devices 898
- sorting managed event, example 39
- sorting table 39
- Source Client dialog box 385, 688
- sparse files 888
 - not in SYSTEM save sets 848
- sparse files, converting 817
- Special Handling dialog box 109
- stacked bar report format 432

- stacking bar chart 433
- staging
 - defined 362
 - filesystem devices 362
 - policies
 - creating 363
 - deleting 365
 - editing 364
 - to cloud 189
- Start Time attribute 251
- starting
 - server process 54
- status
 - viewing
 - group status 464
- stopping Console server 54
- storage nodes
 - affinity
 - problems 835
 - timeouts 835
 - troubleshooting 835
- Store Index Entries attribute, with archive pools 329
- storing database failed 837
- suspected save sets 351
- symptom of problem 836
- System File Protection
 - backup 847
 - defined 847
 - explained 847
 - recover 847
 - SYSTEM FILES save set 847
- system standby 891
- SYSVOL
 - SYSTEM STATE save set 846

T

- tables
 - display or hide columns in 40
 - multicolumn sorting 40
 - rearranging columns 39
 - sorting 39
- Target Sessions 170, 172, 175
- Target sessions 556
- target sessions 176
- TCP/IP
 - certification 833
 - changing NetWorker server address 832
 - DHCP clients 585
 - host name determination 585
 - troubleshooting hostname alias problems 813
- technical support, troubleshooting information 802
- temporary enabler code
 - expired 836, 837
- temporary notes 458
- Terminal Services Licensing
 - recover
 - prerequisites 396
- time attribute of managed event 458
- time range 426
- Together 248

- toolbars
 - User program 45
 - tracking
 - cloned data 349
 - online index information 586
 - traps
 - categories 697
 - SNMP 694
 - troubleshooting 534
 - aborted recover 815
 - AIX
 - STK-9840 904
 - archive pools 834
 - archive requests
 - naming 834
 - archives 834
 - multiple save sets 834
 - nsrarchive program 834
 - remote request failure 834
 - auto media verification 810
 - autochangers
 - AIX considerations 904
 - attributes 820
 - autodetected scsi errors 821
 - control port access 823
 - destination component 821
 - HP-UX 904
 - HP-UX considerations 903
 - maintenance 820
 - backups 807
 - backups levels 813
 - backups, stopping 810
 - bootstrap printing, failure 817
 - client file index
 - size growth 815
 - clients
 - alias 813
 - Solaris, location 898
 - daemons 807
 - devices
 - maintenance 820
 - nonrewinding 825
 - Solaris, unsupported 898
 - disk label errors 816
 - DNS hostname alias 813
 - ECB counter 811
 - file conversion 817
 - firmware 826
 - hosts table 786
 - HP-UX
 - SCSI pass-through driver 904
 - unloading drives 903
 - unsupported media 903
 - illegal characters 812
 - licensing, copy violation 817
 - nsrexec processes 807
 - online indexes 812
 - packet receive buffer 811
 - pathname restrictions 813
 - recoveries 807
 - online indexes 812
 - remote access 831
 - recovering POSIX hard links 116
 - renamed client backups 814
 - retrievals 834, 835
 - routers 826
 - scanner program 811
 - server errors, binding to 832
 - server index 817
 - Solaris 898
 - storage nodes 835
 - technical support information 802
- ## U
- uasm program 812
 - unresponsive browser 836
 - URL
 - Console software 38
 - user
 - authentication 532
 - deleted 452
 - user groups
 - creating 567
 - user interface
 - overview 35
 - setting preferences 44
 - User List Report 447
 - user privileges 45
 - User Reports 423
- ## V
- vanishing managed event 838
 - variable
 - case sensitivity 534
 - for gstd log 839
 - GST_DEBUG 839
 - GST_MAXLOGVERS 839
 - setting 534
 - verification
 - NetWorker User program
 - browse windows 45
 - of files 76
 - Verify button 45
 - verifying ASR recovery 718
 - view
 - document view 434
 - viewing
 - annotation 459
 - enterprise hierarchy 543
 - reports 428
 - volume pool
 - See also pools
 - archive 326
 - defined 304
 - Volume Shadow Copy Service
 - commands 778
 - controlling from Administration window 776
 - controlling from command-prompt 777
 - overview 772

- writers 775
- volume,erase 184
- volumes
 - archive 353
 - client file index
 - removing 593
 - cloning 350
 - archive data 353
 - creating 350
 - recovery 351
 - labeling 318
 - maximum size 322
 - tips 322
 - modes
 - types 221
 - nonrewinding 825
 - recycling 592
 - relabeling 219
 - removing 591
 - verify 76
- VSS. See Volume Shadow Copy Service

W

- waiting priority 458, 469, 473
- warning priority 458, 469, 473
- web
 - browser
 - unresponsive 836
 - Console
 - server name 38
- Windows disaster recovery
 - 2003, XP 709
 - 2008, 2008 R2, Windows 7 719
- WINS (Windows Internet Naming Service) 585
- WINS (Windows Internet Naming Service) database 115, 397
- Wizard
 - Console Configuration 535
- wizard
 - client backup configuration 59
 - device configuration 168
- writer 774

Z

- zone, control 542

